Mismeasuring Productivity in the Energy Transition

Emilien Ravigné*

October 20, 2026

Abstract

The energy transition can raise measured electricity-sector productivity even without a corresponding improvement in productive efficiency. I identify two biases in measured productivity. First, falling imported wind and solar capital costs should appear as capital deepening, not electricity-sector productivity. Applying the United States new capacity-cost index to historical investment flows, I reconstruct US electric-power total factor productivity (TFP) and reclassify 4.6 log points of 2011–2024 gross-output growth. This correction more than doubles the measured productivity decline of the deployment decade. Second, renewable generation is self-correlated and increasingly concentrated in low-value hours. Using hourly market data, I show that official accounts miss this value decline, which grows with penetration and eventually overtakes the capital bias.

Keywords: productivity measurement, renewable energy, capital deflators, value factor, electricity.
JEL codes: O47, E01, C82, Q42, Q55.

1 Introduction

Between 2010 and 2024, the cost of solar PV fell by 87% and onshore wind by 55% (International Renewable Energy Agency, 2025), declines comparable to transistors and computers (Way et al., 2022; Aizcorbe et al., 2008; Byrne et al., 2018; Nordhaus, 2007; Farmer and Lafond, 2016). Rapid cost declines usually reflect technological progress and one might therefore expect them to appear in electricity-sector productivity. Yet measured Total Factor Productivity (TFP) in US electric power roughly stagnated during the past decade. This apparent paradox is familiar from the case of computers: Solow famously observed that computers were everywhere except in the productivity statistics (Solow, 1987; Triplett, 1999). Renewables present the same surface paradox. However, the analogy is incomplete: renewable electricity introduces a second measurement problem with no close counterpart in the computer case.

In this paper, I show that renewable deployment creates two productivity-measurement problems that emerge at different stages of the energy transition. Both overstate measured productivity, and their effects are additive: one understates real capital input, the other overstates real output. Early in the transition, the measurement problem is familiar from the computer case: price indexes fail to fully capture the falling cost of renewable equipment, so some capital deepening is misclassified as productivity growth. Later, as renewable generation becomes increasingly concentrated in the same hours, its market value per MWh falls in a way that annual output indexes do not

*I gratefully acknowledge financial support from the CMCC Foundation (Euro-Mediterranean Center on Climate Change) through a Marsili Research Fellowship. All errors are my own.

capture. I use a 2023 revision by the US Bureau of Economic Analysis (BEA), which switches to a renewable-specific capital deflator, to measure the first bias: how much published US electricity TFP reflected falling capital prices rather than technical change within electric power. I then use hourly electricity prices and generation to measure the emerging output-value bias. The first bias is front-loaded and fixable by statistical convention; the second is back-loaded and requires more granular output measurement.

A capital deflator determines how much nominal investment is recorded as real capital and therefore how measured growth is divided between capital deepening and productivity: if it misses falling equipment prices, it understates capital growth and overstates TFP. Both TFP growth and capital deepening raise labour productivity and real income in a standard growth-accounting framework. However, the distinction matters for cross-country comparisons. Countries using different capital deflators are not directly comparable in measured productivity, creating a risk that differences in statistical convention are attributed to differences in economic performance (Rath et al., 2019). This can bias cross-country productivity comparisons if countries with greater renewable deployment also differ in how renewable capital is deflated. For information and communication technologies (ICT), the debate over growth attribution came only after computers had diffused widely, requiring the productivity record of the 1980s and 1990s to be revised ex post. The renewable transition is still in the early phase of the diffusion process, giving statistical agencies the opportunity to incorporate improved price measurement as the transition takes place.

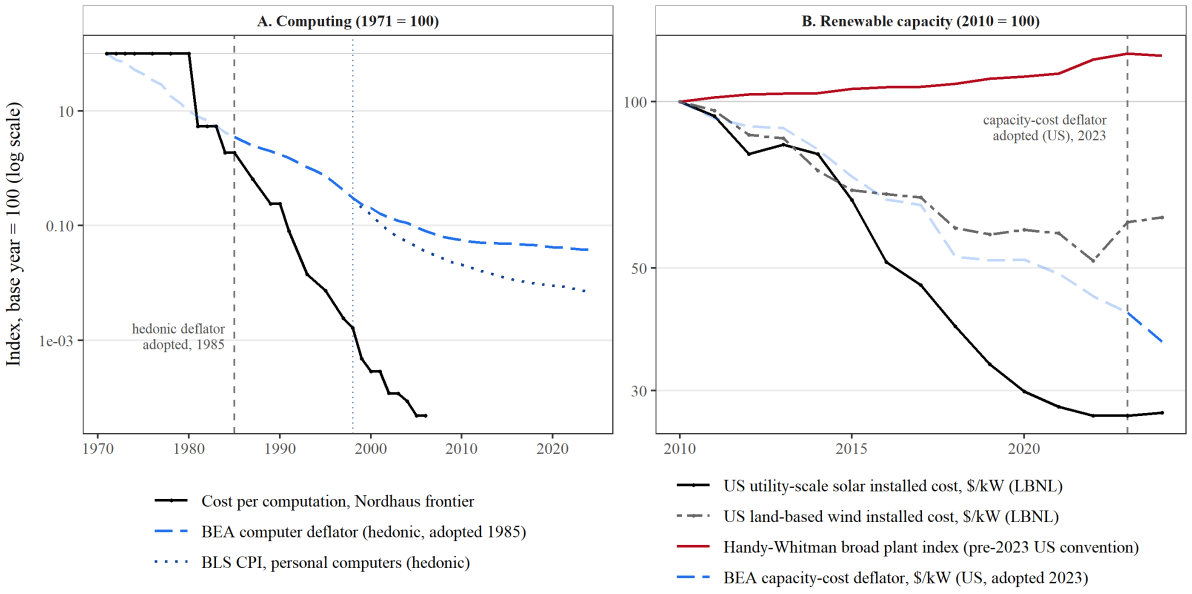


Figure 1: **Two episodes of rapid embodied technological change and statistical adjustment.** All US, real (relative to the GDP price index), indexed to 2010. Panel A compares the cost of computation (Nordhaus, 2007) with official computer price indexes; hedonic adjustment was adopted in 1985, more than a decade into the decline. Panel B compares US utility-scale solar and land-based wind installed costs—capacity-weighted dollars per kilowatt, from the Berkeley Lab utility-scale solar and land-based wind market reports (Seel et al., 2025; Wiser et al., 2025)—with the broad Handy–Whitman electric-plant index (Whitman, Requardt and Associates, LLP, 2026) and the renewable-specific BEA capacity-cost deflator adopted in 2023 and applied back to 2010 (McCulla et al., 2023).

Figure 1 places the renewable transition beside the earlier diffusion of computers, two episodes

in which rapid embodied technological change required technology-specific capital-price measurement. The adjustment came much earlier for renewables: in 2023, the BEA replaced the broad construction indexes used for wind and solar structures with renewable-specific dollars-per-kilowatt deflators, applied back to 2010. The broad indexes had changed little over time, thereby understating capital deepening and leaving too much growth in the electricity-sector productivity residual. Since much of this equipment is imported at world prices, the old convention also misattributes upstream and foreign learning to domestic productivity in the sector that happened to install the equipment, similar to the import-price bias of Houseman et al. (2011). A capacity-cost deflator instead records the gain as cheaper capital for the adopter, as a terms-of-trade gain, and, where appropriate, as manufacturing productivity abroad: green innovation does not disappear, but is booked where it occurs. Cheaper capital still raises the amount of productive capital available to the adopter and thereby raises labour productivity.

Renewables also create a second measurement problem with no close ICT precedent. A faster computer does not become less useful because other firms own faster computers, whereas a solar megawatt-hour becomes less valuable when many solar panels produce simultaneously. Annual output measures can therefore overstate the growth of real electricity output, thus, productivity.

Three results follow. The first is an accounting reclassification identified using the old and new US official price-index conventions applied to the same investment flows; the second is a market-data measurement exercise; the third is a calibrated extrapolation.

First, the capital correction shifts the measured productivity record of the deployment decade. Replacing the broad construction deflator with the renewable-specific capacity-cost index reclassifies 4.6 cumulative log points of 2011–2024 gross-output TFP growth, or 7.1 log points on a value-added basis, from the residual into capital deepening. With renewable capital correctly deflated, a measured gross-output decline of 3.2 log points becomes one of 7.8: the rapid deployment of renewable capacity was primarily capital deepening rather than productivity growth in US electric power.

Second, the output-value bias is already visible in market data and it grows with penetration. Across 138 European market-years, wind and solar generation is valued about 2.9% below the average MWh, with the gap reaching 17.5% in West Denmark in 2024. California and Texas, at roughly half that penetration, show gaps of 11.7% and 5.3%. Official output systems distinguish customer classes, tariffs, or destinations, but not the hour and location of delivery, and therefore miss this decline in the value of electricity output. The full output-side object is a current-market-value benchmark, of which the composition component corresponds to a conventional volume-index correction.

Third, the relative importance of the two biases changes over the transition. The output-value bias grows with renewable penetration and overtakes the fixable capital-price bias at a wind-and-solar generation share of about 34% under the central estimate, and between 29% and 44% across alternative wind–solar mixes. The United States remains well below that point and the EU average just short of it, while Germany has passed it and West Denmark is well beyond it. These crossing points are calibrations rather than forecasts: they hold the historical relationship between renewable penetration and market value fixed, while storage, transmission, demand flexibility, and other forms of adaptation can shift it.

The paper connects three literatures. Relative to the ICT growth-accounting literature, it uses an official change in statistical convention to measure directly the reallocation between capital deepening and sectoral productivity, and shows that renewable electricity adds a second, output-side bias (Jorgenson and Stiroh, 1999; Oliner and Sichel, 2000; Jorgenson et al., 2008; Bils and Klenow, 2001; Syverson, 2017). Relative to the import-price literature, it studies technological progress embodied in a capital good rather than an intermediate input (Feenstra, 1994; Broda and

Weinstein, 2006; Houseman et al., 2011). Relative to the electricity-economics literature, it brings the declining market value of intermittent generation into the measurement of sectoral productivity (Joskow, 2011; Borenstein, 2012; Hirth, 2013; Gowrisankaran et al., 2016; Stiewe et al., 2025). The first measurement bias has already prompted a correction to the US statistical system, though not yet to the other systems examined here. The second bias has not: none of the official systems examined in this paper distinguishes electricity output by the hour and location of delivery.

2 Two biases in the electricity residual

Let Y_E denote electricity-sector output, produced from capital K_E, labour L_E, and intermediates M_E under constant returns and competitive factor pricing, with Hicks-neutral level A_E and nominal cost shares s_K, s_L, s_M summing to one. Denoting log changes by hats, measured TFP growth is the Solow residual:

$$\widehat{A}_E = \widehat{Y}_E - s_K \widehat{K}_E - s_L \widehat{L}_E - s_M \widehat{M}_E. \tag{1}$$

The same algebra applies on a value-added basis with intermediate inputs omitted. I report both measures, taking gross-output TFP as the baseline and the value-added result as the corresponding rescaling.

The input wedge. The investment deflator converts nominal investment into real investment,

$$\widehat{I}_{R,t} = \widehat{I}^{\text{nom}}_{R,t} - \widehat{P}_{K,t}.$$

Let $\mathcal{R}$ denote renewable assets, $\omega_{R,t}$ their share of sector capital services, and

$$\gamma_{K,t} \equiv \widehat{P}^{\text{broad}}_{K,t} - \widehat{P}^{\mathcal{R}}_{K,t}$$

the growth-rate gap between the broad and renewable-specific deflators. A positive $\gamma_{K,t}$ means that the broad deflator records slower growth in real renewable investment.

Proposition 1 (Capital-deflator reallocation)**.** *Suppose nominal renewable investment and both investment deflators grow at constant rates, and the corresponding capital stocks have reached their balanced-growth paths. Holding output, noncapital inputs, and capital-services weights fixed across convention, switching from the broad to the renewable-specific deflator lowers measured productivity growth by*

$$B_K = s_K \, \omega_R \, \gamma_K. \tag{2}$$

The same amount is reallocated from the productivity residual to capital deepening, without changing nominal investment or output.

Under balanced growth, each capital stock grows at the same rate as the real investment used to construct it. The deflator gap therefore translates directly into a capital-growth gap: some capital deepening appears as productivity growth when renewable investment prices are mismeasured.

If renewable prices fall fastest relative to broad construction prices early in deployment, this mechanism receives an early impulse, while the output-value wedge builds with renewable penetration. Surviving investment vintages spread the capital effect over subsequent years, and rising renewable exposure can delay its peak. The relative timing therefore depends on cost declines, capital turnover, and deployment. Section 3 reconstructs both capital histories using perpetual inventory and allows rental weights to respond to the deflator convention.

The output bias. Official electricity-output measures do not distinguish delivered MWh by generating technology, hour, or location. Pooling technologies is appropriate, but pooling time and location conditions creates a bias. At a given time and place, a MWh has the same market value whether produced by wind, gas, or nuclear power. Treating generation technologies as separate products would create composition effects and risk recording falling renewable production costs as additional real output.

Across hours and locations, however, MWh are not equivalent: electricity delivered in a scarce hour or a constrained location is worth more than electricity delivered when production is abundant. Variable renewable output is correlated with itself, so as penetration rises a growing share of renewable generation arrives precisely in the hours of abundant renewable supply, when wholesale prices are low. An output index that weights all MWh equally therefore misses the decline in relative value and overstates real output growth. I measure this bias by comparing generation aggregated with equal MWh weights, Q_t^{phys}, with generation aggregated using observed hour-location prices, Q_t^{val}.

Let $h = (m, n)$ index hour-location cells, where $m = 1, \ldots, M_t$ denotes the hours of year t and $n = 1, \ldots, N$ denotes locations. Let $p_{h,t}$ and $q_{h,t}$ denote the corresponding price and quantity. Define $\bar{p}_t$ as the system-wide average hourly price, weighting hours uniformly and locations by load share.[1] It normalizes the technology-specific capture ratios and cancels from their ratio below.

A value-sensitive output measure distinguishes MWh by their hour and location of delivery, valuing each quantity $q_{h,t}$ at its corresponding price $p_{h,t}$. By contrast, the physical-output benchmark simply counts MWh as

$$Q_t^{\text{phys}} = \sum_h q_{h,t}. \tag{3}$$

This benchmark assigns the same weight to every MWh regardless of when or where it is delivered. It is a conceptual approximation to official constructions that omit the hour-location margin (see Supplemental Appendix C.1).

Let $Q_{j,t} = \sum_h q_{h,t}^j$ denote generation by technology j. Its capture ratio is

$$v_{j,t} = \frac{\sum_h p_{h,t}\, q_{h,t}^j / Q_{j,t}}{\bar{p}_t}$$

the generation-weighted price received by technology j relative to the system-wide average hourly price. Summing generation revenue across technologies gives

$$\sum_h p_{h,t} q_{h,t} = \bar{p}_t \sum_j v_{j,t} Q_{j,t}.$$

Let r denote wind and solar generation and d all other generation, which serves as the comparison group. Deflating total revenue by the average price received by other generation, $\bar{p}_t v_{d,t}$, converts revenue into equivalent MWh of other generation, and gives

$$Q_t^{\text{val}} = Q_{d,t} + v_t Q_{r,t} = Q_t^{\text{phys}}\big[1 - \theta_t(1 - v_t)\big], \qquad v_t \equiv \frac{v_{r,t}}{v_{d,t}}, \qquad \theta_t \equiv \frac{Q_{r,t}}{Q_t^{\text{phys}}}. \tag{4}$$

Here, θ_t is the wind-and-solar share of physical generation, while v_t measures the market value of wind and solar relative to other generation. Both are observable in market data. Because the common price benchmark $\bar{p}_t$ cancels from the ratio v_t, it directly compares the average price received by wind and solar with that received by other generation. Thus, $v_t < 1$ means that a MWh of wind or solar has lower market value, on average, than a MWh of other generation.

[1] For a single zone, $\bar{p}_t$ reduces to the arithmetic mean of hourly prices.

Proposition 2 (Output-value bias)**.** *If the output account assigns the same value to every MWh within a year, the resulting overstatement of generation TFP growth is*

$$B_{V,t}^{\text{gen}} = \widehat{Q}_t^{\text{phys}} - \widehat{Q}_t^{\text{val}} = -\Delta \log\left[1 - \theta_t(1 - v_t)\right] \approx \frac{\overbrace{(1 - v_t)\,\Delta\theta_t}^{\textit{rising renewable share}} \quad - \quad \overbrace{\theta_t\,\Delta v_t}^{\textit{falling relative value}}}{1 - \theta_t(1 - v_t)}. \tag{5}$$

The bias is positive whenever $\theta_t(1 - v_t)$ increases. This occurs when the wind-and-solar share rises while their relative value is below one, when their value relative to other generation falls, or both. To first order, $(1-v_t)\Delta\theta_t$ captures the effect of a rising wind-and-solar share, while $-\theta_t\Delta v_t$ captures a further decline in the average price received by wind and solar relative to other generation.

A natural solution to the output-value bias would be to treat electricity delivered at different hours and locations as separate products. However, it would solve only the composition part: a standard volume index records slower output growth when generation shifts toward lower-value hours or locations, but treats changes in relative prices holding quantities fixed as price changes. Because Q^{val} uses contemporaneous relative prices, it is not a conventional volume index, but captures both effects. I therefore use it as a current-market-value benchmark, asking whether additional generation delivers the same market value as the generation it displaces. For instance, holding quantities fixed a fall in the relative price of midday electricity would leave a volume index unchanged but would lower Q^{val}. Thus, both components of B_V arise on the output side, but only the composition component is unambiguously an output-volume measurement bias; the other captures declining market value.

Proposition 2 applies only to generation, whereas the measured sector also includes transmission and distribution. To express the correction in whole-sector TFP units, I therefore scale the generation bias by w_G, the share of sector output coming from generation rather than from delivery

$$B_{V,t} = w_G B_{V,t}^{\text{gen}}, \qquad w_G \simeq 0.585.$$

I measure w_G as EIA generation revenue divided by BEA gross output for NAICS 2211, matching the output concept used in the sector productivity accounts. The ratio averages 0.585 over 2012–2022. Across years and comparable US and European measures, estimates range from 0.52 to 0.65 (U.S. Energy Information Administration, 2026b; Eurostat, 2026a) (Supplemental Appendix F).

Timing of the two biases. The two biases run on opposite clocks. B_K is front-loaded: it is large while renewable-capital costs are falling rapidly and exposure is rising, and disappears once the capacity-cost deflator is adopted. B_V is back-loaded: it is small at low renewable penetration, then rises as penetration increases and renewable generation becomes increasingly concentrated in low-price hours. Section 5 quantifies the crossover between the two.

3 Capital-price bias and BEA revision reallocation

3.1 Decomposition of the capital-price bias

This subsection quantifies how much measured electricity-sector productivity is reclassified as capital deepening when renewable investment is deflated by capacity cost rather than a broad construction index. All accounts accumulate the same nominal investment history under each price convention by a perpetual-inventory method, and differ only in how capital input is weighted: transparent stock-weighted benchmarks make the arithmetic directly auditable, while the preferred

Hall–Jorgenson account measures capital input using rental-weighted capital services. The preferred account reclassifies 4.6 gross-output and 7.1 value-added log points over 2011–2024, deepening measured declines of 3.2 and 5.2 log points into 7.8 and 12.3. I construct the deflator gap, renewable exposure, and capital share in turn, then show how allowing rental weights to respond produces the preferred estimate.

Price gap γ_K. Before its 2023 comprehensive update, the BEA deflated wind and solar structures using broad Handy–Whitman construction indexes that did not cover alternative-energy structures. The update introduced an "alternative-electric" index based on EIA installed costs per kilowatt and applied it back to 2010. Because installed capacity is the asset being replaced, this technology-specific index better matches the replacement-cost concept used in capital measurement (OECD, 2001).

Between 2010 and 2024, the broad index rose by 52.8 log points while the alternative-electric index fell by 66.5, producing a cumulative gap of $\gamma_K = 119.3$ log points (BEA NIPA Table 5.4.4). The measured decline is not generated by a shift from wind toward cheaper solar. Because the BEA series is a chained price index, a change in the mix of renewable investment does not by itself count as a price decline. The index also does not treat improvements in the services delivered by a kilowatt as additional real investment. To the extent that better turbines or solar installations cost more, it therefore records some quality improvement as a higher price, making the measured price decline smaller rather than larger. Supplemental Appendix A.6 detail both points.

Renewable-capital exposure, ω_R. The relevant exposure is the renewable assets' share of sector capital input. Because s_K covers the whole sector, its denominator must likewise include all NAICS 2211 capital, including equipment, transmission, and distribution. Official capital-services weights are not published at the required asset split. Table 1 therefore uses the preferred Hall–Jorgenson account, in which exposure is the annual two-period rental share of the repriced assets, $\bar{\omega}^{\text{cap}}_{R,t}$. It rises from 0.038 in 2011 to 0.107 in 2024, above the corresponding share of current-cost stock value because fast-depreciating renewable assets carry a higher user cost than their share of the stock. Applied to the accumulated-stock growth gap, it implies a cumulative capital-input wedge of 6.05 log points. Supplemental Appendix A.4.2 reports the stock-weighted alternative.

The renewable category is BEA asset SU60, alternative-electric structures, whose current-cost share of total NAICS 2211 capital rose from 3.6 percent in 2011 to 11.2 percent in 2024. This input exposure is distinct from wind-and-solar generation penetration, an output share of 17.2 percent in 2024. Section 5 uses the observable stock share for the forward exercise; because it produces the larger capital wedge, it yields a later crossover and is conservative for when the output-value wedge overtakes the capital-price wedge.

Capital share, s_K. The deflator choice changes measured real capital but not the nominal current-price factor payments from which $s_{K,t}$ is calculated. In the net-stock benchmark, I fix the capital share under both deflators at its 2011–2024 arithmetic mean,

$$s^*_K = \frac{1}{T}\sum_t \bar{s}_{K,t}.$$

The preferred account instead retains annual variation, applying the standard two-period share $\bar{s}_{K,t} = (s_{K,t-1} + s_{K,t})/2$ under both deflator conventions.

From the direct effect to the preferred capital-services account. Equation (2) holds fixed the asset weights used to construct capital input. The preferred Hall–Jorgenson account instead recomputes them under each deflator convention (Hall and Jorgenson, 1967). Intuitively, these weights are rental weights measure each asset's contribution to the cost of using capital, rather than its share of the capital stock. Because measured asset-price changes affect both real capital growth and the cost of using the asset, the deflator enters twice: it changes the growth assigned to renewable capital and the weight that renewable capital receives in aggregate capital services.

Let $j \in \{\text{capacity}, \text{broad}\}$ denote the convention and $\bar{v}^{j}_{a,t}$ the corresponding two-period average rental weight. The same weights define renewable exposure when summed over the repriced assets and aggregate capital-input growth when applied to all assets:

$$\bar{\omega}^{j}_{R,t} = \sum_{a \in \mathcal{R}} \bar{v}^{j}_{a,t}, \qquad \widehat{K}^{j}_{E,t} = \sum_{a \in \mathcal{A}} \bar{v}^{j}_{a,t} \widehat{K}^{j}_{a,t}.$$

The resulting TFP reclassification is the capital term of the Solow equation (1)

$$B^{\text{services}}_{K,t} = \bar{s}_{K,t} \left(\widehat{K}^{\text{capacity}}_{E,t} - \widehat{K}^{\text{broad}}_{E,t} \right). \tag{6}$$

Using the capacity-cost weights as the reference gives

$$B^{\text{services}}_{K,t} = \bar{s}_{K,t} \left[\bar{\omega}^{\text{capacity}}_{R,t} \gamma_{K,t} + \sum_{a \in \mathcal{A}} \left(\bar{v}^{\text{capacity}}_{a,t} - \bar{v}^{\text{broad}}_{a,t} \right) \widehat{K}^{\text{broad}}_{a,t} \right]. \tag{7}$$

The first term is the direct fixed-weight effect; the second is the response of the rental weights to the deflator convention. If those weights are invariant, the second term is zero and Equation (7) reduces to Equation (2).

Capital-deflator correction. In the preferred account, a direct renewable-stock growth effect of 3.13 log points is reinforced by a +1.48 rental-weight response, yielding the headline reclassification of 4.60 on a gross-output basis and 7.09 on a value-added basis. Falling costs embodied in renewable capital therefore belong in capital deepening rather than in the electricity-sector productivity residual.

The size of the correction, though not its sign, depends on how capital input is weighted: weighting assets by current-cost stock value instead of rental price reclassifies 1.00 to 1.18 gross-output log points rather than 4.60. Rental weighting is the appropriate concept for capital services, so the preferred estimate is the one carried forward; Supplemental Appendix A.4.2 sets out why the two weighting schemes diverge and what each implies.

3.2 Why we cannot directly compare official TFP vintages

A simple comparison of the published TFP of utilities before and after the 2023 comprehensive update cannot identify this effect directly. Indeed, over 2011–2021, the update raised measured capital input by 1.9 log points, as the new deflator would imply, but raised Utilities TFP by 5.9 log points because intermediate-input revisions dominated (Supplemental Appendix A.1). Treating the comprehensive update as a single-change experiment would therefore give the wrong sign.

Table 1 instead isolates the effect of the deflator convention directly. It applies the old and new deflators to the same nominal investment flow and holds output and all other inputs fixed. The magnitude of B_K is not driven by changes in nominal investment or by how assets are classified: the nominal investment series is essentially unchanged across vintages, and reallocating 10% of renewable investment to or from conventional structures changes the preferred estimate by about 0.5 log points (Supplemental Appendix A.5).

Table 1: **Capital-deflator revision and U.S. electric-power productivity.** Both panels report the preferred capital-services account, which accumulates stocks under each price convention. Panel A builds the correction up from the deflator gap; Panel B reports it as a productivity revision. In both, direct renewable-stock growth effect + rental-weight response = total B_K = broad-deflator productivity − capacity-cost productivity.

Panel A. From the deflator gap to the capital-input bias, 2011–2024

	Component	Capital services (annual, preferred)
	Broad electric-infrastructure deflator growth, $\Delta \ln P_K^{\text{broad}}$	+52.8
−	Capacity-cost deflator growth, $\Delta \ln P_K^{\mathcal{R}}$	−66.5
=	Deflator gap, $\gamma_K = \sum_t \gamma_{K,t} \equiv \Delta \ln P_K^{\text{broad}} - \Delta \ln P_K^{\mathcal{R}}$	**+119.3**
⇝	Accumulated renewable-stock gap, $\kappa = \sum_t \kappa_t$ damping $\kappa/\gamma_K = 0.589$	**+70.3**
×	Renewable-capital rental exposure, $\bar{\omega}_{R,t}$ 2011 → 2024	0.038 → 0.107
=	Cumulative capital-input bias, $\sum_t \bar{\omega}_{R,t}\,\kappa_t$ (log points)	6.05
×	Sector capital share, gross output	0.449 → 0.569
=	Direct effect, $\sum_t \bar{s}_{K,t}\,\bar{\omega}_{R,t}\,\kappa_t$ (log points)	3.13
+	Response of the rental weights to the deflator convention (log points)	+1.48
=	Total B_K, gross output (log points)	**4.60**

Panel B. Capital-deflator reclassification of measured TFP, 2011–2024

			Deflator reclassification, B_K			
Productivity basis	Sector capital share	Broad-deflator TFP	Direct effect	+ Weight response	= Total B_K	Capacity-cost TFP
Gross output	0.449 → 0.569	−3.15	**3.13**	+1.48	**4.60**	**−7.75**
Value added	0.771 → 0.788	−5.17	**4.75**	+2.34	**7.09**	**−12.27**

Notes: Deflator changes and productivity effects are in log points; exposures and factor shares are fractions. Cumulative values sum the changes 2011-2024. Calculations use unrounded values, so displayed components may differ slightly because of rounding. Arrows report 2011 and 2024 values; parentheses report the sample range. Note that 2024 values are provisional and might be revised. Panel A reports gross-output effects and builds up the gross-output row of Panel B, with which it reconciles exactly. Its central step is the passage from prices to stocks. The two conventions price each year's investment $\gamma_K = 119.3$ log points apart, but the *accumulated* renewable stocks grow apart by only $\kappa = 70.3$ log points, a damping of 0.589, because only surviving vintages carry the price difference forward. Substituting the contemporaneous gap γ_K for κ, as a one-period repricing calculation does, overstates the correction by a factor of roughly 1.7.
Sources are BEA NIPA Table 5.4.4, BEA Detailed Fixed Assets and industry accounts, BLS, and EU KLEMS. Supplemental Appendices A.1, A.3, and A.4 provide construction details, appendix A.4.2 compares the capital-input weighting schemes, and appendix A.4.3 discusses the 2024 endpoint uncertainty.

3.3 Who gets the learning curve?

In the framework of Section 2, sourcing creates no third measurement bias; it determines the price path entering γ_K in Proposition 1. Imported equipment follows the global learning curve, whereas

production for a smaller domestic market may follow a shallower local curve. With the broad construction deflator, this difference passes into B_K through the accumulated renewable-stock gap κ that Table 1 builds up, and changes how much upstream learning is recorded as domestic electricity TFP. With the capacity-cost deflator, it changes measured capital deepening rather than electricity-sector TFP. Proposition 2 is unaffected because the output-value bias depends on when and where electricity is delivered, not where its equipment was produced.

I develop an illustrative model to trace how sourcing reallocates learning gains between electricity and manufacturing. To a first approximation, installed renewable-capital prices follow an experience curve, falling by $b\,\Delta \ln Q_{\text{cum}}$ as cumulative deployment rises, where b is the learning elasticity. For example, a 20% cost decline with each doubling of cumulative deployment implies $b = 0.322$. The calibration below allows this elasticity to differ between solar and wind.

For technology $j \in \{s, w\}$, let m_j denote the growth of global cumulative capacity, $s_{d,j}$ country d's share of global additions, b_j the elasticity of installed cost with respect to cumulative experience, and w_j its weight on technology j in the BEA wind-and-solar capacity-cost deflator. Starting both sourcing regimes from the same inherited experience stock, the reduction in the broad-deflator residual under full reshoring is

$$B_K^{\text{off}} - B_{K,d}^{\text{re}} = s_K \omega_R \sum_j w_j b_j \ln\left(\frac{m_j}{1 + (m_j - 1)s_{d,j}}\right). \tag{8}$$

Equation (8) is a no-spillover counterfactual: with complete international knowledge spillovers, domestic producers continue to track the global frontier and the difference vanishes.

The technology-specific curves reproduce the measured 66.5-log-point decline in the renewable capacity-cost index under global experience. Restricting subsequent learning to domestic experience reproduces only 27.9 log points in the United States and 34.6 in the EU. Full reshoring therefore reduces the broad-deflator electricity residual by 1.4–2.3 value-added log points in the United States and 1.2–1.9 in the EU. Reshoring also creates additional domestic manufacturing learning. Even under an upper-bound calculation that attributes the entire resulting cost improvement to manufacturing TFP, however, the gain is only 0.36 and 0.44 electricity-sector-equivalent log points, respectively. These sourcing magnitudes scale the correction in Table 1 in proportion to the change in the cumulative price gap, a linear approximation to the vintage account, in which a change in the price path reaches measured capital input only through the surviving vintages. Supplemental Appendix B provides the experience-stock, domestic-content, and manufacturing calculations.

The broad-deflator residual is therefore largest under global sourcing: electricity appears most productive when its capital embodies the deepest learning curve, even when that learning occurs upstream or abroad. This is the capital-input analogue of the offshoring bias in Houseman et al. (2011): unmeasured declines in imported intermediate-input prices appear there as domestic manufacturing productivity, while unmeasured declines in renewable-capital prices appear here as electricity TFP. Reshoring relocates some learning to domestic manufacturing, but the estimated manufacturing gain is substantially smaller than the resulting reduction in the electricity-sector residual.

4 The output-value bias

The previous section showed how the 2023 BEA deflator revision resolves the input-side measurement problem for renewable investment, but it does not resolve the second bias defined in Section 2. I first establish that official output indexes omit the relevant hour-location margin, then measure it from European and US market data and state the benchmark under which it is additive to B_K.

4.1 Why official output does not resolve the value bias

Official systems construct real electricity output in three main ways: direct quantity indexes, deflation of current values, or quantity extrapolation. The methods differ, but the documented US and European implementations share one feature: they aggregate electricity by customer class, tariff, product, or destination, not by the hour and location of delivery. This is explicit in direct-quantity systems such as Norway's, where physical electricity volumes enter the output measure without differentiation by delivery time. As a result, they cannot distinguish a MWh delivered when electricity is scarce from one delivered when renewable supply is abundant.

The implication depends on how real output is constructed. Direct-quantity and extrapolation methods assign equal weight to each MWh and therefore miss shifts toward lower-value hours. In deflated-value systems, part of the decline may already enter measured real output if lower capture prices reduce recorded revenues more than the output deflator. In US electric power, regulated tariffs, long-term contracts, and internal transfers limit the direct link between wholesale capture prices and recorded revenues, making full absorption of the bias unlikely. A calibration of this pass-through by settlement channel, with revenue weights from EIA Form 861 and purchase classes from FERC Form 1, gives 0.39, and 0.49 in a high-pass-through scenario, so about half or more of the wedge would remain even under deflated value. The channel elasticities are assumed rather than estimated, so this is a sensitivity analysis, not a bound. Supplemental Appendix D.2 derives the incidence condition and Supplemental Appendix D.7 reports the calibration.

This missing distinction is not foreign to national-accounting principles. ESA 2010 explicitly recognizes time of delivery as a potential quality characteristic of electricity. The same principle already appears in electricity tariffs: interruptible supply is priced below firm supply because the buyer gives up the right to consume at a moment of its choosing. Finer differentiation by delivery time changes the resolution, not the principle.[2] The problem is therefore not the underlying concept but the level of detail at which it is implemented. By contrast, the accounts are right not to distinguish electricity by generation technology: conditional on its time and place of delivery, a MWh is the same product whether generated by wind, solar, nuclear, or gas. Using source-specific production costs as output deflators would instead create artificial real-output growth.

The US aggregates illustrate why the omitted margin cannot be reconstructed after the fact: Between 2017 and 2024, nominal NAICS 2211 output and its deflator rose by 25.3 and 25.5%, leaving official real output essentially flat, while the Federal Reserve's physical-output index for the same NAICS cut rose by 6.3% (Board of Governors of the Federal Reserve System, 2026). The official measure thus failed to register the increase in physical electricity supplied, and neither series resolves the time-of-delivery margin (Supplemental Appendix D.6). This creates the same identification problem encountered on the capital side, with a different solution. B_K cannot be inferred from changes across official TFP vintages because other inputs were revised at the same time. Similarly, B_V cannot be recovered by manipulating annual output aggregates because the relevant hourly and locational information has already been averaged out within the published strata. It must instead be constructed directly from hourly market prices and generation. Section 4.2 provides the relevant test: holding the underlying data fixed, the bias appears when output is stratified by delivery time or system state.

4.2 Measuring the output-value bias

I implement the value-weighted output index in Equation (4) by pricing each MWh at the day-ahead price for its bidding zone and hour. This price measures the marginal market value of electricity at

[2] See Appendix C.2 develops the full conceptual case and Appendix E for the data approach.

the time and location of delivery. It therefore differs from the revenue received by the generator, which may also include support payments or contract terms such as feed-in tariffs, contracts for difference, and power purchase agreements.

I estimate the output-value bias using hourly market data from a European panel (ENTSO-E) containing 138 bidding-zone-years from 15 zones over 2015–2024. I apply the same construction to CAISO over 2020–2025 and ERCOT over 2018–2025. For each technology $j \in s, w$, penetration θ_j is generation divided by load and the value factor v_j is its generation-weighted capture price relative to the time-weighted baseload price (Hirth, 2013). A value factor above one means that the technology tends to generate in relatively high-price hours. Supplemental Appendix E provides the full construction. The output-value gap is

$$D = \theta_s(1 - v_s) + \theta_w(1 - v_w). \tag{9}$$

Holding penetration fixed, a lower value factor increases D; more generally, the gap is largest when a technology both supplies a large share of load and earns a low price relative to baseload. Equivalently, $D = 1 - Q^{\text{val}}/Q^{\text{phys}}$. Thus D is the proportional shortfall of value-weighted output relative to physical output: it measures how much lower output is when each MWh is valued at its market price rather than counted equally.

Figure 2 shows the level of the output-value gap. It averages 2.9% across European market-years, rising from 0.6% among low-penetration, solar-dominated observations to 3.4% among high-penetration, wind-dominated observations, and reaches 17.5% in West Denmark in 2024. Multiplying by $w_G = 0.585$, the sample mean corresponds to 1.6% of whole-sector output and the West Denmark observation to 9.6%. CAISO and ERCOT lie on the same penetration–gap relationship: at 38 and 35% wind-and-solar penetration in 2024, their gaps were 11.7 and 5.3%. Alternative treatments of negative prices and curtailment leave the European results essentially unchanged (Supplemental Appendix E.5). These are differences in output levels, not annual TFP biases.

A chained Törnqvist index applied to the same hourly data isolates the quantity-composition channel. The result depends on how delivery products are defined: month-by-hour cells produce a 6.2-log-point bias in solar-dominated CAISO but little in wind-dominated ERCOT, whereas residual-load strata produce 6.9 and 7.2 log points, respectively. Across Europe, the residual-load estimate is positive in 13 of 15 zones, with a median of 0.76 log points. The exact decomposition of (5) attributes 34% of the cumulative European bias and 41% of the US bias to quantity composition, with relative-value decline accounting for the remainder. Supplemental Appendix E.6 reports the full results and price-shock sensitivities.

This contrast matches the mechanism: solar output is closely tied to clock time, whereas wind varies substantially across days within the same clock-time cell. The bias is therefore not weaker in wind systems; it requires strata that capture the relevant system state (Supplemental Appendix D).

The Figure 2 is descriptive rather than causal. Because penetration enters D directly, its upward slope combines the greater weight placed on renewable generation with any accompanying decline in its value factor. It therefore cannot identify the forward value-factor schedule used in Section 5.

The level gap D maps into annual generation-TFP overstatement through (5), $B^{\text{gen}}_{V,t} = -\Delta \log(1 - D_t)$, and into sector units as $B_{V,t} = w_G B^{\text{gen}}_{V,t}$. Thus, when the output-value gap D widens over time, value-weighted generation grows more slowly than physical generation, raising measured generation-output growth. Holding measured inputs fixed, it therefore overstates generation TFP growth by $B^{\text{gen}}_{V,t}$. In whole-sector units $B_{V,t}$ averages 0.29 percentage points per year across the European panel—−0.03 before 2020 against +0.53 after—and removing common annual price movements lowers it to 0.06, erasing part of the cannibalisation and making that a lower bound. In the two US markets where the construction is possible it averages 0.90 percentage points per year in CAISO

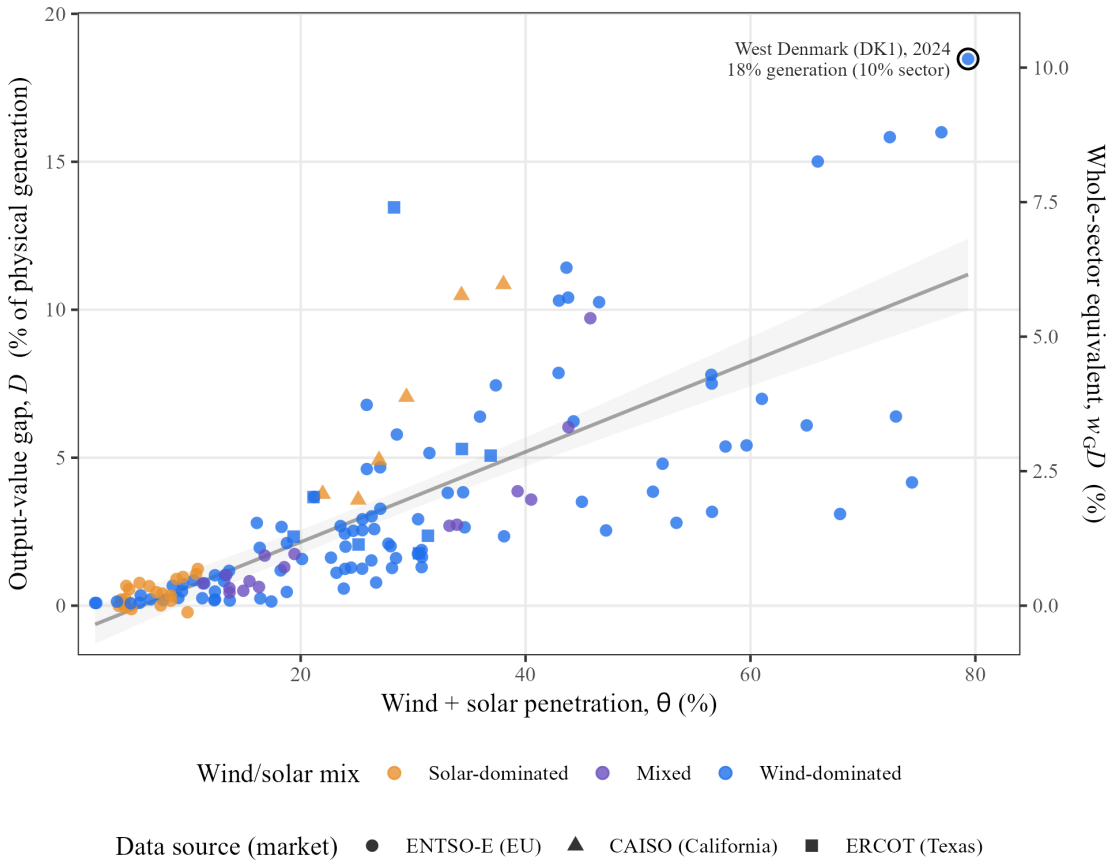


Figure 2: **The output-value gap rises with wind and solar penetration.** Each point represents one market-year. The horizontal axis reports total wind and solar penetration, $\theta_s + \theta_w$. The left vertical axis reports the combined output-value gap, $D = \theta_s(1 - v_s) + \theta_w(1 - v_w)$: the proportional shortfall of market-value-weighted generation relative to physical generation. The right axis expresses the same gap in whole-sector units, $w_G D$, using $w_G = 0.585$. Colour denotes whether renewable generation is solar-dominated, mixed, or wind-dominated, and marker shape identifies the market. The solid line is an OLS fit through all observations, with a pointwise 95% confidence band. The sample contains 152 market-years: 138 from 15 European bidding zones over 2015–2024 (ENTSO-E, 2026), six from CAISO (2020–2025), and eight from ERCOT (2018–2025). The data are constructed from hourly day-ahead prices, generation, and load, with negative prices retained. West Denmark in 2024 has an output-value gap of 17.5%, equivalent to 9.6% in whole-sector units. The elevated ERCOT observation in 2021 reflects Winter Storm Uri.

and 0.23 in ERCOT (Supplemental Appendix E.4). Section 5 compares these magnitudes with the capital bias at matched penetration.

4.3 Additivity and non-overlap of the biases

The additivity of the effects of the two biases on TFP,

$$\text{TFP bias}_t = B_{K,t} + B_{V,t} \tag{10}$$

holds under a specific benchmark: the input deflator prices what a kilowatt costs to build, while the output measure values what the delivered electricity is worth. The BEA capacity-cost index

satisfies the first half of this benchmark. It is a replacement-cost measure per installed kilowatt, not an asset-price index capitalizing the expected capture prices of future output. Lower expected capture prices may affect how much capacity is built and the rents that it earns, but they do not constitute a decline in construction costs. A dollars-per-expected-lifetime-MWh deflator would import the value factor into capital services and mechanically offset the output correction. Such an index could be used under a different accounting benchmark, but it could not then be added to B_V.

The output-side restriction is parallel: electricity is differentiated by its hour and location of delivery, not by its generating technology. Supplemental Appendix C.3 shows that assigning source-specific cost deflators to clean and conventional generation creates two artificial output terms even when the output-value bias is set to zero: a composition effect when the generation mix shifts across technologies with different imputed costs, and a learning effect when clean output is deflated by its falling production cost. At the current US clean share, the latter would add roughly 2 percentage points per year to measured real output growth. Neither term represents additional electricity delivered or a change in its market value.

These restrictions keep the two biases separate. B_K measures how the capital quantity associated with a given nominal investment flow changes when renewable capacity is deflated by its replacement cost. B_V measures how the value of delivered generation changes with its timing and location. The first corrects the measurement of inputs; the second corrects the valuation of output. Neither correction contains the other.

5 The crossover between the two biases

5.1 Current systems' biases

The two biases evolve differently over the transition, so Figure 3 compares them over a common path of wind-and-solar penetration on a common whole-sector gross-output TFP basis.

To extend the two biases beyond the observed data, I express both as functions of wind-and-solar penetration, θ. For the capital bias, I set current-cost renewable-capital exposure equal to penetration, $\omega_R(\theta) = \theta$, beyond the range observed in the United States, and scale it by the sector's 2024 gross-output capital share, $s_K = 0.57$, so that both biases are measured against gross output. For the output-value bias, I let technology-specific value factors decline linearly with penetration, $v_j = v_{0,j} - a_j\theta_j$, and use Hirth (2013) as the central calibration, (Stiewe et al., 2025) for robustness, and alternative wind–solar schedules to form the band. On this common whole-sector basis, the two biases cross at $\theta^\star \simeq 0.34$.

The biases have different bases and trajectories. B_K scales with renewable capital exposure, which is bounded because sector capital also includes transmission, distribution, and other unaffected assets, and with the price gap γ_K, which narrows as technologies mature. B_V instead scales with renewable generation, which can keep rising toward most of output, and with the value discount $1 - v$, which deepens with penetration. The capital bias is therefore front-loaded, while the output-value bias grows in importance later in the transition.

The 2024 positions (black diamonds) on Figure 3 illustrate the change in the dominant bias. The United States ($\theta = 0.17$) remains in the region where the capital bias dominates and correcting the capital deflator removes most of the measured bias, and the EU average (0.29) is just short of the central crossover. Germany (0.44) has passed it, and sits at the crossover of the wind-heavy schedule closest to its own generation mix; West Denmark (about 0.75) is well beyond it, in the region where the larger measurement problem is the valuation of output. CAISO at 0.38 and ERCOT at 0.35 sit at or past the central crossover, so leading regional markets reach it well before

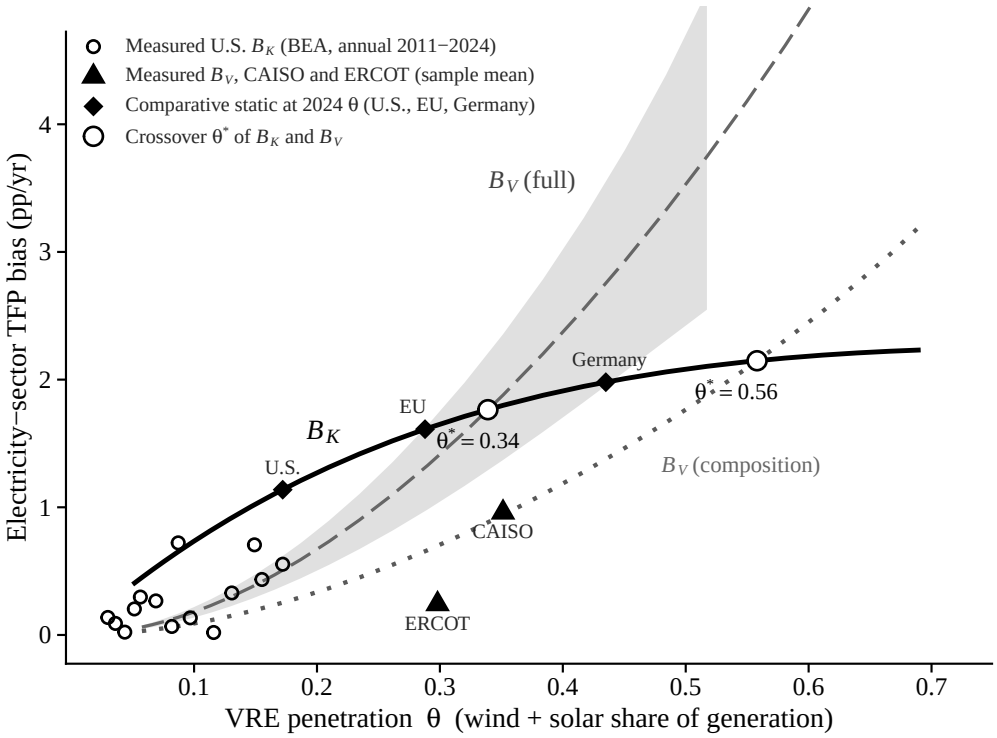


Figure 3: **The dominant measurement bias shifts with renewable penetration.** The solid curve is the capital-price bias $B_K(\theta)$, calibrated to the IRENA installed-cost decline and the measured US deflator gap; open circles use the observed current-cost renewable share of NAICS 2211 capital, $\widetilde{\omega}_R$. The dashed curve is the full output-value bias $B_V^{\text{full}}(\theta) = w_G B_V^{\text{gen}}(\theta)$, with the shaded band spanning alternative wind–solar mixes. The dotted curve retains only the quantity-composition channel of (5); it is a lower-bound specification, not the lower edge of the band. All biases are in whole-sector gross-output units, and filled diamonds mark 2024 system positions at the observed wind-and-solar share of generation (Ember, 2025). Filled triangles mark the average measured output-value bias in CAISO and ERCOT, the two US markets (annual values are volatile, ranging from −0.29 to +2.85 in CAISO and −7.44 to +7.26 in ERCOT). The central full bias crosses B_K at $\theta^\star \simeq 0.34$, compared with $\theta^\star \simeq 0.56$ under the composition-only specification. The large difference in crossover penetration reflects the flatness of B_K at high penetration. The crossover is a penetration threshold, not a date. Supplemental Appendix F provides the construction and sensitivities.

the US national average. Their measured output-value biases nonetheless lie below the central schedule, and still below the capital bias at their penetration. ERCOT is a relatively isolated market, so the neighbouring-penetration component of the cross-border slope does not apply to it. CAISO's shortfall is more plausibly the growth of its battery fleet, which is one of the margins considered next.

Two assumptions affect the reported crossover in opposite directions. Setting $\widetilde{\omega}_R = \theta$ raises the forward capital wedge relative to the observed 2024 US relationship: the repriced assets account for 11.2% of sector capital, compared with wind-and-solar penetration of 17.2%. If this relationship persists, equating exposure with penetration delays the crossover. This assumption is conservative on the exposure margin and does not affect the historical US estimates, which use observed exposures. Conversely, omitting future adaptation through storage and other flexibility may over-

state value-factor declines and bring the crossover forward.[3] Capital vintages introduce a further timing qualification: past deflator divergence can sustain the capital wedge after the annual price gap narrows, but its effect relative to the contemporaneous-gap benchmark depends on investment and stock turnover. Its effect on the crossover therefore requires a forward perpetual-inventory calculation. These effects are not netted.

The 34% crossover is a calibration rather than a forecast: alternative wind–solar value-factor schedules move the full-bias threshold from roughly 29% to 44%, and the gross-output capital share that scales the capital bias moves it from 28% to 40% over its plausible range. Supplemental Appendix F and Table F.1 report the construction and sensitivities.

5.2 Implications for measurement.

Figure 3 summarizes the central result: the two biases evolve in opposite directions over the transition, while their effects on measured TFP are additive (Equation 10). Hence, the capital correction remains necessary but becomes insufficient as penetration rises and the output-value bias overtakes it.

Accounting differences complicate current cross-country comparisons. A system that still uses broad construction deflators records less capital deepening and more electricity TFP than the post-2023 United States, even under identical technology. For every ten percentage points of renewable-capital exposure, the measured US deflator gap raises value-added TFP by about 0.6 percentage points per year using its smoothed 2010–2024 rate, and about 0.8 points using the larger recent rate. A US-anchored harmonization attributes 17–34 cumulative log points of the published EU–US Utilities gap over 2011–2024 and 3.3–8.5 log points of the Germany–US Section-D gap over 2010–2017 to the deflator convention. These ranges make some published rankings fragile, but do not establish a definitive reversal (Supplemental Appendix G). They apply the contemporaneous deflator gap to measured exposure rather than the damped accumulated-stock gap of Table 1, and are in that respect generous estimates of what harmonising the capital account would reallocate. A closely related harmonisation strategy was used in the ICT literature, where differences in national computer deflators led researchers to impose common, often US-based, price indexes for international productivity comparisons (Wyckoff, 1995; Schreyer, 2000).

The statistical priorities follow directly. First, renewable capital should be deflated with capacity-specific cost indexes, following the post-2023 US treatment. Second, delivered electricity should not be divided into separate output products according to generation technology, because the market does not value otherwise identical MWh differently by source. Third, statistical agencies should begin constructing value-weighted electricity output that distinguishes delivery by time, location, and system state.

6 Conclusion

The energy transition reproduces the ICT measurement bias on the capital side but creates a distinct bias on the output side. The BEA's 2023 deflator revision shows that broad construction indexes overstated US electricity productivity growth by 2.8 log points over 2011–2024, or 4.3 log points on a value-added basis. It recorded falling renewable investment costs as productivity rather than capital deepening. Correct capital deflation removes this misclassification.

[3] Storage could also generate a capital-price wedge if its investment deflator fails to capture battery-cost declines (Way et al., 2022); this additional channel is not quantified here.

It does not, however, address the declining value of renewable generation as solar and wind deploy. Hourly evidence from European and US wholesale markets shows that wind and solar generation is increasingly concentrated in lower-value hours. As the energy transition unfolds and renewable energy becomes a larger share of electricity generation, the output-value bias becomes increasingly important and eventually overtakes the capital-price bias.

Productivity measurement in the energy transition therefore requires two distinct corrections. The United States has adopted the first, but others have not[4], creating a discrepancy on green productivity. None of the official statistical systems examined distinguishes electricity output by the hour and location of delivery, even though ESA 2010 recognizes time of delivery as a potential quality characteristic.

[4] As of August 2026.

References

Aizcorbe, A., Oliner, S. D., and Sichel, D. E. (2008). Shifting trends in semiconductor prices and the pace of technological progress. *Business Economics*, 43(3):23–39.

Bils, M. and Klenow, P. J. (2001). Quantifying quality growth. *American Economic Review*, 91(4):1006–1030.

Board of Governors of the Federal Reserve System (2026). Industrial production and capacity utilization — G.17. Statistical release, https://www.federalreserve.gov/releases/g17/. Series `IPG2211A`, physical industrial-production index for electric power (NAICS 2211), annual, 2017=100; generation and technology splits `IPG22111A`, `IPN221114T8A` and `IPN221112A`. Retrieved via FRED. Accessed 17 July 2026.

Bontadini, F., Corrado, C., Haskel, J., and Iommi, M. (2023). Euklems and intanprod: industry productivity accounts with intangibles. Technical report, University Library of Munich, Germany.

Borenstein, S. (2012). The private and public economics of renewable electricity generation. *Journal of Economic Perspectives*, 26(1):67–92.

Bougroug, A. (2021). Quarterly national accounts: Methods and sources of the quarterly and monthly national accounts compilations for norway. Documents 2021/27, Statistics Norway.

Broda, C. and Weinstein, D. E. (2006). Globalization and the gains from variety. *Quarterly Journal of Economics*, 121(2):541–585.

Bureau of Economic Analysis (2024a). *Concepts and Methods of the U.S. National Income and Product Accounts*. U.S. Bureau of Economic Analysis. All-chapters edition, December 2024.

Bureau of Economic Analysis (2024b). *Concepts and Methods of the U.S. National Income and Product Accounts: Chapter 6, Private Fixed Investment*. U.S. Bureau of Economic Analysis.

Bureau of Economic Analysis (2024c). Updated summary of NIPA methodologies. *Survey of Current Business*.

Bureau of Economic Analysis (2025). Updated summary of NIPA methodologies. *Survey of Current Business*.

Bureau of Economic Analysis (2026). Table 5.4.4U. Price Indexes for Private Fixed Investment in Structures by Type. NIPA Underlying Detail, series IA001174-A, Alternative electric, via DBnomics.

Bureau of Economic Analysis (n.d.a). BEA depreciation estimates. U.S. Department of Commerce, Bureau of Economic Analysis, methodology note. Undated document; accessed 4 August 2026.

Bureau of Economic Analysis (n.d.b). Guide to the interactive GDP-by-industry accounts tables. U.S. Bureau of Economic Analysis. Living web guide; accessed 4 August 2026.

Bureau of Economic Analysis and Bureau of Labor Statistics (2026). Integrated industry-level production account (KLEMS). U.S. Bureau of Economic Analysis. Data tables cover 1997–2024; accessed 4 August 2026.

Bureau of Labor Statistics (2005). New producer price indexes for electric power generation, NAICS 221110, and electric bulk power transmission and control, NAICS 221121. BLS Producer Price Index fact sheet. The indexes were introduced with the release of January 2004 data.

Bureau of Labor Statistics (2017). Industry productivity measures: Overview. BLS Handbook of Methods. Archived chapter; superseded by the Office of Productivity and Technology Handbook of Methods.

Bureau of Labor Statistics (2026a). Measuring price change in the CPI: Household energy. BLS Consumer Price Index fact sheet.

Bureau of Labor Statistics (2026b). Producer price indexes: Overview. BLS Handbook of Methods.

Bureau of Labor Statistics (n.d.). Data sources and methods for measuring multifactor productivity in the utility services industries (SIC 49). BLS Office of Productivity and Technology, technical note. Undated technical note; accessed 4 August 2026.

Byrne, D. M., Oliner, S. D., and Sichel, D. E. (2018). How fast are semiconductor prices falling? *Review of Income and Wealth*, 64(3):679–702.

California Independent System Operator (2026a). Open access same-time information system (OASIS). CAISO. Accessed 23 July 2026.

California Independent System Operator (2026b). Production and curtailments data. CAISO data library. The comprehensive production report was discontinued from 1 June 2025; historical files remain available; accessed 23 July 2026.

California Independent System Operator (2026c). Today's outlook: Demand. CAISO. Accessed 23 July 2026.

Chatzipanagi, A., Jaeger-Waldau, A., Letout, S., Mountraki, A., Gea Bermudez, J., et al. (2025). Clean energy technology observatory: Photovoltaics in the european union — 2025 status report on technology development, trends, value chains and markets. Technical Report JRC144141, EUR 40529, Publications Office of the European Union, Luxembourg. Manufacturing benchmark in Table 9, p. 59.

Electric Reliability Council of Texas (2026a). Dam settlement point prices. ERCOT Data Product NP4-190-CD, https://www.ercot.com/mp/data-products/data-product-details?id=NP4-190-CD. Accessed July 23, 2026.

Electric Reliability Council of Texas (2026b). Generation: Fuel mix. ERCOT Generation data page. Fuel Mix reports provide actual generation by fuel type for each 15-minute settlement interval; accessed 23 July 2026.

Electric Reliability Council of Texas (2026c). Hourly load data archives. https://www.ercot.com/gridinfo/load/load_hist/. Accessed July 23, 2026.

Electric Reliability Council of Texas (2026d). System-wide prices. https://www.ercot.com/gridmktinfo/dashboards/systemwideprices. Accessed July 23, 2026.

Ember (2025). Yearly electricity data. https://ember-energy.org/data/yearly-electricity-data/. Yearly full release, long format. Variable "Wind and Solar" in the electricity-generation-share block, percent of total generation; used for the EU, Germany and United States penetration path $\theta(t)$, and the capacity block for the sourcing shares. Compiled from EIA, Eurostat, ENTSO-E and national statistical sources. Accessed 22 June 2026.

ENTSO-E (2026). Transparency platform. European Network of Transmission System Operators for Electricity, https://transparency.entsoe.eu/. Hourly day-ahead prices (Art. 12.1.D), actual generation per production type (Art. 16.1.B&C) and actual total load (Art. 6.1.A), retrieved through the Web API for 15 bidding zones over 2015–2024. Accessed 16 July 2026.

European Commission (2020). Commission implementing regulation (EU) 2020/1197 of 30 july 2020 laying down technical specifications and arrangements pursuant to regulation (EU) 2019/2152 of the european parliament and of the council on european business statistics repealing 10 legal acts in the field of business statistics. Official Journal of the European Union, L 271, 18 August 2020, pp. 1–170. See Annex IV for variable 130201 (producer prices).

European Parliament and Council of the European Union (2019). Regulation (EU) 2019/2152 of the european parliament and of the council of 27 november 2019 on european business statistics, repealing 10 legal acts in the field of business statistics. Official Journal of the European Union, L 327, 17 December 2019, pp. 1–35.

Eurostat (2012). Handbook on industrial producer price indices (PPI). Technical Report KS-RA-12-020-EN-N, Publications Office of the European Union, Luxembourg.

Eurostat (2013). European system of accounts — ESA 2010. Technical report, Publications Office of the European Union, Luxembourg. Annex A of Regulation (EU) No 549/2013. Catalogue KS-02-13-269-EN. Chapter 10, "Price and volume measures"; paragraphs 10.16 and 10.32.

Eurostat (2016). Handbook on prices and volume measures in national accounts. Technical report, Publications Office of the European Union, Luxembourg. Catalogue KS-GQ-14-005-EN-N. Section 4.4, "CPA D — Electricity, gas, steam and air conditioning," pp. 79–80.

Eurostat (2021). *European Business Statistics Methodological Manual for Short-Term Business Statistics*. Publications Office of the European Union, Luxembourg, 2021 edition.

Eurostat (2026a). Electricity prices components for household and non-household consumers, annual data. Datasets `nrg_pc_204_c` and `nrg_pc_205_c`, https://ec.europa.eu/eurostat/databrowser/. Components `NRG_SUP` (energy and supply), `NETC` (network costs) and taxes, in EUR per kWh, all consumption bands. EU27 aggregate accessed 3 August 2026.

Eurostat (2026b). Short-term business statistics (STS). Reference Metadata in Euro SDMX Metadata Structure (ESMS). National reference metadata are provided in the annexes; metadata last updated 7 May 2026.

Farmer, J. D. and Lafond, F. (2016). How predictable is technological progress? *Research policy*, 45(3):647–665.

Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *American Economic Review*, 84(1):157–177.

Gowrisankaran, G., Reynolds, S. S., and Samano, M. (2016). Intermittency and the value of renewable energy. *Journal of Political Economy*, 124(4):1187–1234.

Grid Status (2026). Developers: Grid Status API documentation. Grid Status Documentation. Accessed 23 July 2026.

Hall, R. E. and Jorgenson, D. W. (1967). Tax policy and investment behavior. *American Economic Review*, 57(3):391–414.

Hirth, L. (2013). The market value of variable renewables: The effect of solar wind power variability on their relative price. *Energy Economics*, 38:218–236.

Houseman, S., Kurz, C., Lengermann, P., and Mandel, B. (2011). Offshoring bias in U.S. manufacturing. *Journal of Economic Perspectives*, 25(2):111–132.

International Renewable Energy Agency (2025). Renewable power generation costs in 2024. Technical report, International Renewable Energy Agency, Abu Dhabi.

Jakhmola, A., Jewell, J., Vinichenko, V., and Cherp, A. (2026). Probabilistic projections of global wind and solar power growth based on historical national experience. *Nature Energy*, 11(5):743–755. An author correction to the Figure 3 legend was published 20 April 2026, doi:10.1038/s41560-026-02064-z.

Jorgenson, D. W., Ho, M. S., and Stiroh, K. J. (2008). A retrospective look at the U.S. productivity growth resurgence. *Journal of Economic Perspectives*, 22(1):3–24.

Jorgenson, D. W. and Stiroh, K. J. (1999). Information technology and growth. *American Economic Review*, 89(2):109–115.

Joskow, P. L. (2011). Comparing the costs of intermittent and dispatchable electricity generating technologies. *American Economic Review*, 101(3):238–241.

Mataloni, L. S. and McCulla, S. H. (2025). The 2025 annual update of the national economic accounts. *Survey of Current Business*.

McCulla, S. H., Turner, D. L., and Mataloni, L. S. (2023). Preview of the 2023 comprehensive update of the national economic accounts: Changes in methods, definitions, and presentations. *Survey of Current Business*.

Nordhaus, W. D. (2007). Two centuries of productivity growth in computing. *The Journal of Economic History*, 67(1):128–159.

OECD (2001). *Measuring Productivity - OECD Manual: Measurement of Aggregate and Industry-level Productivity Growth*. OECD Publishing, Paris.

OECD (2009). *Measuring Capital - OECD Manual 2009: Second Edition*. OECD Publishing, Paris.

Office for National Statistics (2024). Index of production, UK QMI. Quality and methodology information report.

Office for National Statistics (2025). Multi-factor productivity (MFP) QMI. Quality and methodology information report.

Oliner, S. D. and Sichel, D. E. (2000). The resurgence of growth in the late 1990s: Is information technology the story? *Journal of Economic Perspectives*, 14(4):3–22.

Rath, B. N., Akram, V., Bal, D. P., and Mahalik, M. K. (2019). Do fossil fuel and renewable energy consumption affect total factor productivity growth? evidence from cross-country data with policy insights. *Energy Policy*, 127:186–199.

Schreyer, P. (2000). The contribution of information and communication technology to output growth: A study of the G7 countries. OECD Science, Technology and Industry Working Papers 2000/02, OECD Publishing, Paris.

Seel, J., Mulvaney Kemp, J., Cheyette, A., Gorman, W., Darghouth, N. R., Robson, D., Rand, J., and Jeong, S. (2025). U.S. utility-scale solar, 2025 data update. Technical report, Lawrence Berkeley National Laboratory, Berkeley, CA. Data file accompanying the 2025 edition, tab "CapEx Trend (PV-only)": capacity-weighted average installed cost of US ground-mounted photovoltaic plants above 5 MW-AC, dollars per kilowatt, AC and DC bases. Plant-level data also posted at https://data.openei.org/submissions/8541. Accessed 4 August 2026.

Solow, R. M. (1987). We'd better watch out. *The New York Times Book Review*, page 36.

Statistics Finland (2026). Producer prices in industry. National Reference Metadata in Euro SDMX Metadata Structure (ESMS), hosted by Eurostat. Metadata last certified, posted, and updated 11 June 2026.

Statistics Norway (2026a). Electricity prices. Official statistics and "About the statistics" documentation. Statistics updated 18 May 2026; "About the statistics" last updated 2 February 2026; accessed 4 August 2026.

Statistics Norway (2026b). External trade in goods. Official statistics and "About the statistics" documentation. Methods documentation states that electricity valuation changed from hourly to 15-minute intervals from October 2025, creating a discontinuity; "About the statistics" last updated 4 June 2026.

Stiewe, C., Xu, A. L., Eicke, A., and Hirth, L. (2025). Cross-border cannibalization: Spillover effects of wind and solar energy on interconnected European electricity markets. *Energy Economics*, 143:108251. A corrigendum was published in Energy Economics 155 (2026), article 109167, doi:10.1016/j.eneco.2026.109167.

Syverson, C. (2017). Challenges to mismeasurement explanations for the U.S. productivity slowdown. *Journal of Economic Perspectives*, 31(2):165–186.

Tapoglou, E., Mc Govern, L., Georgakaki, A., Mountraki, A., Letout, S., et al. (2025). Clean energy technology observatory: Wind energy in the european union — 2025 status report on technology development, trends, value chains and markets. Technical Report JRC144351, EUR 40559, Publications Office of the European Union, Luxembourg. Component coverage in Table 5, p. 50; manufacturing benchmark in Table 6, p. 51.

Triplett, J. E. (1999). The Solow productivity paradox: What do computers do to productivity? *Canadian Journal of Economics*, 32(2):309–334.

U.S. Bureau of Reclamation (2026). Construction cost trends. Technical Service Center, U.S. Bureau of Reclamation. Construction-cost index series originally developed for 35 construction categories. Current cost models use labor, equipment, fuel, and material inputs drawn primarily from BLS Producer Price Indexes and supplemented by other industry sources and engineering judgment. Page last updated 27 May 2026; accessed 4 August 2026.

U.S. Energy Information Administration (2025). Solar photovoltaic module shipments report, form eia-63b. https://www.eia.gov/renewable/monthly/solar_photo/. Table 6, source and disposition of photovoltaic module shipments, monthly, peak kilowatts. Accessed 4 August 2026.

U.S. Energy Information Administration (2026a). Annual electric power industry report, form EIA-861 detailed data files. Live data page; accessed 4 August 2026.

U.S. Energy Information Administration (2026b). Annual energy outlook, table 8: Electricity supply, disposition, prices, and emissions. Reference case tables, https://www.eia.gov/outlooks/aeo/tables_ref.php; archived vintages at https://www.eia.gov/outlooks/archive/. Block "Prices by Service Category" reports end-use prices split into generation, transmission and distribution. The first data column of each vintage is the history year; later columns are National Energy Modeling System projections. Vintages AEO2015 to the current release accessed 3 August 2026.

Way, R., Ives, M. C., Mealy, P., and Farmer, J. D. (2022). Empirically grounded technology forecasts and the energy transition. *Joule*, 6(9):2057–2082.

Whitman, Requardt and Associates, LLP (2026). The handy-whitman index of public utility construction costs. Continuing semiannual proprietary publication. 2026 vintage consulted. The index is a proprietary family of public-utility construction-cost indexes; publication is semiannual.

Wiser, R., Bolinger, M., et al. (2024). Land-based wind market report: 2024 edition. Technical report, Lawrence Berkeley National Laboratory. Data file, sheets "Wind Equip Imports over Time" (Berkeley Lab analysis of Census USA Trade Online) and "Wind Turbine Prices". Verify the author list against the report cover before submission.

Wiser, R. H., Millstein, D., Hoen, B., Gorman, W., Rand, J., Seel, J., Barbose, G. L., Cheyette, A., Darghouth, N. R., Jeong, S., Mulvaney Kemp, J., O'Shaughnessy, E., and Paulos, B. (2025). Land-based wind energy technology update: 2025 edition. Technical report, Lawrence Berkeley National Laboratory, Berkeley, CA. Data file accompanying the 2025 edition, tab "CapEx Over Time": capacity-weighted average installed cost of US land-based wind projects, dollars per kilowatt nameplate. Successor title to the *Land-Based Wind Market Report* series. Accessed 4 August 2026.

Wyckoff, A. W. (1995). The impact of computer prices on international comparisons of labour productivity. *Economics of Innovation and New Technology*, 3(3-4):277–293.

This supplemental appendix is organized to follow the paper. Appendix B documents the measurement and robustness of the capital-deflator reclassification (Section 3 of the paper). Appendix C calibrates the sourcing comparison of Section 3.3. Appendix D states the two aggregation disciplines behind the output-side analysis: electricity output is differentiated by the time and location of delivery, and not by generating technology. Appendix E documents where the output-value bias lives in official statistics. Appendix F contains the market-data measurement of the bias. Appendices G and H document the crossover calibration and the cross-country harmonization of Section 5. Appendix I presents storage as a second member of the self-correlated class, and Appendix J records the official-source documentation.

A From investment-price gaps to capital-input growth

This appendix establishes the balanced-growth result in Proposition 1 and its counterpart outside balanced growth. Let $N_{R,t}$ denote nominal investment in renewable assets. Under deflator convention $j \in \{b, c\}$, where b denotes the broad convention and c the renewable-specific convention, real investment is

$$I^j_{R,t} = \frac{N_{R,t}}{P^j_{K,t}}, \tag{S.11}$$

and the corresponding productive stock evolves according to

$$K^j_{R,t} = (1 - \delta_R)K^j_{R,t-1} + I^j_{R,t}. \tag{S.12}$$

The two accounts therefore use the same nominal investment and depreciation rate; they differ only in the price index used to deflate investment.

Balanced-growth benchmark. Suppose nominal investment and each deflator grow at constant log rates:

$$\Delta \ln N_{R,t} = n, \qquad \Delta \ln P^j_{K,t} = \pi^j.$$

Real investment under convention j then grows at the constant rate

$$i^j \equiv \Delta \ln I^j_{R,t} = n - \pi^j.$$

Assume $(1 - \delta_R)\exp(-i^j) < 1$ for each convention. On the corresponding balanced-growth path, once the inherited-stock transition has dissipated, the solution to Equation (S.12) is

$$K^j_{R,t} = \frac{I^j_{R,t}}{1 - (1 - \delta_R)\exp(-i^j)}. \tag{S.13}$$

The investment–stock ratio is then constant, and the stock grows at the same rate as real investment:

$$\Delta \ln K^j_{R,t} = i^j.$$

It follows that

$$\begin{aligned} \Delta \ln K^c_{R,t} - \Delta \ln K^b_{R,t} &= i^c - i^b \\ &= \pi^b - \pi^c \\ &= \gamma_K. \end{aligned} \tag{S.14}$$

Thus, under balanced growth, the difference between the two measured capital-growth rates equals the investment-deflator growth gap.

Let ω_R denote renewable assets' share of sector capital services. If capital-services weights are held fixed across conventions, the difference in sector capital-input growth is

$$g_K^c - g_K^b = \omega_R \gamma_K.$$

Holding output and noncapital inputs fixed, the corresponding difference in measured TFP growth is therefore

$$B_K \equiv \widehat{A}^b - \widehat{A}^c = s_K \left(g_K^c - g_K^b \right) = s_K \omega_R \gamma_K,$$

which proves Proposition 1.

Capital accumulation during the transition. The balanced-growth equality need not hold over a finite sample. With a common inherited stock $K_{R,0}$, the two measured stocks are

$$K_{R,t}^j = (1-\delta_R)^t K_{R,0} + \sum_{\tau=1}^{t} (1-\delta_R)^{t-\tau} \frac{N_{R,\tau}}{P_{K,\tau}^j}. \tag{S.15}$$

Normalising the two deflators to the same initial value, their relative levels satisfy

$$\ln\left(\frac{P_{K,t}^b}{P_{K,t}^c}\right) = \sum_{\tau=1}^{t} \gamma_{K,\tau}. \tag{S.16}$$

Equations (S.15) and (S.16) show that every difference between the two stocks is generated by applying different deflators to the same nominal investment flows.

Define the resulting renewable-stock growth gap as

$$\kappa_{R,t} \equiv \Delta \ln K_{R,t}^c - \Delta \ln K_{R,t}^b. \tag{S.17}$$

In general, $\kappa_{R,t} \neq \gamma_{K,t}$, because the capital stock contains the inherited stock and surviving investment from earlier years. Whereas $\gamma_{K,t}$ measures the difference in current real-investment growth, $\kappa_{R,t}$ reflects the accumulated effects of current and past deflator gaps on surviving investment vintages.

At fixed capital-services weights, the exact finite-sample reallocation is therefore

$$B_{K,t}^{\text{fixed}} = s_{K,t}\, \omega_{R,t}\, \kappa_{R,t}. \tag{S.18}$$

Equation (S.18) reduces to Equation (2) when the two stocks are on their respective balanced-growth paths.

Convention-specific rental weights. In the preferred capital-services account, changing the investment deflator also changes current-cost asset values, capital gains, rental prices, and hence the asset weights. Let

$$g_{K,t}^j = \sum_{a \in \mathcal{A}} \bar{v}_{a,t}^j \, \Delta \ln K_{a,t}^j$$

denote sector capital-input growth under convention j, where $\bar{v}_{a,t}^j$ is asset a's two-period capital-services share. Only the renewable stock is reconstructed differently, so the difference between the two capital-input indexes can be written as

$$\begin{aligned} g_{K,t}^c - g_{K,t}^b = {} & \bar{v}_{R,t}^c \, \kappa_{R,t} \\ & + \sum_{a \in \mathcal{A}} \left(\bar{v}_{a,t}^c - \bar{v}_{a,t}^b \right) \Delta \ln K_{a,t}^b. \end{aligned} \tag{S.19}$$

The corresponding productivity reallocation is

$$B_{K,t} = s_{K,t}\left[\bar{v}^{c}_{R,t}\,\kappa_{R,t} + \sum_{a\in\mathcal{A}}\left(\bar{v}^{c}_{a,t} - \bar{v}^{b}_{a,t}\right)\Delta \ln K^{b}_{a,t}\right]. \tag{S.20}$$

The first term is the direct effect of the deflator on measured renewable-capital growth. The second is the change in capital-input growth caused by the response of rental weights. If the weights are held fixed across conventions, the second term is zero and Equation (S.20) reduces to Equation (S.18).

B The capital-deflator reclassification: measurement and robustness

This appendix explains how the capital-deflator reclassification in Table 1 is constructed. The empirical question is narrow: how much does measured capital-input growth change when the renewable assets are deflated by their capacity cost rather than by the broad electric-infrastructure index? The comparison holds nominal investment, output, labour, and intermediate inputs fixed. It therefore changes only the amount of real capital obtained from the same nominal investment.

The calculation proceeds in four steps, building on Section A, which shows how the two deflators, applied to the same nominal investment, generate different measured capital stocks. Section B.1 first explains why a comparison of published statistical vintages cannot isolate this effect. Section B.2 then defines the share of sector capital exposed to the deflator change and explains why a one-period repricing of the published stock is not used as a benchmark. Section B.3 replaces stock values with Hall–Jorgenson rental values. Finally, Section B.4 accumulates the same nominal investment into separate perpetual-inventory stocks under each deflator convention, aggregates them with those rental values, applies the decomposition of Section A into the direct effect and the rental-weight response, and reports the robustness checks. Subsequent sections check the asset classification, the nominal investment base, the measured decline in capacity costs, and the timing of the correction.

B.1 Why the published-vintage difference does not identify the effect

The 2023 Comprehensive Update appears to offer a simple before-and-after test: BEA introduced the renewable-specific deflator in that update, so one might compare published Utilities productivity before and after it. That comparison does not isolate the deflator change because the update revised several parts of the production account at the same time.

The capital revision has the expected sign. A lower investment deflator converts the same nominal expenditure into more real capital, so measured capital input rises. Holding the other terms in the growth-accounting equation fixed, measured TFP should consequently fall. Figure B.1 instead shows that published TFP rises. The vintage difference is therefore not a deflator-only experiment.

Figure B.2 identifies the source of the discrepancy. The output revision is small, and the capital and labour revisions lower TFP. The intermediate-input block contributes about +7.0 log points and more than offsets them. That contribution is recovered as a residual, so it cannot be assigned to a particular intermediate input; moreover, NAICS 22 is broader than the NAICS 2211 electric-power account used in Table 1. The relevant conclusion is simply that the comprehensive update bundled several changes. This is why the paper instead applies both deflators to the same nominal investment flow and holds every other measured input fixed.

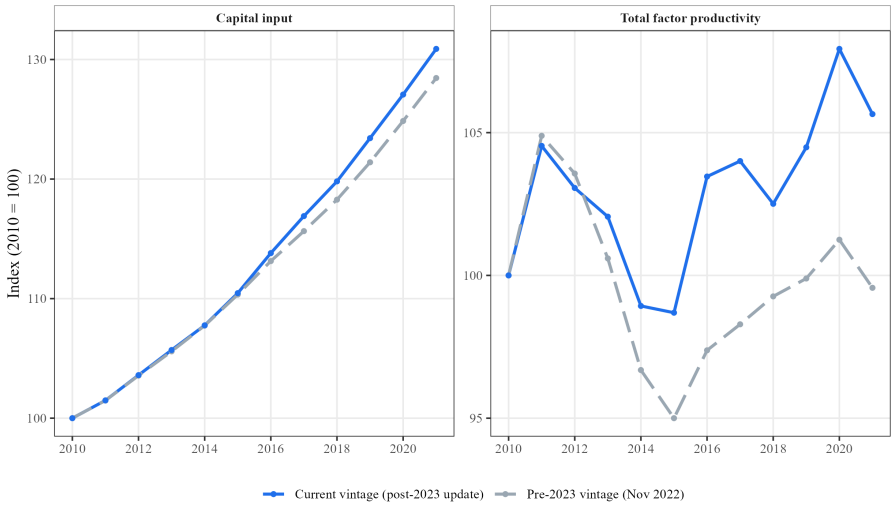


Figure B.1: **Why the 2023 vintage comparison does not identify the renewable-capital deflator effect.** US Utilities (NAICS 22), BLS–BEA KLEMS, 2011–2021 cumulative log growth (Bontadini et al., 2023). The grey long-dashed line reports the pre-2023 vintage and the blue solid line the current vintage. The update raised measured capital input by 1.9 log points, from 25.0 to 26.9, but also raised measured TFP by 5.9 log points, from −0.4 to 5.5. Because the deflator change by itself would raise capital input and lower TFP, the opposite movement in published TFP shows that other revisions dominate the vintage comparison.

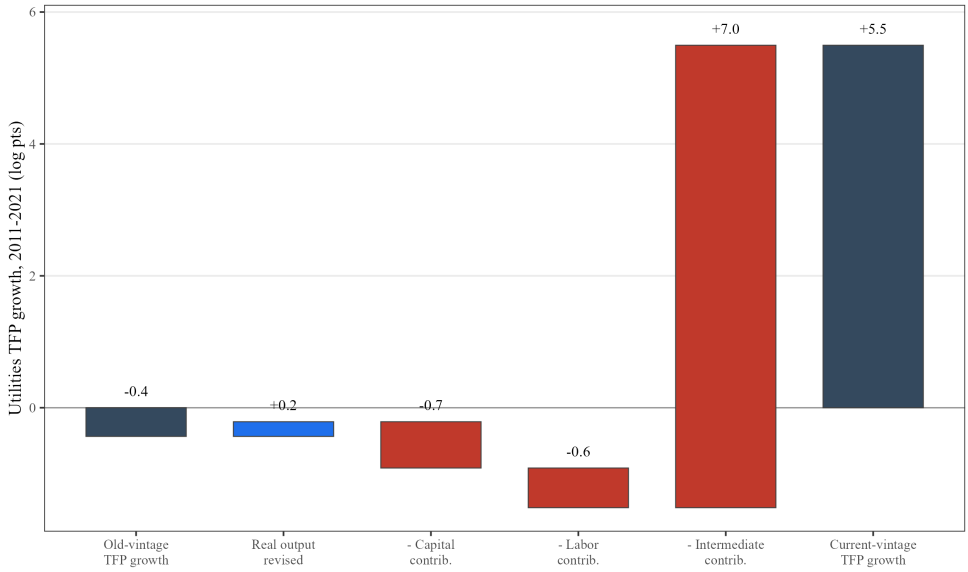


Figure B.2: **Decomposition of the 2023 Utilities TFP vintage revision.** US Utilities (NAICS 22), cumulative log growth over 2011–2021. Read from left to right, the revision moves TFP growth from −0.4 to +5.5 log points: +0.2 from output, −0.7 from capital, −0.6 from labour, and +7.0 from intermediate inputs. The capital contribution has the sign implied by the new deflator, but it is dominated by the intermediate-input revision.

B.2 Renewable-capital exposure

This subsection defines the exposure required by Equation (2) and distinguishes it from the observable stock-value measure. The distinction matters because a stock value asks how much an asset is worth, whereas a capital-services value asks how much of the sector's annual cost of using capital is attributable to that asset.

Let $\mathcal{A}$ denote all assets used in NAICS 2211 and let $\mathcal{R} \subset \mathcal{A}$ denote the assets whose deflator is changed (renewables), that we called "repriced assets". For asset a, let $K_{a,t}$ denote its real productive stock and $R_{a,t}$ its rental price. The exposure required by the capital-input formula is the repriced assets' share of sector capital services,

$$\omega_{R,t} = \frac{\sum_{a\in\mathcal{R}} R_{a,t}K_{a,t}}{\sum_{a\in\mathcal{A}} R_{a,t}K_{a,t}}. \tag{S.21}$$

The denominator covers all sector assets (structure, equipment, and intellectual property products) because the capital share in Equation (2) also covers all sector capital. A structures-only denominator would therefore be inconsistent with the factor share multiplying the exposure term.

Official data do not publish the rental values in Equation (S.21) at this asset detail. They do publish current-cost stocks, which give the observable proxy

$$\widetilde{\omega}_{R,t} = \frac{\sum_{a\in\mathcal{R}} P^K_{a,t}K_{a,t}}{\sum_{a\in\mathcal{A}} P^K_{a,t}K_{a,t}}, \tag{S.22}$$

where $P^K_{a,t}K_{a,t}$ is the current replacement-cost value of the stock. The numerator is BEA asset SU60, "alternative electric" structures. It contains wind and solar structures together with small dry-waste and geothermal components, matching the assets to which BEA applies the alternative-electric deflator. The denominator includes all NAICS 2211 structures, equipment, and intellectual property products. On this sector-wide base, $\widetilde{\omega}_{R,t}$ rises from 3.6 percent in 2011 to 11.2 percent in 2024. The corresponding structures-only ratio reaches 18.2 percent, but it is not used because it omits other assets that supply sector capital services. The published stocks are valued on the broad price basis. The perpetual-inventory reconstruction in Section B.4 reproduces them closely under that convention, with a stock share rising from 3.4 to 11.8 percent, and gives 3.2 to 7.6 percent once the same investment is valued at capacity cost.

Why a one-period repricing is not used as a benchmark. The simplest calculation would multiply observed exposure, the capital share and the cumulative price gap, $s_K\widetilde{\omega}_R\sum_t \gamma_{K,t}$, which is Equation (2) applied as if each year's price gap repriced the whole renewable stock. It does not. With unchanged nominal investment, the two conventions differ by $\gamma_{K,t}$ in real *investment*, but the real *stock* also contains every surviving earlier vintage, so measured stock growth differs by a smaller and smoother amount, κ_t (Section B.4). Over 2011–2024 the accumulated renewable stocks grow apart by $\kappa = 70.3$ log points against a price gap of $\gamma_K = 119.3$. Substituting the price gap for the stock gap overstates the correction by a factor of about 1.7: with the preferred rental weights, the gross-output direct effect would be 5.26 log points instead of 3.13. The paper therefore reports no one-period benchmark. Its transparent comparison is instead the stock-weighted perpetual-inventory account of Section B.4.2, which uses the same accumulated stocks as the preferred account and differs from it only in the weights.

Note that labour-productivity growth can be decomposed into TFP growth and capital deepening:

$$\widehat{Y/L} = \widehat{A} + s_K, \widehat{K/L},$$

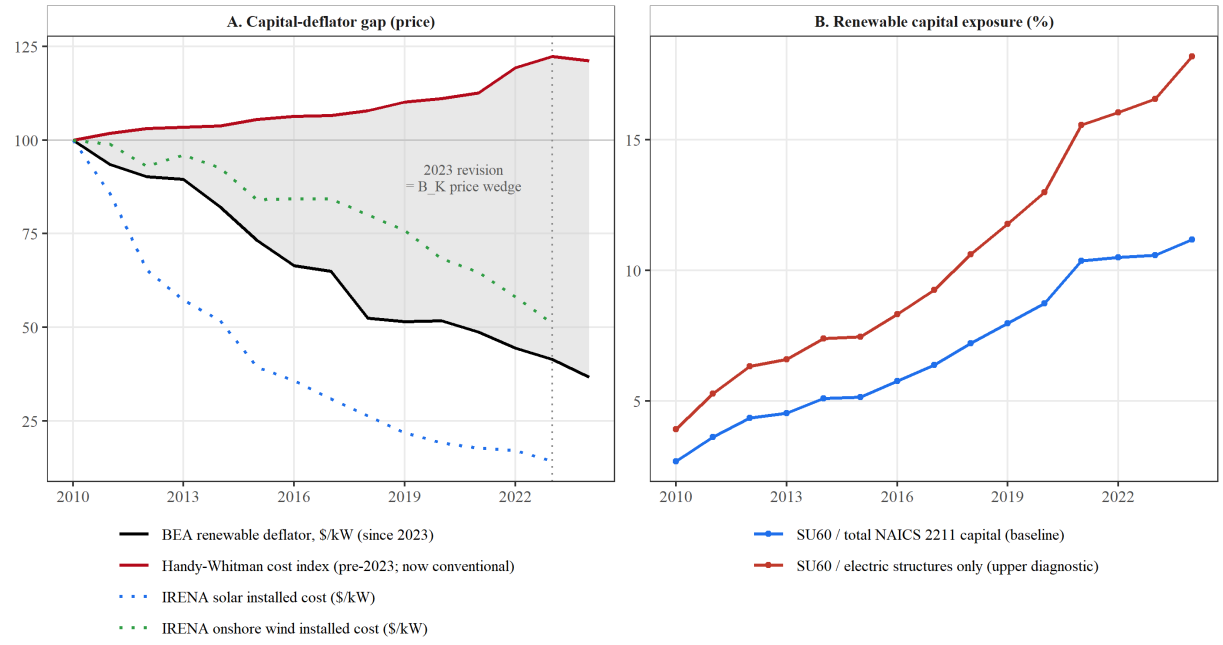


Figure B.3: **The price and exposure components of the capital reclassification.** *Panel A* compares the renewable-specific capacity-cost deflator introduced in the 2023 Comprehensive Update with the broad Handy–Whitman construction index previously applied to alternative-electric structures. The shaded difference is the deflator gap γ_K; IRENA solar and onshore-wind installed costs provide external comparisons. All price series are relative to the GDP deflator and indexed to 2010 = 100. *Panel B* reports the current-cost stock exposure in Equation (S.22), measured against total NAICS 2211 capital.

where s_K is capital's income share. Thus, a correction to the capital deflator that reallocates measured growth from TFP to capital deepening does not reduce measured labour-productivity growth itself; it changes only the channel through which that growth is attributed. Both higher TFP and a higher capital stock per worker can raise output per worker and, under competitive factor pricing, real wages.

B.3 Constructing rental values

The preferred account asks how much service each asset supplies during the year, not only how much its stock is worth. The Hall–Jorgenson rental price answers that question: it is the annual cost of tying up funds in the asset, replacing the part that wears out, and accounting for the gain or loss from a change in the asset's price (Hall and Jorgenson, 1967; OECD, 2009).

For asset a, define own-price inflation as $\pi_{a,t} = \Delta \ln P^K_{a,t}$ and let δ_a denote geometric depreciation. The ex-post rental price is

$$R^{\text{raw}}_{a,t} = P^K_{a,t} \left(r_t + \delta_a - \pi_{a,t} \right). \tag{S.23}$$

Note that the value of $R^{\text{raw}}_{a,t}$ can be negative, I deal with it with nonnegative-weight restriction described below.

The three terms inside parentheses are the required return, depreciation, and the holding loss. A falling asset price makes $\pi_{a,t} < 0$, so the capital-loss term $-\pi_{a,t}$ raises its rental price, holding

the required return and depreciation fixed. Asset-price inflation $\pi_{a,t}$ is indeed the only component that varies both across assets and over time. If its price falls faster than those of other assets, its rental value rises relative to theirs and its normalized capital-services weight changes. This is the key reason the deflator affects rental weights as well as measured real asset growth. Equation (S.23) is the continuous-time form of the standard discrete user-cost formula.

Closing the rental account. The rate of return r_t is not chosen to obtain a particular exposure. It is solved separately in each year so that the raw rental values sum to measured sector capital income. Let $\Pi_t \equiv$ value added$_t$ − labour compensation$_t$ be the capital income. The identifying condition is

$$\Pi_t = \sum_{a \in \mathcal{A}} P^K_{a,t} K_{a,t} \left(r_t + \delta_a - \pi_{a,t} \right). \tag{S.24}$$

Solving this identity gives

$$r_t = \frac{\Pi_t - \sum_{a \in \mathcal{A}} P^K_{a,t} K_{a,t} \left(\delta_a - \pi_{a,t} \right)}{\sum_{a \in \mathcal{A}} P^K_{a,t} K_{a,t}}. \tag{S.25}$$

Thus, observed capital income determines the overall return, while depreciation and asset-price changes determine how that income is allocated across assets. Because the asset prices differ across conventions, the return is solved separately for each. Over 2011–2024 it averages 8.0 percent under the capacity-cost convention, ranging from 3.0 to 12.8 percent, and 8.3 percent under the broad convention.

From rental values to weights. The Törnqvist aggregation requires the asset-level capital-services weights $v_{a,t}$ to be nonnegative. Because a negative rental rate would produce a negative rental value and hence a negative $v_{a,t}$, I impose a lower bound on the net rental rate as

$$\rho_{a,t} = \max\{r_t + \delta_a - \pi_{a,t},\, 0.005\}, \qquad R_{a,t} = P^K_{a,t} \rho_{a,t}, \tag{S.26}$$

and normalize each asset's rental value to obtain

$$v_{a,t} = \frac{R_{a,t} K_{a,t}}{\sum_{b \in \mathcal{A}} R_{b,t} K_{b,t}}. \tag{S.27}$$

Each $v_{a,t}$ is therefore the asset's share of the sector's annual rental value of capital. The floor is applied only after Equation (S.25) is solved: the raw user costs satisfy the capital-income identity, and the floor changes only the weights used for aggregation. Section B.4 reports how much this restriction matters.

Depreciation inputs. I hold asset-specific depreciation rates fixed over the sample. For equipment, conventional structures, and intellectual property products, I use the 2024 BEA implied rates for NAICS 2211,

$$\delta_{\text{equipment}} = 0.0645, \quad \delta_{\text{conventional structures}} = 0.0236, \quad \delta_{\text{IPP}} = 0.3880. \tag{S.28}$$

Conventional structures are total structures less SU60. BEA does not report a separate NAICS 2211 implied rate for SU60. I therefore derive $\delta_{\text{SU60}} = 0.0303$ from BEA's declining-balance parameter of 0.909 and 30-year service life (Bureau of Economic Analysis, nda, p. 10),

$$\delta_{\text{SU60}} = 0.0303$$

The robustness analysis varies this less precisely established rate over $[0.025, 0.060]$ and repeats the full capital-services calculation.

B.4 Capital-services aggregation and the exact reclassification

This subsection connects the rental construction to the preferred numbers in Table 1. There are two effects to keep separate that I detail below, i) repricing SU60 directly raises its measured real growth, ii) It also changes changes the rental weights used to aggregate all assets. The preferred account includes both.

Let $j \in \{\text{capacity}, \text{broad}\}$ denote the deflator convention, let $\widehat{K}^j_{a,t} \equiv \Delta \ln K^j_{a,t}$, and define the two-period average rental weight

$$\bar{v}^j_{a,t} = \frac{1}{2}\left(v^j_{a,t-1} + v^j_{a,t}\right). \tag{S.29}$$

The same weights provide the two objects needed for the calculation: summing the weights on the repriced assets $\mathcal{R}$ measures renewable exposure, while applying all the weights (all of $\mathcal{A}$) to asset growth measures sector capital-input growth,

$$\bar{\omega}^j_{R,t} = \sum_{a\in\mathcal{R}} \bar{v}^j_{a,t}, \qquad \widehat{K}^j_{E,t} = \sum_{a\in\mathcal{A}} \bar{v}^j_{a,t}\widehat{K}^j_{a,t}. \tag{S.30}$$

The first expression selects the renewable part of capital services; the second uses the entire asset distribution to construct the Törnqvist capital-input index.

The sector capital share is computed from nominal current-price payments, so it is identical under the two deflator conventions. It may nevertheless vary over time. Let $\bar{s}_{K,t}$ denote its two-period average. The preferred annual TFP reclassification is the difference between the two capital-input indexes, multiplied by that common share:

$$B^{\text{services}}_{K,t} = \bar{s}_{K,t}\left(\widehat{K}^{\text{capacity}}_{E,t} - \widehat{K}^{\text{broad}}_{E,t}\right). \tag{S.31}$$

Equation (S.31) is the formula used for the total preferred B_K in Panel B of Table 1. Thus, the two rows of that panel differ in their treatment of factor shares: the net-stock benchmark uses sample means, whereas the preferred account uses the annual two-period shares.

Accumulating the stocks. Both conventions start from a common inherited real stock in 2010 and apply the same nominal investment history, $N_{a,t}$, taken from the FAA NAICS 2211 detail. Real stocks then accumulate by geometric perpetual inventory,

$$K^j_{a,t} = (1-\delta_a)K^j_{a,t-1} + \frac{N_{a,t}}{P^j_{a,t}}, \tag{S.32}$$

with end-of-year timing in the preferred account. Current-cost values, capital gains and rental prices are computed separately under each convention, so every input to Equation (S.23) is convention-specific.

Separating the two effects. To show how Equation (S.31) relates to the direct formula in Equation (2), add and subtract broad-deflator asset growth evaluated at the capacity-cost weights,

$$\begin{aligned}\widehat{K}^{\text{capacity}}_{E,t} - \widehat{K}^{\text{broad}}_{E,t} = &\sum_{a\in\mathcal{A}} \bar{v}^{\text{capacity}}_{a,t}\left(\widehat{K}^{\text{capacity}}_{a,t} - \widehat{K}^{\text{broad}}_{a,t}\right) \\ &+ \sum_{a\in\mathcal{A}}\left(\bar{v}^{\text{capacity}}_{a,t} - \bar{v}^{\text{broad}}_{a,t}\right)\widehat{K}^{\text{broad}}_{a,t}.\end{aligned} \tag{S.33}$$

Only the assets in $\mathcal{R}$ are repriced, so only their measured real growth differs across conventions. That difference is not the deflator gap. Because nominal investment is the same, the two conventions' real *investment* differs by exactly $\gamma_{K,t}$ in each year. The real *stock* in Equation (S.32), however, combines that year's investment with every surviving earlier vintage, so the gap in stock growth is a damped and smoothed function of the current and past price gaps. Define it as

$$\widehat{K}_{a,t}^{\text{capacity}} - \widehat{K}_{a,t}^{\text{broad}} = \begin{cases} \kappa_t, & a \in \mathcal{R}, \\ 0, & a \notin \mathcal{R}. \end{cases} \tag{S.34}$$

Over 2011–2024 the price gaps sum to $\gamma_K = 119.3$ log points but the stock gaps to only $\kappa = 70.3$, a damping of 0.589. The two also differ in timing: $\gamma_{K,t}$ is volatile and trendless, from 0.5 log points in 2020 to 22.7 in 2018, whereas κ_t is smooth and rising, between 1.1 and 3.8 log points a year over 2011–2015 and between 5.7 and 8.0 from 2018 onward. Substituting Equation (S.34) into Equation (S.33), and then multiplying by the capital share, gives the exact decomposition

$$B_{K,t}^{\text{services}} = \bar{s}_{K,t} \left[\underbrace{\bar{\omega}_{R,t}^{\text{capacity}} \kappa_t}_{\text{direct effect}} + \underbrace{\sum_{a \in \mathcal{A}} \left(\bar{v}_{a,t}^{\text{capacity}} - \bar{v}_{a,t}^{\text{broad}} \right) \widehat{K}_{a,t}^{\text{broad}}}_{\text{rental-weight response}} \right]. \tag{S.35}$$

The first term is the main-text exposure formula evaluated with the capacity-cost rental weights used in the preferred index. The superscript "capacity" is necessary because the weights are recomputed under each deflator. The second term records the resulting change in aggregation weights; it runs over all assets because the weights are normalized to sum to one. Using the broad weights as the reference would give a different split between these two terms but exactly the same total in Equation (S.31).

Equation (2) is recovered under two restrictions. The first is that the weights do not respond to the deflator convention,

$$\bar{v}_{a,t}^{\text{capacity}} = \bar{v}_{a,t}^{\text{broad}} \qquad \text{for every } a, \tag{S.36}$$

in which case the second term in Equation (S.35) is zero and

$$B_{K,t}^{\text{services}} = \bar{s}_{K,t} \bar{\omega}_{R,t} \kappa_t. \tag{S.37}$$

The weights may still change over time; the restriction is only that changing the deflator does not change them within the same year. The second restriction is $\kappa_t = \gamma_{K,t}$, which holds when each year's price gap reprices the entire renewable stock, as it would for a stock made up of current-year investment alone. Proposition 1 states this contemporaneous, fixed-weight case. With a stock of surviving vintages the second restriction fails, and it matters quantitatively: Equation (S.37) with $\gamma_{K,t}$ in place of κ_t would overstate the direct effect by about two thirds.

Mapping the decomposition to Table 1. The two-period rental exposure rises from 0.038 in 2011 to 0.107 in 2024. It is not monotone: it falls back in several years, most sharply in 2017 and 2020, and its sample maximum, 0.113, occurs in 2022. Summing the annual products $\bar{\omega}_{R,t}^{\text{cap}} \kappa_t$ over 2011–2024 gives the 6.05-log-point capital-input bias in Panel A of Table 1; weighting each year's product by the gross-output capital share gives the direct effect of 3.13 log points.

Table B.1 separates the two steps that take a one-period repricing calculation to the preferred total. Both rows use the preferred rental weights and annual shares, so the comparison isolates the

Table B.1: **From a one-period repricing to the preferred capital-input bias.** NAICS 2211, cumulative 2011–2024 log points, preferred rental weights and annual capital shares throughout. The row marked ⋆ is the direct effect in Panel A of Table 1; the final row is the total B_K in Panel B.

	Gross output		Value added	
	Change	Level	Change	Level
One-period repricing, $\sum_t \bar{s}_{K,t}\bar{\omega}_{R,t}\gamma_{K,t}$	—	5.26	—	8.11
accumulated-stock gap κ_t in place of $\gamma_{K,t}$	−2.13	3.13	−3.36	4.75
Preferred direct effect⋆		**3.13**		**4.75**
rental-weight response	+1.48	4.60	+2.34	7.09
Preferred total B_K		**4.60**		**7.09**

stock-flow step from the weighting step. Replacing the price gap by the stock gap removes about two fifths of the one-period figure. The rental-weight response then adds back part of it.

In full precision,

$$\begin{aligned} B_K^{\mathrm{GO}} &= \underbrace{3.129}_{\text{direct}} + \underbrace{1.476}_{\text{rental-weight response}} = \underbrace{4.605}_{\text{total}}, \\ B_K^{\mathrm{VA}} &= \underbrace{4.754}_{\text{direct}} + \underbrace{2.339}_{\text{rental-weight response}} = \underbrace{7.093}_{\text{total}}. \end{aligned} \tag{S.38}$$

The positive response means that recomputing rental weights reinforces the direct effect. Under the capacity-cost convention the renewable price falls, which raises the renewable user cost through the capital-loss term in Equation (S.23) and so raises the weight on the fastest-growing asset. These totals move gross-output TFP from −3.15 to −7.75 log points and value-added TFP from −5.17 to −12.27 log points. Section B.4.2 shows that the sign of this response, and much of the size of B_K, depend on weighting by rental rather than stock value.

B.4.1 The nonnegative-rental-weight restriction

This check asks whether the preferred result depends on the 0.005 floor in Equation (S.26). A raw ex-post rental rate becomes negative when measured asset-price inflation exceeds the required return plus depreciation. Such a value would give the asset a negative weight in the Törnqvist index, which is why the floor exists.

In the preferred account it never binds. With the return solved annually from Equation (S.25), no asset has a negative raw rental rate in any year under either convention, including the 2010 base year, and this holds in all eight combinations of SU60 depreciation and investment timing reported in Section B.4.3. Floored and raw B_K are therefore identical. The endogenous closure is what keeps the weights well defined: because it lets the return absorb years of unusually high asset-price inflation, such as 2022, it does not leave the holding-gain term to push individual user costs below zero.

The floor does matter once the return is fixed exogenously. The binding cases cluster in 2021–2023, when construction-price inflation spiked, and fall on conventional structures under both conventions and on SU60 under the broad convention only. That asymmetry follows from the two price paths. Under capacity-cost deflation a falling SU60 price is a holding loss and raises its user cost; under the broad deflator, measured SU60 price growth is much higher, so the holding gain lowers its user cost and can push it below zero. Section B.4.3 reports those cases, and every binding asset-year is listed in the replication files.

B.4.2 How capital input is weighted

The accounts in Table 1 share everything except the weights applied to capital input. All of them accumulate the same FAA NAICS 2211 SU60 nominal investment history from a common 2010 inherited real stock, under each of the two annual price conventions. They differ only in whether assets are weighted by their current-cost stock value or by their Hall–Jorgenson rental price. That single choice moves the gross-output reclassification from 1.00 to 4.60 log points.

Table B.2: Capital-deflator reclassification under alternative capital-input weights, 2011–2024

Capital-input account	Direct effect	Weight response	B_K	Capacity-cost TFP
Gross-output basis				
Stock-weighted, maintained shares	1.73	−0.73	1.00	+0.02
Stock-weighted, measured annual shares	2.04	−0.86	1.18	−2.70
Capital services, Hall–Jorgenson	3.13	+1.48	**4.60**	−7.75
Value-added basis				
Stock-weighted, maintained shares	2.74	−1.16	1.58	−1.36
Stock-weighted, measured annual shares	3.08	−1.30	1.78	−4.33
Capital services, Hall–Jorgenson	4.75	+2.34	**7.09**	−12.27

Notes: Log points. In each account the direct effect is $\sum_t \bar{s}_{K,t}\,\bar{\omega}_{R,t}\,\kappa_t$ evaluated at that account's own renewable weight $\bar{\omega}_{R,t}$, and the weight response is $\sum_t \bar{s}_{K,t} \sum_a (\bar{\omega}_{a,t}^{\text{capacity}} - \bar{\omega}_{a,t}^{\text{broad}})\widehat{K}_{a,t}^{\text{broad}}$. The maintained-share rows fix the gross-output and value-added capital shares at 0.4417 and 0.70. Capacity-cost TFP equals broad-deflator TFP minus B_K.

Two forces separate the accounts, and they work in the same direction.

The first is the level of renewable exposure. Wind and solar structures depreciate faster than the transmission, distribution and generation assets that dominate the rest of the sector, so their user cost is high relative to their stock value. Renewable exposure is accordingly 0.038 → 0.107 on rental weights against 0.029 → 0.074 on stock-value weights, and the cumulative capital-input bias is 6.05 rather than 3.92 log points.

The second is the sign of the weight response, which is the more interesting of the two. Repricing renewable investment at capacity cost does not only raise the measured real growth of SU60; it changes the weight every asset receives in the aggregate. Under the capacity-cost convention the renewable price level falls, and a falling asset price has opposite implications for the two weighting concepts. It reduces the replacement-cost value of the installed stock, and therefore the weight that a stock-weighted index gives to the fastest-growing asset. But it raises the user cost of that asset, through the capital-loss term, and therefore the weight that a rental-weighted index gives to it. The response is −0.86 in the stock-weighted account and +1.48 in the preferred one: in the first it partly cancels the direct effect, in the second it reinforces it.

Rental weighting is the appropriate concept here. The quantity being measured is the flow of capital services, and the price of that flow is the user cost, not the asset price. The preferred estimate is therefore 4.60 gross-output and 7.09 value-added log points. It should nevertheless be read as a rental-weighting result rather than one that is invariant to the capital-input construction, and the accounts differ in what they imply for the level of measured productivity as well as for its reallocation: capacity-cost gross-output TFP is −7.75 log points in the preferred account, −2.70 under measured-share stock weighting, and indistinguishable from zero in the maintained-share benchmark. Only the sign of B_K, and the qualitative conclusion that the deployment decade reallocates measured growth from the residual into capital deepening, survives every construction.

B.4.3 Depreciation, the rate of return, and the sample endpoint

The remaining checks address three possible sources of sensitivity. The first is the user-cost construction: SU60 depreciation is the least precisely measured depreciation input, while the preferred rate of return is obtained from the annual capital-income closure rather than imposed externally. The second is the perpetual-inventory starting point. The third is the sample endpoint, because the 2024 capacity-cost observation is provisional.

Table B.3 addresses the first. It varies SU60 depreciation and replaces the annual endogenous return with a fixed nominal return applied to every asset, every year and both conventions. The OECD capital-measurement framework treats such an exogenous return as a standard alternative to the endogenous approach and reports long-run real-return calibrations between 3 and 5 percent across countries (OECD, 2009, pp. 138–140). Because the user-cost formula pairs r with nominal asset-price inflation, the 3 to 5 percent columns should be read as deliberately low nominal returns rather than as exact real-to-nominal conversions. The final column fixes the return at the 8.0 percent sample mean of the capacity-convention endogenous series.

Table B.3: **Sensitivity to the rate of return and renewable depreciation.** NAICS 2211, cumulative 2011–2024 gross-output log points.

	Endogenous r_t	Fixed nominal return			
δ_{SU60}	(preferred)	r=0.03	r=0.04	r=0.05	r=0.080
Capital-input reclassification, B_K					
0.025	4.63	8.33	7.55	6.72	4.86
0.0303 (BEA)	**4.60**	8.37	7.49	6.68	4.85
0.040	4.58	8.32	7.43	6.65	4.84
0.060	4.58	8.26	7.42	6.65	4.84
Capacity-cost TFP					
0.025	−7.71	−14.32	−12.86	−11.36	−8.19
0.0303 (BEA)	**−7.75**	−14.54	−13.06	−11.50	−8.24
0.040	−7.82	−14.93	−13.36	−11.73	−8.32
0.060	−7.91	−15.56	−13.83	−12.07	−8.43
Floor binds (cap./broad)	0/0	2/3–2/5	2/3–2/4	1/2	0/0–0/1

Notes: End-of-year investment timing. Fixed-return columns hold r constant across years, assets and conventions; the final column uses 0.0798, the mean of the capacity-convention endogenous return. The last row counts asset-years in 2011–2024 at which the 0.005 rental-rate floor binds under the capacity-cost and broad conventions, as a range across the four depreciation rates. Mid-year timing lowers the endogenous-return B_K to between 4.44 and 4.57. Fixing the return at the 8.3 percent broad-convention mean instead gives 4.72 to 4.74.

Depreciation hardly matters under the endogenous closure. With the return solved annually, raising δ_{SU60} from 0.025 to 0.060 moves gross-output B_K only from 4.63 to 4.58 log points, and switching to mid-year timing lowers it to between 4.44 and 4.57. Across all eight depreciation-by-timing cells, B_K lies between 4.44 and 4.63.

The return convention matters, through the floor. Fixed returns at or near the endogenous mean give similar answers, 4.72 to 4.86. Low fixed returns do not: at 3 to 5 percent, B_K rises to between 6.65 and 8.37. These are also the cases in which the floor binds in every cell, in up to seven asset-years, and the floor is then doing substantial work. At a 3 percent return its effect on B_K ranges from −6.4 to +4.6 log points across depreciation and timing, and its sign varies because flooring changes each convention's Törnqvist weights rather than adding a fixed amount. Without

the floor these specifications would have negative rental weights, so they are diagnostics of how much an exogenous return must be regularised rather than alternative capital-services indexes. The low-return columns also show why the return matters for the level of TFP: a low common return leaves asset-specific depreciation and capital gains to dominate the rental weights, which raises measured capital-input growth and lowers capacity-cost TFP to between −11.4 and −15.6 log points.

The perpetual-inventory starting point. All accounts start from a common 2010 inherited real stock. Halving or doubling it moves gross-output B_K from 4.70 to 4.43 log points, because a larger inherited stock dilutes the share of new, differently priced vintages in measured stock growth.

Excluding the provisional endpoint. Finally, the 2024 capacity-cost observation is provisional and accounts for 9.7 of the 66.5 log-point decline in that deflator. Ending the calculation in 2023 reduces the cumulative price gap from 119.3 to 108.1 log points and the accumulated-stock gap from 70.3 to 62.3. It lowers the preferred reclassification from 4.60 to 3.96 gross-output log points and from 7.09 to 6.21 value-added log points, and gives capacity-cost TFP of −6.39 and −10.39 log points. The detailed comparison with independent installed-cost series is reported in Section B.6.

Taken together, these checks delimit the claim. The exact size of the correction depends on the capital-input weights and on the return convention, and the rental-weight response is not negligible. Its sign, however, survives stock versus rental weighting, the rental-rate floor, the return convention, depreciation, the inherited stock, and the exclusion of 2024. Under every reported construction, capacity-cost deflation shifts measured growth out of TFP and into capital deepening.

B.5 Exposure sensitivity and nominal-investment stability

This subsection reports two separate checks on the capital correction. First, I stress-test renewable exposure, ω_R, which determines the weight placed on the repriced assets. It asks how much the result changes if renewable exposure is measured differently Second, I test the stability of the nominal investment base by comparing total nominal investment in electric structures across BEA vintages. This comparison shows whether BEA revised the dollars being deflated when it introduced the new price index.

Sensitivity to the asset classification. SU60 also contains small dry-waste and geothermal components, and the boundary between alternative-electric and conventional structures is not observed project by project. I therefore scale SU60 nominal investment up or down by 10 percent from 2011 onward, offset by an equal nominal amount in conventional structures so that total investment is unchanged, and rebuild both perpetual-inventory accounts. Gross-output B_K moves from 4.60 to 4.13 and 5.08 log points, about ±0.48, in the preferred account, and by about ±0.09 in the stock-weighted account. A plausible classification error cannot overturn the result and leaves the sign unchanged.

Did the update change nominal investment? The second check asks whether the two deflators are being applied to a stable nominal series. Over the common 2011–2021 window, cumulative nominal investment in electric structures is $792.6 billion in the current vintage and $796.8 billion in the September 2022 vintage, a revision of only −0.5 percent. Nine of the eleven annual observations are unchanged, and the largest annual revision is 4.5 percent in 2021.

Note that this check has a narrower purpose than the vintage comparison in Section B.1. Here the question is only whether the nominal investment flow underlying the deflator experiment

changed materially. As it did not, the broad and capacity-cost deflators can therefore be applied to effectively the same nominal investment base.

B.6 Does the changing wind–solar mix explain the fall in $/kW?

The BEA alternative-electric deflator is a chained Fisher price index. A pure shift in investment between wind and solar therefore cannot lower the index if the prices of both technologies are unchanged. This subsection addresses the separate empirical question of whether the measured decline is corroborated by independent installed-cost data. I construct a comparable chained index from LBNL data and contrast it with raw average $/kW, which can move when the wind–solar mix changes. These comparisons validate the price decline; they do not enter the construction of γ_K or B_K.

Comparable windows. The BEA deflator is observed from 2010 through 2024. For external validation, the preferred comparison ends in 2023 because the 2024 observations are provisional in both the BEA and LBNL data. The shorter EIA series is available only for 2018–2024. The global IRENA comparison uses 2010–2023. The endpoint check in Section B.4.3 shows separately that excluding the provisional 2024 observation does not change the sign of the capital reclassification.

Constructing a like-for-like comparison. Independent U.S. series on installed costs per kilowatt for utility-scale solar and land-based wind are published by Lawrence Berkeley National Laboratory (LBNL) in its annual market reports (Seel et al., 2025; Wiser et al., 2025). Solar capacity can be measured on a direct-current (DC) basis, using the nameplate capacity of the photovoltaic modules before conversion, or on an alternating-current (AC) basis, using the inverter-limited capacity after conversion. LBNL reports installed costs on both bases but reports the capacity additions used here only in MW-AC. The BEA deflator is built from those EIA data, whose records utility-scale capacity which are expressed AC terms. I therefore construct the benchmark on the observed AC basis. The DC specification uses LBNL's $/W-DC costs and inferred DC capacity and is reported only as a capacity-basis sensitivity. Let $p_{j,t}$ and $q_{j,t}$ denote installed cost per kilowatt and new capacity for technology $j \in \{\text{solar}, \text{wind}\}$. I aggregate the two price series with a chained Törnqvist index,$\bar{s}_{j,t}$ is technology j's two-period average expenditure share,

$$\Delta \ln P_t^{\text{chain}} = \sum_j \bar{s}_{j,t} \, \Delta \ln p_{j,t}.$$

Like the BEA Fisher index, the Törnqvist index cannot move when quantities shift at unchanged technology-specific prices. This is a well known property of chained index.[5]

[5] A chained price index moves only with prices, never with the quantity mix. The Fisher index BEA uses is $P_t^F = \left(P_t^L \, P_t^P\right)^{1/2}$, with $P_t^L = \dfrac{\sum_j p_{jt} q_{j,t-1}}{\sum_j p_{j,t-1} q_{j,t-1}}$ and $P_t^P = \dfrac{\sum_j p_{jt} q_{jt}}{\sum_j p_{j,t-1} q_{jt}}$; under constant prices $p_{jt} = p_{j,t-1}$ every summand of each numerator equals its denominator counterpart, so $P_t^L = P_t^P = 1$ and hence $P_t^F = 1$ for *any* quantities q—the mix cannot move it (the identity test). The Törnqvist index satisfies the same property transparently: $\Delta \ln P_t = \sum_j \bar{s}_{jt} \, \Delta \ln p_{jt}$, with expenditure-share weights $\bar{s}_{jt} = \frac{1}{2}(s_{jt} + s_{j,t-1})$ and $s_{jt} = p_{jt} q_{jt} / \sum_k p_{kt} q_{kt}$; a pure composition shift moves the quantities q_{jt}—hence the shares—at unchanged prices, so $\Delta \ln p_{jt} = 0$ for every j and $\Delta \ln P_t = 0$ whatever the mix, with chaining ($\ln P_T = \sum_t \Delta \ln P_t$) preserving this. The shares enter only as weights on price changes. By contrast the capacity-weighted unit value $V_t = \sum_j p_{jt} q_{jt} \big/ \sum_j q_{jt}$ has $\Delta \ln V_t \neq 0$ under the same shift; the residual $\Delta \ln V - \Delta \ln P$ is exactly the composition (unit-value) bias. That residual is the object decomposed below, and the expenditure-share weighting is what makes the two capacity bases comparable.

For reference, I also calculate average expenditure per kilowatt,

$$V_t = \frac{\sum_j p_{j,t} q_{j,t}}{\sum_j q_{j,t}},$$

which can change solely because the wind–solar mix changes. In Table B.4, the difference between the change in average $/kW and the chained price change is labelled the mix contribution. A positive value means that the changing mix raised average $/kW relative to the chained price path; a negative value means that it lowered it. This decomposition is reported only to interpret the independent LBNL data.

Comparison with an independent U.S. price series. The resulting LBNL index tracks the BEA alternative-electric deflator closely through 2023, while the broad Handy–Whitman plant index rises (Figure B.4). On the nominal basis used in Table B.4, the BEA and LBNL indexes fall by 56.8 and 53.0 log points over 2010–2023, a difference of only 3.8 log points. The much larger gap through 2024 reflects the provisional final observation: the BEA index falls sharply while the LBNL series rises slightly. The close agreement over 2011-2023 therefore provides independent support for both the direction and approximate magnitude of the BEA capacity-cost decline. The 2024 divergence is treated as endpoint uncertainty in Section B.4.3; it does not overturn the earlier agreement. The construction and adoption of the BEA deflator are documented in Section J.

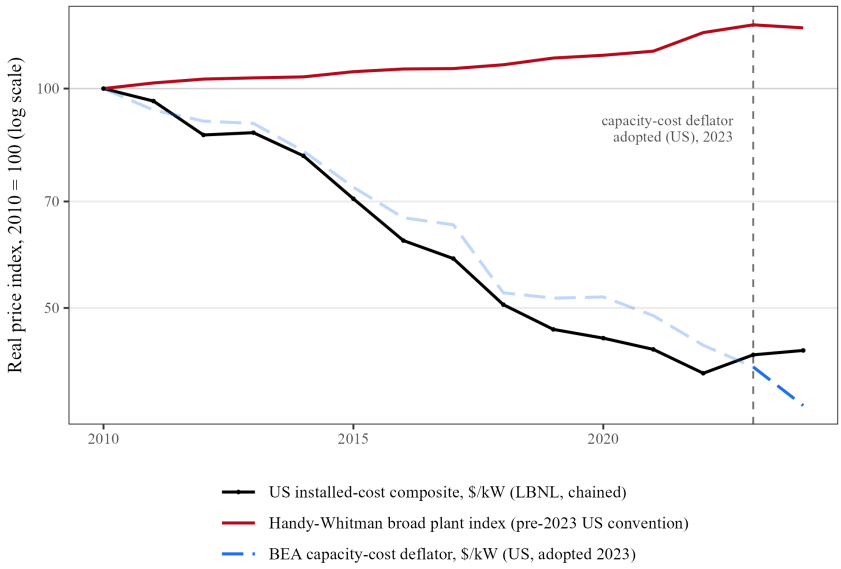


Figure B.4: **The BEA capacity-cost deflator tracks an independent U.S. installed-cost index.** LBNL utility-scale solar and land-based wind installed costs, aggregated with a chained Törnqvist index (black); the BEA alternative-electric deflator (blue, faded before its adoption in 2023); and the broad Handy–Whitman electric-plant index used previously (red). The figure expresses all three series relative to the GDP price index and sets 2010 to 100. The nominal BEA–LBNL difference reported in the text is unchanged by this common deflation. The 2024 divergence reflects a provisional BEA decline not yet matched in the LBNL data.

What the U.S. decomposition shows. On the directly comparable AC basis, the LBNL chained price falls by 53.0 log points over 2010–2023, close to the 56.8-log-point BEA decline.

Table B.4: **Separating within-technology prices from the changing wind–solar mix.** U.S. utility-scale solar and land-based wind installed costs, cumulative log points.

		Within technology		Wind–solar composite			BEA comparison	
Capacity basis	Window	Solar	Wind	Chained price	Mix contribution	Average $/kW	BEA	BEA − chained
AC	2010–2023	−99.8	−19.0	−53.0	+21.1	−31.8	−56.8	−3.8
AC	2010–2024	−96.1	−14.7	−49.1	+19.2	−29.9	−66.5	−17.4
DC	2010–2023	−108.5	−19.0	−55.1	+2.9	−52.3	−56.8	−1.7
DC	2010–2024	−105.2	−14.7	−51.6	−2.5	−54.1	−66.5	−14.9

Notes: Nominal cumulative log points. LBNL reports installed costs in real 2024 dollars; I convert them to nominal dollars with the BEA GDP price deflator before comparison with the nominal BEA investment deflator. The mix contribution is average $/kW minus the chained price; the final column is the BEA change minus the chained price. Components may not add exactly at the displayed precision. LBNL reports solar additions in MW-AC. For the DC rows, DC capacity is inferred as $q_s^{DC} = q_s^{AC}\widehat{\text{ILR}}$, where $\widehat{\text{ILR}} = p_s^{AC}/p_s^{DC}$ is the implied inverter loading ratio. This rescaling preserves solar expenditure and therefore the expenditure weights. The DC rows are an alternative capacity basis, not an independent data source. The 2024 observations are provisional.

The mix contribution is +21.1 log points. Thus, the shift between wind and solar made average $/kW fall less than the underlying within-technology price index, meaning that the mix shift did not produce the measured price decline.

The DC rows ask whether the conclusion depends on measuring solar capacity by module capacity rather than inverter output. This change increases measured solar capacity, especially late in the sample when solar is a large share of additions. It therefore has a large effect on average $/kW but only a small effect on the chained price: the settled-window decline changes from −53.0 to −55.1 log points. The DC construction is a rescaling of the same projects, not independent evidence. Even so, its settled-window mix contribution remains positive at +2.9 log points. With the provisional 2024 observation included, it becomes −2.5 log points, less than 5 percent of the 51.6-log-point chained decline. The conclusion therefore does not depend on the capacity basis.

Checks with other cost data. Two additional comparisons lead to the same conclusion. The EIA generator-cost available data begin only in 2018, so they cannot validate the full 2010–2023 decline. Over 2018–2024, a period dominated by the post-2021 cost increase, their wind–solar composite rises by 17.0 log points: 11.1 points come from within-technology price increases and 5.9 from the changing mix. The mix is therefore not the main component even over this shorter different window.

IRENA provides a longer global check. Over 2010–2023, its installed-cost series fall by 194.7 log points for solar and 67.2 for onshore wind. Holding technology shares fixed gives a composite decline between 131.0 log points under equal wind–solar weights and 180.0 log points when solar receives its 2024 U.S. share of capacity additions. Because these are global rather than U.S. data, they are less directly comparable than the LBNL evidence. They nevertheless show that within-technology declines alone are more than large enough to explain the BEA price decline.

B.7 Subperiod composition of the reclassification

Table 1 reports the accounting bridge from the apparent to the corrected productivity record over the full 2011–2024 window. This subsection asks when within the deployment decade that reclassification accrues. Because $B_K = \sum_t \bar{s}_{K,t}(\widehat{K}_{E,t}^{\text{capacity}} - \widehat{K}_{E,t}^{\text{broad}})$ is a sum of annual terms, any partition of the window is additive: the 2011–2017 and 2018–2024 subperiods sum exactly to the

full-period figure under each construction, with no interaction term to allocate.

The reclassification is concentrated in the second half. Under the preferred capital-services construction, 3.08 of the 4.60 cumulative gross-output log points, two thirds, accrue over 2018–2024, against 1.53 over the first seven years; on a value-added basis the corresponding split is 4.58 against 2.51. Under the two stock-weighted accounts of Section B.4.2 the concentration is stronger, at between 77 and 80 percent of the cumulative correction in the later window on both bases. Two features of the data produce the late concentration. Renewable exposure rises over the decade: the published stock share roughly triples, from 3.6 percent of NAICS 2211 capital in 2011 to 11.2 percent in 2024, and the rental exposure in the preferred account rises from 0.038 to 0.107. The accumulated-stock gap also trends upward, from between 1.1 and 3.8 log points a year over 2011–2015 to between 5.7 and 8.0 from 2018 onward, because each year's price gap persists in the surviving vintages. The annual price gap itself does not trend: $\gamma_{K,t}$ is volatile, ranging from 0.5 log points in 2020 to 22.7 in 2018. The deflator convention therefore matters little for the early 2010s and substantially for the years since 2018.

C Sourcing, local learning, and the net productivity comparison

This appendix calibrates both sides of the sourcing comparison of Section 3.3: the compression of measured electricity-sector TFP when renewable capital is repriced on a local rather than a global learning curve, Equation (8), and the manufacturing TFP generated by the additional domestic production associated with reshoring. Throughout, prices follow experience curves: a technology's installed cost falls by b_j log points per log point of cumulative experience, equivalently by

$$\Lambda = 1 - 2^{-b_j}$$

per doubling. The elasticities are measured below.

Common experience stock and global descent. Let $Q_{j,0}$ denote the effective cumulative experience stock inherited in 2011 for technology j, common to domestic and foreign producers. The renewable capacity-cost index declined by 66.5 log points over 2011–2024; the construction below attributes that descent to the growth of global cumulative capacity in the two technologies separately, and the local-only counterfactual then replaces global with domestic experience while leaving the inherited stock unchanged.

Two technologies. The paper considers a composite renewable good. However I calibrate a learning curve for solar and wind since each technology has different characteristics. Global solar cumulative capacity grew forty-five-fold over the window and global wind that grew six-fold, with learning elasticities that differ by a factor of three.

For technology $j \in \{s, w\}$, let m_j be the growth of global cumulative capacity over 2011–2024, $s_{d,j}$ the domestic share of global additions of the technology j, and b_j the elasticity of the *domestic* installed cost with respect to *global* experience.[6] The local-only counterfactual cumulative price descent in log points is Λ_d and the price descent along the global curve is Λ_g, such that $\Lambda > 0$ means a price decrease,

$$\Lambda_d = \sum_j w_j\, b_j \ln(1 + (m_j - 1)s_{d,j}), \qquad \Lambda_g = \sum_j w_j\, b_j \ln m_j, \tag{S.39}$$

[6] I chose this specific elasticity as i) γ_K is a deflator gap for investment installed in the United States, so domestic, and ii) over 2011–2024 the US bought on the world market. The experience stock actually pricing US-installed equipment was global. Hence, the observed −66.5 is a movement of a domestic price along global experience.

with w_j the technology weight in the composite renewable index.

The elasticities b_j are calibrated on US installed capacity costs as it is the object priced by the deflator. Between 2010 and 2024, these costs fell by 96.1 log points for utility-scale solar and 14.7 for land-based wind (Table B.4), while global installed capacity increased by factors of 45.5 and 6.3, respectively. The implied elasticities are therefore $b_s = 0.252$ and $b_w = 0.080$.

These elasticities are deliberately smaller than learning rates estimated for the underlying technologies at the global level. IRENA reports installed-cost learning rates of about 34% for solar and 25% for onshore wind (International Renewable Energy Agency, 2025, Table 1.1 and Box 1.8), while Way et al. (2022) estimate an experience exponent of 0.319 for solar-PV LCOE against cumulative electricity generation rather than installed capacity. The difference is important here. US installed cost includes not only equipment exposed to global technological learning, but also construction, labour, interconnection, and other soft costs that need not fall along the global technology curve. Applying the global learning rates directly to US cumulative capacity would therefore imply a 142–182 log-point decline, far larger than the 66.5 log points observed for the composite US capacity-cost index.

The solar weight $w_s = 0.636$ is the single calibrated composition parameter: it is chosen so that Equation (S.39) reproduces the observed composite descent. An independent check is that a chained index constructed from the two US installed-cost series tracks the BEA deflator to within 2–4 log points over 2010–2023 (Figure B.4). In the absence of a European counterpart to the LBNL installed-cost series, I apply the same technology-specific elasticities to the EU.

At measured shares of global additions (0.095 in the United States and 0.152 in the EU for solar, and 0.119 and 0.160 for wind) the corresponding local experience curves generate only 27.9 log points of cost decline in the United States and 34.6 in the EU. These are 42% and 52%, respectively, of the 66.5-log-point global descent. The remaining 38.6 and 31.9 log points are therefore the cost decline that would be lost if installed renewable capital followed only the smaller domestic experience stock rather than the global one. Note that this is a comparison between two counterfactual price paths, not a decomposition of observed cost declines into domestic and foreign learning.

Domestic content is primarily a property of the technology. The manufacturing bound requires $\alpha_{d,j}$, the share of demand that can already be met by domestic producers. Across the values used here, domestic-content proxies vary by roughly a factor of four across technologies but much less across regions (Table C.1). Wind turbines are large and costly to transport, encouraging production close to deployment; solar modules are standardised, relatively cheap to ship, and produced at large scale in China. Variation in $\alpha_{d,j}$ therefore reflects technology much more than geography. The values used here should be interpreted as upper-bound proxies: they measure domestic manufacturing capacity relative to projected 2030 deployment needs, rather than current domestic-content shares.

Electricity-sector residual. To translate the capital-price gap into electricity-sector value added, I define the exposure terms $\chi \equiv s_K \omega_R$.

Under import sourcing, the relevant deflator gap is $\gamma_K^{\text{off}} = \Delta \ln P^{\text{broad}} + \Lambda = 52.8 + 66.5 = 119.3$ log points. I use $\chi \in [0.0360, 0.0595]$, so that $B_K^{\text{off}} = \chi \gamma_K^{\text{off}} \simeq 4.3$–$7.1$ value-added log points. The upper end is the preferred value-added correction in Table 1 per log point of price gap, 7.09/119.3. Treating χ as a constant scales that correction in proportion to the cumulative price gap, which is a linear approximation to the perpetual-inventory account of Section B.4: there, a change in the price path reaches measured capital input only through the surviving vintages.

Under full US reshoring, only 27.9 of the 66.5 log points of renewable-capital cost decline are

Table C.1: Domestic content of renewable-equipment demand, $\alpha_{d,j}$.

Region	Technology	α	Basis
US	solar	0.17	module shipments, imports over total, 2022–2024
US	wind	0.72	equipment imports over additions at turbine prices, 2012–2023
EU	solar	0.15	module manufacturing capacity over 2030 deployment needs
EU	wind	0.80	turbine-component capacity over 2030 deployment needs

Sources: US solar, U.S. Energy Information Administration (2025), Table 6, "Source and disposition of photovoltaic module shipments" (0.06 in 2022, 0.21 in 2023, 0.23 over the first nine months of 2024). US wind, Wiser et al. (2024), data file, sheet "Wind Equip Imports over Time" (Berkeley Lab analysis of Census USA Trade Online) over annual additions valued at the OEM turbine transaction prices of sheet "Wind Turbine Prices"; the 2012–2023 range is 0.57 to 0.82 with no trend. EU solar, Chatzipanagi et al. (2025), Table 9, p. 59: EU module manufacturing capacity of 12 $\mathrm{GW_{DC}}$ against 76 $\mathrm{GW_{DC}}$ of 2030 deployment needs; upstream the EU is at 3% of needs for cells and zero for wafers and ingots. EU wind, Tapoglou et al. (2025), Table 6, p. 51: nacelle assembly 36.3/44.7 GW, blades 27.9/44.7, towers 41.7/44.7, weighted by turbine value; deployment needs follow REPowerEU, 510 GW by 2030 against 241.7 GW installed at end-2024 spread uniformly over the remaining years. The components covered are set out in Table 5, p. 50, and exclude the drivetrain, for which no data are available. *All four are upper bounds:* US module assembly from imported cells counts as domestic; the EU figures are manufacturing capacity rather than production, and the source notes that deployment will be exponential rather than linear, which understates the denominator.

reproduced by domestic experience. The deflator gap therefore falls to $\gamma^{\mathrm{re}}_{K,\mathrm{US}} = 52.8 + 27.9 = 80.7$ log points, implying $B^{\mathrm{re}}_{K,\mathrm{US}} = \chi(80.7) \simeq 2.9$–$4.8$ value-added log points. Relative to import sourcing, reshoring compresses the residual by

$$\begin{aligned} B^{\mathrm{off}}_K - B^{\mathrm{re}}_{K,\mathrm{US}} &= \chi(66.5 - 27.9) \\ &= \chi(38.6) \\ &\simeq 1.4\text{–}2.3 \text{ log points.} \end{aligned}$$

For the EU, domestic experience reproduces 34.6 log points of the global descent. Full reshoring therefore gives $\gamma^{\mathrm{re}}_{K,\mathrm{EU}} = 52.8 + 34.6 = 87.4$ log points and $B^{\mathrm{re}}_{K,\mathrm{EU}} = \chi(87.4) \simeq 3.1$–$5.2$ value-added log points. The corresponding compression relative to import sourcing is

$$\begin{aligned} B^{\mathrm{off}}_K - B^{\mathrm{re}}_{K,\mathrm{EU}} &= \chi(66.5 - 34.6) \\ &= \chi(31.9) \\ &\simeq 1.2\text{–}1.9 \text{ log points.} \end{aligned}$$

These calculations are the numerical counterpart of $B^{\mathrm{off}}_K - B^{\mathrm{re}}_K = s_K \omega_R \sum_j w_j b_j \ln\!\left(\frac{m_j}{1+(m_j-1)s_{d,j}}\right)$, which measures, technology by technology, how much of the global experience-driven cost decline would be lost if renewable capital followed only the domestic rather than the global experience stock.

Comparison on a common value-added scale. The same learning-curve logic gives an upper bound on the manufacturing productivity gain from reshoring. Let $\alpha_{d,j}$ denote the share of domestic demand already supplied by domestic manufacturers. Under current sourcing, domestic producers therefore accumulate experience on only that share of post-2011 domestic demand; under full reshoring, they accumulate experience on all of it. The implied gain is

$$\Delta A^{\mathrm{re}}_{M,d} = \sum_j w_j \, b_j \ln\!\left(\frac{1 + (m_j - 1)s_{d,j}}{1 + (m_j - 1)\alpha_{d,j} s_{d,j}}\right). \tag{S.40}$$

If the entire associated unit-cost decline is attributed to domestic manufacturing TFP[7] the gain is at most 18.1 log points in the United States and 21.9 in the EU.

These manufacturing gains are not directly comparable with B_K, which is expressed in electricity-sector value-added TFP units. To put both on the same scale, I define $\nu_d \equiv \frac{VA_{M,d}}{VA_{E,d}}$, the ratio of renewable-manufacturing value added to electricity-sector value added. The electricity-sector-equivalent manufacturing gain is then $\widetilde{\Delta A}^{\text{re}}_{M,d} = \nu_d \, \Delta A^{\text{re}}_{M,d}$.

For the United States, solar-module and wind-equipment demand valued at shipment and turbine transaction prices was \$17.4 bn in 2023. Applying the BEA all-manufacturing value-added share of gross output, 0.394, gives \$6.9 bn of manufacturing value added, compared with \$368.9 bn for NAICS 2211. This implies $\nu_{\text{US}} = 0.0186$.[8] The relevant ratio is the reshored one, since the productivity gain in Equation (S.40) is realised when domestic producers supply the full domestic market; under current sourcing the corresponding ratio is only about 0.007. Moreover, the estimate omits inverters, trackers, mounting equipment and other balance-of-system manufacturing, upstream cells, wafers and polysilicon, and offshore wind. I therefore retain the round value $\nu_d = 0.02$.

On this common scale, the upper-bound manufacturing gains are

$$\widetilde{\Delta A}^{\text{re}}_{M,\text{US}} = 0.02(18.1) = 0.36$$

and

$$\widetilde{\Delta A}^{\text{re}}_{M,\text{EU}} = 0.02(21.9) = 0.44$$

electricity-sector-equivalent log points.

These gains are small relative to the corresponding reduction in the electricity-sector measurement residual. Reshoring compresses that residual by

$$\chi(66.5 - D^{\text{re}}_d),$$

or 1.4–2.2 log points in the United States and 1.2–1.8 in the EU, as summarised in Table C.2.

Table C.2: Reshoring: manufacturing gain versus electricity-residual compression

	Manufacturing gain	Electricity-residual compression
United States	0.36	1.4–2.3
EU	0.44	1.2–1.9

Notes: Entries are in electricity-sector-equivalent value-added log points, evaluated at $\nu_d = 0.02$ and $\chi \in [0.0360, 0.0595]$.

The electricity-sector compression is about three to six times larger than the manufacturing gain, despite assumptions chosen throughout to favour the manufacturing side. The manufacturing bound sets $\rho_d = 1$, attributing the entire learning-curve improvement to domestic manufacturing TFP. It therefore allows no role for international knowledge spillovers, imported machinery and intermediate inputs, input-price declines, scale economies not recorded as TFP, changes in product

[7] Equivalently, if only a fraction $\rho_d \in [0, 1]$ of the learning-curve improvement reflects domestic manufacturing TFP, the manufacturing gain is $\rho_d \Delta A^{\text{re}}_{M,d}$, then attributing the associated unit-cost decline to domestic manufacturing TFP means setting $\rho_d = 1$ throughout, so the reported gains are upper bounds.

[8] Equipment demand from U.S. Energy Information Administration (2025), Table 3, and Wiser et al. (2024); value-added share from BEA GDP-by-Industry, sector 31G; NAICS 2211 value added from BEA Underlying GDP-by-Industry. The corresponding 2022 ratio is 0.0191.

composition, or margins. The domestic-content shares $\alpha_{d,j}$ are also upper-bound proxies; higher initial domestic content leaves less additional experience to be gained from reshoring. By contrast, the electricity-sector compression is bracketed using the measured range of capital exposure κ. Reshoring therefore shifts some learning into domestic manufacturing, but the resulting productivity gain is too small to offset the reduction in measured electricity-sector TFP when the global cost curve is replaced by the shallower domestic one.

Single-curve alternatives. The two-technology specification is not essential for this result. Collapsing solar and wind into a single learning curve requires choosing one elasticity and one experience stock, but plausible alternatives give the same qualitative conclusion. First, using the solar-PV experience exponent $b = 0.319$ from Way et al. (2022, Table S24), the observed 66.3-log-point cost decline corresponds to three doublings of experience, since $3b \ln 2 = 0.663$. Applying the corresponding domestic shares of experience, the local curves reproduce 17.3 log points of decline in the United States and 23.4 in the EU. Second, observed global wind-and-solar capacity rises from 222 to 3,013 gigawatts, a factor of 13.6. Holding the observed cost decline fixed and using this growth factor to calibrate the common elasticity gives $b = 0.255$, with local descents of 21.2 log points in the United States and 27.5 in the EU. The preferred two-technology construction gives 27.9 and 34.6 log points, respectively. Under all three specifications, domestic experience reproduces only a minority of the global cost decline, and the upper-bound manufacturing gain remains well below the corresponding compression of the electricity-sector residual.

D Output aggregation: by delivery time, not by generation source

The output-side analysis rests on two aggregation disciplines that pull in opposite directions and are both required. Electricity output should be differentiated by the hour and location of its delivery, because that is the margin on which its market value differs; and it should not be differentiated by the technology that produced it, because within a delivery cell a MWh is a homogeneous product. This appendix states both disciplines. Section D.1 defines the reference price behind the capture ratios of Section 2. Section D.2 makes the conceptual case that a shift of generation toward low-value hours is a fall in real output, not a price effect. Section D.3 shows that the opposite convention—splitting output by generation source and deflating each stream by its own production cost—creates artificial real-output growth.

D.1 The reference price and the value-weighted benchmark

Let

$$\bar{p}_t = \sum_h \tau_h \, p_{h,t}, \qquad \tau_{(m,n)} = \frac{\lambda_n}{M_t}, \qquad \sum_h \tau_h = 1,$$

be the reference price, the load-weighted average of each zone's baseload price, where λ_n is location n's share of load. One can check that $\sum_h \tau_h = \sum_n \sum_m \lambda_n / M_t = \sum_n \lambda_n = 1$: hours enter uniformly, locations by economic size. Because $v_t \equiv v_{r,t}/v_{d,t}$ is a ratio of capture ratios sharing $\bar{p}_t$ as a denominator, the reference price cancels: the weights λ_n affect neither Q_t^{val} nor $B_{V,t}$, and fix only the units in which capture ratios are quoted. For a single zone this is the arithmetic mean of hourly prices.

Q^{phys} and Q^{val} are conceptual benchmarks, not published series. Official real output, generically Q_t^{off}, is constructed in one of three ways: as a direct quantity index over strata such as customer classes, as current-dollar output deflated by a price index, or as a quantity extrapolation of a

benchmark level. None of the documented constructions carries the hour or location of delivery as a stratum of the aggregation, so each behaves like Q^{phys} on exactly the margin along which renewable generation differs, while correctly refusing to split output by source. Section 4.1 states the premise; Supplemental Appendix E.1 documents the constructions and derives the condition under which the underlying output-value wedge survives as bias in a deflated-value series.

D.2 Why a shift toward low-value hours is a fall in real output

Proposition 2 records lower real output when the same annual MWh are delivered in lower-value hours. The natural objection is that a MWh is a MWh: if physical generation is unchanged, how can real output fall? The answer is that a volume index measures the basket, not the tonnage. Electricity delivered in different hours is not the same good, so a basket that shifts toward hours buyers value less is a smaller basket even at constant MWh. Three features of existing accounting practice already embody this principle, in ascending order of proximity to the electricity case.

First, seasonal goods. Strawberries in December and strawberries in July are different products in the accounts; a grower who keeps tonnage constant but shifts sales from December to July has produced less valuable output, and the volume index says so. The Eurostat handbook places summer versus winter gas in the same sentence as day versus night electricity and cross-references its own treatment of seasonal products (Eurostat, 2016, §4.4). Electricity, the least storable good in the accounts, is the strongest case for time-differentiation, not an awkward exception to it.

Second, the accounts already differentiate electricity by reliability, through tariffs. Interruptible industrial supply is priced below firm supply and treated as a different product, because the buyer has surrendered the right to consume at a moment of its choosing. That is time-and-system-state differentiation already inside the accounts; moving from "firm versus interruptible" to "hour by hour" is a change of resolution, not of principle.

Third, curtailment provides a useful limiting case. Curtailed electricity is not counted as output because it cannot be absorbed by the system. The same logic applies, more gradually, to generation delivered in hours of very low or negative prices: its economic value is lower because supply is abundant relative to demand. This does not require treating electricity from different hours as physically different goods. It only requires recognising that time of delivery is a value-relevant characteristic. An annual index that assigns the same weight to every MWh abstracts from that margin. The issue is therefore one of aggregation: as generation becomes increasingly concentrated in low-value hours, annual equal-MWh weighting becomes a progressively poorer approximation to the value of delivered electricity.

The accounting standard gives the same conclusion. ESA 2010 distinguishes quality differences from price discrimination by asking whether identical products are sold at different prices "under exactly the same circumstances" (ESA 2010, §10.16). Electricity delivered at different times does not meet that condition: the circumstances of delivery, including the state of the system, differ across hours. Consistently, the Eurostat handbook treats electricity supplied at night and during the day as an example of a quality difference that can be recorded as different products. Appendix E.8 documents the relevant provisions.

The physical unit is the same across hours, but the service is not. A kilowatt-hour can boil the same kettle at noon or at midnight, yet the consumer generally cares when the kettle is boiled. That timing value is reflected in the price: consumers willing to give up control over when electricity is supplied can already buy interruptible service at a discount. The same logic applies on the production side. As renewable penetration rises, more generation arrives in hours when electricity is abundant and therefore less valuable. This does not mean that more sunshine reduces welfare, or even that value-weighted electricity output falls in level terms. It means only that an additional

midday MWh contributes less to output than a MWh delivered when electricity is scarce. The benefit of cheap midday power is real, but part of it appears outside measured electricity-sector output, through lower input costs for firms and lower effective energy costs for households. As noted in the discussion of Proposition 2, B_V is a market-value benchmark, not a welfare measure.

D.3 The homogeneous-output discipline: why output is not split by source

Within a node and settlement interval, electricity is homogeneous across generation technologies. To isolate this point, this section sets the output-value gap to zero: all MWh delivered in the same market cell receive the same transaction price p, regardless of whether they were generated by wind, solar, nuclear, or fossil fuel. Differences in value across hours and locations are therefore suppressed here and treated separately in Section 4.

The distinction between technologies remains relevant on the input side. Renewable and conventional generating structures are physically different capital goods with separately observed replacement costs. It does not follow that their electricity injections are separate output products. Once delivered in the same market cell, a MWh is a MWh. Splitting output by generation source and assigning each stream its own production-cost index would therefore treat a production process as product variety.

Let the sector produce q_d conventional and q_c clean MWh, with

$$Q = q_d + q_c, \qquad \theta = \frac{q_c}{Q}.$$

Because both streams receive the common output price p, nominal electricity sales are

$$S = p(q_d + q_c) = pQ,$$

and deflating nominal sales by this common transaction price returns physical output:

$$Q^{\text{true}} = \frac{S}{p} = Q.$$

Consider instead an analyst who assigns the two technologies separate imputed prices c_d and c_c, such as source-specific levelized costs, and constructs the cost-valued aggregate

$$\widetilde{S} = c_d q_d + c_c q_c = \bar{c} Q, \qquad \bar{c} = (1-\theta)c_d + \theta c_c.$$

These imputed prices are production costs, not transaction prices for distinct output products.

Proposition 3 (Composition artifact from source splitting)**.** *Suppose the analyst treats clean and conventional electricity as separate products and deflates $\widetilde{S}$ with a separated-goods price index satisfying*

$$d\log P^{\text{sep}} = \omega_d \, d\log c_d + \omega_c \, d\log c_c, \qquad \omega_i = \frac{c_i q_i}{\widetilde{S}}.$$

To first order, the resulting real-output index differs from physical output by

$$d\log Q^{\text{sep}} - d\log Q = d\log \bar{c} - d\log P^{\text{sep}} = \frac{c_c - c_d}{\bar{c}} \, d\theta. \tag{S.41}$$

A shift toward the technology with the higher imputed cost therefore appears as additional real output even when total physical generation is unchanged.

Proof. Because

$$\widetilde{S} = \bar{c}Q,$$

we have

$$d\log \widetilde{S} = d\log Q + d\log \bar{c}.$$

The separated-goods quantity index satisfies

$$d\log Q^{\text{sep}} = d\log \widetilde{S} - d\log P^{\text{sep}}.$$

It follows that

$$d\log Q^{\text{sep}} - d\log Q = d\log \bar{c} - d\log P^{\text{sep}}.$$

Differentiating

$$\bar{c} = (1-\theta)c_d + \theta c_c$$

gives

$$d\log \bar{c} = \omega_d\, d\log c_d + \omega_c\, d\log c_c + \frac{c_c - c_d}{\bar{c}}\, d\theta.$$

Subtracting the separated-goods price index yields Equation (S.41). □

The result is a pure composition effect. If the imputed costs are fixed, a standard separated-goods price index records no inflation. A shift toward the higher-cost source nevertheless raises the analyst's cost-valued nominal aggregate, so the entire increase is recorded as real output.

Worked illustration. Hold total physical output fixed and let the clean share rise from 0.40 to 0.85. Suppose the imputed costs remain constant at

$$c_d = 50, \qquad c_c = 60.$$

The cost-valued unit value rises from

$$\bar{c}_0 = 0.60(50) + 0.40(60) = 54$$

to

$$\bar{c}_1 = 0.15(50) + 0.85(60) = 58.5.$$

Because the two source-specific cost indexes have not changed,

$$\frac{P_1^{\text{sep}}}{P_0^{\text{sep}}} = 1.$$

The separated account therefore reports

$$\log\left(\frac{58.5}{54}\right) \simeq 0.080,$$

or approximately 8 log points of real-output growth, despite unchanged physical generation. Deflation by the unit value of the cost-valued aggregate would remove this composition term and return zero quantity growth.

Remark 1 (The LCOE trap). A second artifact arises when actual electricity revenues are divided by source-specific production-cost indexes. Let

$$S_i = pq_i$$

denote the actual revenue associated with source i. Because the common output price is p, revenue shares equal generation shares. An analyst who deflates each revenue stream by its source-specific cost index obtains

$$d\log Q^{\mathrm{LCOE}} = \sum_{i\in\{c,d\}} \theta_i \left(d\log S_i - d\log c_i\right).$$

Using

$$d\log S_i = d\log p + d\log q_i,$$

the difference from physical-output growth is

$$d\log Q^{\mathrm{LCOE}} - d\log Q = \sum_{i\in\{c,d\}} \theta_i \left(d\log p - d\log c_i\right). \tag{S.42}$$

Suppose the conventional cost index moves with the common output price, while the clean cost index falls relative to it at rate

$$\gamma_c = d\log p - d\log c_c > 0.$$

Equation (S.42) then reduces to

$$d\log Q^{\mathrm{LCOE}} - d\log Q = \theta\gamma_c. \tag{S.43}$$

This term exists even at a fixed generation mix. Falling clean production costs are recorded as greater real electricity output, although neither the physical quantity nor the market value of electricity has increased.

The relative cost decline γ_c is not published, but it is bounded by the gap between output-price growth and the observed decline in clean generation cost. Over 2010–2024 the NAICS 2211 output deflator rises 2.2% per year, while clean cost falls by between 4.8% per year on the BEA renewable capacity-cost index and 18.0% per year on IRENA solar LCOE; the implied γ_c therefore spans 7 to 20% per year, with the onshore-wind LCOE estimator at 11.6.[9] Taking a clean generation share of $\theta \simeq 0.17$ and the mid estimator $\gamma_c \simeq 12\%$ per year,

$$\theta\gamma_c \simeq 0.17 \times 0.12 = 0.0204.$$

Source-cost deflation would therefore add approximately 2 percentage points per year to measured sectoral output growth, and between 1.2 and 3.5 points across the range of γ_c above. At a clean share of one third, the same convention would add approximately 4 percentage points per year. Over a longer period, the cumulative artifact is

$$\sum_t \theta_t\gamma_{c,t}.$$

Because the clean share rises over the transition, the current annual rate should not be extrapolated backward over the full sample.

[9] Estimators and sources in `Results/provenance_gap_closure.md`. The arithmetic below is linear in γ_c, so the span carries through proportionally.

The two artifacts are distinct: Equation (S.41) is a composition effect created when an invented relative price is combined with a changing generation mix, whereas Equation (S.43) is a learning effect created when a falling input cost is used as an output deflator; both are separate from B_V, which reflects genuine differences in the market value of electricity delivered across hours and locations.

The official output series studied in Appendix E avoid both phantom-output terms because they neither split electricity by generation source nor deflate renewable output by LCOE; their limitation instead is that aggregation removes the time and location variation required to identify B_V.

E Where the output-value bias lives in official statistics

Proposition 2 compares two aggregation rules for delivered electricity: one that weights every MWh within the year equally, and one that weights MWh by the price prevailing in the hour and location of delivery. The proposition is a statement about aggregation and holds by construction. What requires documentation is the premise connecting it to published statistics: that official real electricity output behaves like the first rule rather than the second on exactly the margin along which renewable generation differs. This appendix establishes that premise from current agency methodology.

Section E.1 classifies the documented official constructions into three families and previews the form the premise takes in each. Section E.2 derives the accounting condition required in the deflated-value family, where the premise cannot be read off the deflator alone. Sections E.3–E.5 document the US constructions; Section E.6 examines what the available price indexes can and cannot register; Section E.7 calibrates the one quantity that remains unidentified, the settlement pass-through. Sections E.8 and E.9 document the European framework and the two direct-quantity systems that implement its alternative branch. Section E.10 states what would overturn the premise, and Section E.11 summarizes. The empirical counterpart—the demonstration that the gap emerges from the same hourly data once hour and system state are introduced as strata—is in Appendix F.6.

E.1 Three families of official construction

Documented official measures of real electricity output fall into three families. In the first, *direct quantity aggregation*, the agency aggregates physical quantities across strata with value weights; nominal output and the deflator play no role in the volume index, whose properties follow entirely from the strata and weights. In the second, *deflated value*, real output is current-dollar output divided by a price index, and the volume properties depend jointly on what reaches the nominal numerator and what the price index absorbs. In the third, *quantity extrapolation*, a benchmark nominal level is carried forward by a physical quantity indicator and a price index, and real output tracks the quantity indicator by construction. The harmonised European framework treats the first two as equivalent A-methods for electricity, requiring for the direct branch quantities *sufficiently detailed by tariff and customer type* (Eurostat, 2016, CPA 35).

No stratum in any row of Table E.1 is defined by the hour or location of delivery. In the direct-quantity and extrapolation families that observation settles the premise outright: within each stratum-year every MWh carries the same weight, so the index inherits the equal-value-per-MWh property of Q^{phys} on the time-of-delivery margin, with no further condition. In the deflated-value family the same observation is necessary but not sufficient, for the reason the next subsection makes precise.

Table E.1: Documented constructions of official real electricity output.

System	Family	Strata of the volume aggregation
BLS industry productivity (documented utility method)	direct quantity (Törnqvist)	kWh sold to ultimate customers, by class of service (Bureau of Labor Statistics, nd)
BLS industry productivity (general method)	deflated value	revenues deflated at product-class level; product classes are customer classes and the utility/non-utility split (Bureau of Labor Statistics, 2017)
BEA industry accounts	deflated value	commodity cells deflated with electric-power PPIs, Fisher-aggregated (Bureau of Economic Analysis, ndb)
BEA personal consumption (household electricity)	quantity extrapolation	EIA residential kWh carried by the BLS electricity CPI (Bureau of Economic Analysis, 2025)
Eurostat / ESA 2010 (harmonised)	either A-method	PPI deflation, or quantities by tariff and customer type (ESA 2010, §10.32) (Eurostat, 2016)
UK ONS Index of Production, Section D	direct volume	energy-supply volume data supplied by DESNZ (Office for National Statistics, 2024)
Norway SSB national accounts	quantity extrapolation	monthly physical generation, inflated by recipient-group price indexes (Bougroug, 2021; Statistics Norway, 2026a)

E.2 The deflated-value family: the premise is an accounting condition, not a property of the deflator

Write the capture factor of Section 2 as

$$C_t \equiv 1 - \theta_t(1 - v_t), \qquad Q_t^{\text{val}} = Q_t^{\text{phys}} C_t, \tag{S.44}$$

so that $B_{V,t}^{\text{gen}} = -\Delta \log C_t$. Under deflated value, official real output is a ratio,

$$Q_t^{\text{off}} \equiv \frac{V_t}{P_t}, \tag{S.45}$$

of current-dollar output to a price index. Suppressing price components orthogonal to the value factor, let $\bar{p}_t$ denote the reference price of Section 2 and index the incidence of the capture factor on each side of Equation (S.45) by two exponents,

$$V_t = \bar{p}_t \, Q_t^{\text{phys}} \, C_t^{\varphi}, \qquad P_t = \bar{p}_t \, C_t^{\psi}, \tag{S.46}$$

where $\varphi \in [0, 1]$ is the *settlement pass-through*: the share of the capture-value change reaching recorded current-dollar output; and $\psi \in [0, 1]$ is the share the price index absorbs. Then $Q_t^{\text{off}} = Q_t^{\text{phys}} C_t^{\varphi - \psi}$ and the bias remaining in the published series is

$$\widehat{Q}_t^{\text{off}} - \widehat{Q}_t^{\text{val}} = (1 + \psi - \varphi) \, B_{V,t}^{\text{gen}}. \tag{S.47}$$

Two consequences discipline everything that follows. First, within this family the premise requires $\varphi \simeq \psi$, not $\psi \simeq 0$. Establishing that the deflator carries no capture-price signal is by itself

Table E.2: Incidence of the capture factor and the bias remaining in published real output, from Equation (S.47).

φ	ψ	Published real output behaves like	Residual bias
0	0	Q^{phys}	B_V^{gen}
1	0	Q^{val}	0
1	1	Q^{phys}	B_V^{gen}
0	1	below Q^{phys}	$2B_V^{\text{gen}}$

uninformative, and taken in isolation points the wrong way: the second row of Table E.2 shows that if the capture-value loss reaches recorded revenue while the deflator ignores it, the loss remains in measured real output, which then tracks Q^{val} and leaves no bias at all. An argument resting only on the deflator therefore cannot support Proposition 2. Second, Equation (S.45) is not how every official series is built. Where the agency constructs real output directly from physical quantities, there is no ratio, φ and ψ do not arise, and the premise follows from the aggregation strata alone. That is why this appendix documents constructions before deflators.

E.3 US industry productivity: the documented method is a quantity index over customer classes

The BLS industry productivity program defines output on a sectoral basis: "Output is the total amount of goods and services produced in an industry for sale either to consumers or to businesses outside that industry." On construction, the program distinguishes two methods:

> "Real output is most often derived by deflating nominal sales or values of production using BLS price indexes, but for a few industries it is measured by physical quantities of output."
> "Where possible, physical quantity output indexes are Törnqvist aggregations of the quantities of individual products."[10]

Utilities are named among the physical-quantity sectors, and the program's electric-utility methodology specifies the quantity and the weights: output is "based on KWHs sold to ultimate consumers, not KWHs generated"; "KWHs sold to each class of service are weighted by the price of that service"; and class-of-service weighting "prevents bias in productivity measures due to changes in the distribution of sales among service classes" (Bureau of Labor Statistics, nd). This is the premise of Proposition 2 in the agency's own terms: real output is a Törnqvist index of kilowatt-hours sold, and the weights are class-of-service prices. The aggregation resolves the customer to whom a MWh is delivered, not the hour or location at which it is delivered. Every MWh sold to a given class within a given year enters with the same weight, irrespective of when it was produced — the equal-value-per-MWh property, holding on exactly the margin along which wind and solar differ from dispatchable generation, while correctly refusing to distinguish output by generating technology (Supplemental Appendix D.3). Under this construction the incidence problem of Section E.2 does not arise: recorded revenue is not an input to the quantity index, and the published output price index is implicit, recovered as current-dollar output divided by real output.

[10] Bureau of Labor Statistics (2017, Calculation). The physical-quantity method is used for "some industries within the mining, utilities, and transportation sectors, and at community hospitals, commercial banking, and postal services."

Uncertainty (i): industry and vintage coverage. Two gaps should be stated plainly. The handbook names "some industries within the . . . utilities . . . sectors" without enumerating them at the six-digit level, and the electric-utility methodology note is written on the SIC basis (SIC 491–493) that preceded the NAICS conversion, with source data from the Energy Information Administration. I have not located a current document stating that NAICS 2211 specifically is among the physical-quantity industries today. The alternative is that 2211 output is now built by the deflated-value method, which Section E.4 treats on its own terms. The premise survives under that method as well, for a different and weaker reason, and the two cases are kept separate throughout.

Uncertainty (ii): no official 2211 multifactor series. BLS publishes labor productivity and real sectoral output at NAICS 2211, but its published multifactor productivity measure for utilities is at NAICS 22, which includes natural gas distribution and water and sewage. The official series used as a comparator in this paper (`MPU0022`) is therefore broader in scope than the electricity sector and is labelled as a proxy; the electricity-only multifactor series is constructed and validated against `MPU0022` on a concept-matched gross-output basis. The BEA–BLS integrated production account likewise reports an aggregated Utilities industry rather than 2211 (Bureau of Economic Analysis and Bureau of Labor Statistics, 2026). No claim in this paper requires an official 2211 multifactor series to exist.

E.4 The deflated-value alternative, and why the premise weakens but survives

Where the deflated-value method is used, BLS describes it as follows: "Deflated value output indexes are derived from data on the value of industry output adjusted for price change"; "Revenues are deflated at the detailed product-class level to take into account changes in the mix of products over time"; and "The resulting real output indexes are conceptually equivalent to indexes that are developed using data derived from physical quantities of output" (Bureau of Labor Statistics, 2017, Calculation). The equivalence claim is informative about the cell structure: deflation is performed at the product-class level, and for electric power the product classes available in the price statistics are the customer classes of the distribution indexes and the utility/non-utility split of the generation index (Section E.6). No product class in that scheme is defined by hour or location of delivery.

Consequently, under the deflated-value method $\psi \simeq 0$ on the time-of-delivery margin: a shift of generation into low-price hours changes no product class's price, so the price index does not move on that account. From Equation (S.47) the residual bias is then

$$(1 - \varphi)\, B_{V,t}^{\text{gen}}, \tag{S.48}$$

and the premise requires φ to be small rather than following automatically; Section E.7 calibrates φ by settlement channel. The premise is therefore robust to which construction BLS currently applies, in the sense that neither contains a time-of-delivery weight, but the two are not equally strong: under the physical-quantity method the bias is B_V^{gen} unconditionally, and under deflated value it is B_V^{gen} scaled by $(1 - \varphi)$. Results in the main text are reported on the understanding that Equation (S.48) is the conservative case.

E.5 US national accounts: quantity extrapolation on the household component

The BEA industry accounts build industry real gross output by deflation: chain-type quantity indexes are prepared by deflating the current-dollar commodity composition of gross output with the corresponding commodity price indexes and combining the resulting commodity quantity indexes, with value added obtained by double deflation (Bureau of Economic Analysis, ndb). On the

household side, however, the construction is a quantity extrapolation. The current methodology summary gives, for the electricity and natural gas component of personal consumption expenditures, the same source data at every vintage — benchmark years, non-benchmark years, the two most recent years, and the current quarterly estimates: "annual EIA data on kilowatt-hour and cubic-foot quantities used and BLS CPIs for electricity and natural gas" (Bureau of Economic Analysis, 2025, PCE, housing and utilities); the quarterly key-source documentation lists the same combination of EIA residential sales in MWh and the BLS electricity CPI. Current dollars are physical quantities carried by a price index, and real consumption is obtained by deflating with that same index, so real electricity consumption is identically the EIA physical quantity,

$$\text{real}_t = \frac{K_t \times C_t}{C_t} = K_t, \tag{S.49}$$

up to benchmark scaling and BEA's timing and coverage adjustments. Recorded revenue does not enter the construction, so on this component $\varphi = \psi = 0$ by construction and the equal-value-per-MWh property holds exactly. The same document reports that electric-utility components of government consumption and of private inventories are estimated from monthly EIA physical quantities combined with BLS PPIs, with real values obtained by direct valuation from quantities and prices.

Uncertainty (iii): a superseded handbook description. The December 2024 NIPA Handbook describes the same component differently, with annual current-dollar estimates built from EIA residential *revenue* and the most recent estimates from monthly EIA revenue, together with EIA kilowatt-hours and cents per kilowatt-hour adjusted from a billing to a usage basis and the BLS electricity CPI (Bureau of Economic Analysis, 2024a, Ch. 5). Under that description current dollars are observed revenue and Equation (S.49) does not follow. BEA states that the handbook reflects the 2024 annual update and was not revised for the 2025 update; the September 2025 summary is the current statement of source data and is the one relied on here. The discrepancy is recorded because it bears directly on whether $\varphi = 0$ on this component, and a reader consulting the handbook alone would reach the opposite conclusion.

Uncertainty (iv): scope. Personal consumption expenditures on electricity cover household purchases only, roughly a third of US electricity sales, and are a use of the electricity commodity rather than the industry output entering the productivity residual. Equation (S.49) is corroborating evidence on the household component, not a statement about real gross output of NAICS 2211.

E.6 What the price indexes can and cannot register

Whichever construction applies, the price statistics available for US electric power have no hour-of-delivery dimension. The PPI price concept is a specified transaction: "the net revenue accruing to a specified producing establishment from a specified kind of buyer for a specified product shipped, or service provided, under specified transaction terms on a specified day of the month" (Bureau of Labor Statistics, 2026b, Concepts), collected as realistic transaction prices including discounts, premiums, rebates and allowances, with quality adjustment when the specification changes. The industry cells are defined by seller type and customer class: the generation index prices the initial commercial transaction of generating establishments, split into utility and non-utility cells; the bulk transmission index prices transmission fees and charges; the distribution indexes are published separately for residential, commercial, and industrial power, subdivided by Census division (Bureau of Labor Statistics, 2005). Aggregation holds transaction specifications and their weights fixed between weight revisions, so a change in the composition of sales across cells does not move the index; only a change in the price of a sampled specification does.

Three implications follow. First, a shift of generation into low-price hours alters no sampled specification's price and changes no product class, so it cannot register as a price movement: $\psi \simeq 0$ on the time-of-delivery margin. Second, a relative price movement across hours is invisible for the same reason — there is no hourly stratum whose price could move. Third, the capture ratio $v_{j,t}$ of Section 4.2 is a generation-weighted object, and nothing in the PPI is generation-weighted at sub-monthly frequency. One clarification closes a natural objection: a PPI quote for a regulated utility is locally a ratio of dollars to kilowatt-hours — revenue for a specified customer class and rate schedule at a specified usage level. But the ratio is taken *within a held-constant specification*, priced on a single day of the month, and is never re-weighted across hours or across the sales mix. A quote that is locally a unit price is not thereby an aggregate unit value, and it is the re-weighting, not the ratio form, that would carry the capture signal.

The deflator is not an average revenue per MWh. It is worth establishing separately that the published output price index is not a unit value, since a unit-value deflator would force real output to equal physical output identically and make the premise trivially true. It does not. Over 2017–2024 the published NAICS 2211 output price deflator rose 25.5 percent, while current-dollar output divided by the Federal Reserve physical-production index — the implied average revenue per unit — rose 17.9 percent, a gap of 7.6 percentage points in seven years; the correlation of their annual log changes is 0.79, not unity.[11] A unit value in this sense is published, but by a different agency and not as a deflator: EIA's average retail price is reported revenue divided by reported kilowatt-hours sold (U.S. Energy Information Administration, 2026a).[12]

The generation cut. Over 2017–2024 the composite generation PPI rose 26.4% and the merchant cell — the one official series pricing non-utility wholesale generation — rose 47.0 percent, both peaking in 2022 with the gas shock, while wind and solar generation rose 127.9 percent, all

[11] Author's calculation from BLS series `IPUCN2211T010000000` and `IPUCN2211T050000000`, with the physical benchmark from Federal Reserve G.17 `IPG2211A`. Over 1987–2024 the deflator rises from 65.3 to 125.5 and the implied unit value from 57.0 to 117.9 on a 2017=100 scale.

[12] Three statistics should not be conflated: the CPI for electricity, which prices selected rate schedules at consumption amounts drawn from Department of Energy information (Bureau of Labor Statistics, 2026a); the BLS average price per kilowatt-hour, a dollar level computed from the CPI utility sample; and EIA's average revenue per kilowatt-hour, the only one of the three mechanically equal to revenue divided by quantity.

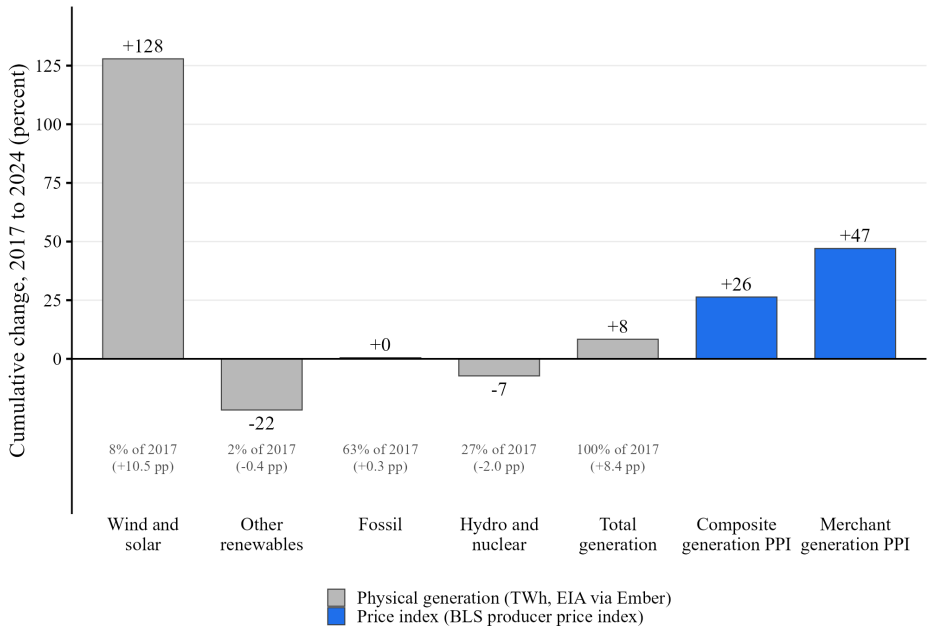


Figure E.1: Physical renewable generation roughly doubled while generation price indices rose (NAICS 22111, 2017–2024). Grey bars are physical generation in TWh (EIA, via Ember); the four component categories are mutually exclusive and exhaustive, so their contributions sum to the change in total generation, and the 2017 share and percentage-point contribution of each are printed below the axis. Blue bars are price indices (BLS producer price indices, monthly, annualized with February 2021 excluded, rebased 2017=100), which rose on fuel cycles peaking in 2022. The two are not commensurable: both are cumulative percent changes over the same window, but the juxtaposition is of direction, not of a common unit. The merchant cell prices wholesale rather than bundled retail output and is the natural home for a cannibalization signal. Neither PPI is the NAICS 2211 output deflator, and under a physical-quantity output construction neither deflates anything. Because the indices are monthly, matched-model and nationally aggregated, they resolve neither the intraday nor the regional structure on which $v(\theta)$ turns, and no capture-price decline is separately identified in them.

non-hydro renewables 98.3 percent, and total generation 8.4 percent (Figure E.1).[13] The movements are dominated by fuel prices. The correct reading is narrow: no secular decline attributable to capture prices is *separately identified* in these series. It would be an overstatement to conclude that the capture component is zero, since a negative capture contribution could be masked by fuel, contract, and congestion movements of larger magnitude; the 47% rise in the merchant cell does not by itself bound the capture contribution. What the generation cut establishes is that B_V cannot be recovered by differencing published price series, which is the output-side counterpart of the reason B_K cannot be read off a vintage difference (Supplemental Appendix B.1).

E.7 Settlement pass-through: a channel-based sensitivity analysis

Under the deflated-value construction, with $\psi = 0$ on the time-of-delivery margin, the residual in Equation (S.48) depends on φ, the share of the capture-value change that reaches recorded current-dollar sectoral output. This is not the usual question of how wholesale price changes feed through to retail tariffs (**??**): the two objects have different numerators and denominators, so I do not identify φ from retail-price regressions. Instead, let $G = \sum_c G_c$ be generation-related revenue, split by settlement channel c. Holding physical output and the reference price fixed,

$$\varphi = \sum_c \omega_c \varphi_c, \qquad \omega_c = \frac{G_c}{G}, \qquad \varphi_c = \frac{\partial \log G_c}{\partial \log C}. \tag{S.50}$$

The identity holds locally. The weights are shares of generation revenue, not of total sectoral revenue; the generation weight w_G converts one into the other. A weighted average of local elasticities is not, in general, a constant elasticity over a large transition. The weights are data-based; the channel elasticities φ_c are assumptions, which I state explicitly.

Channel weights. EIA Form 861 separates bundled sales (Part A), competitive energy-only sales (Part B), delivery-only service for the same load (Part C), and the bundled billing used by retail power marketers in Texas (Part D) (U.S. Energy Information Administration, 2026a). Revenues are observed; their generation component is partly allocated. Part C contains no generation revenue and leaves the base, as it does in w_G. Part B is already energy-only. For Parts A and D I apply the AEO generation share of a bundled bill. As a check, Part B revenue divided by Parts B and C revenue is 0.597, close to the AEO estimate of 0.612.

For Part A, I match the Form 861 operational records to utilities that report bundled retail sales in the same year; the matched records cover 97.8 percent of Part A sales. Purchases net of sales for resale average 0.350 of these utilities' supply over 2013–2024. Netting removes offsetting volumes: in 2024 the gross and net purchase shares are 0.50 and 0.37. It does not necessarily remove offsetting value. A utility that sells one MWh in a low-price hour and buys one in a high-price hour has zero net purchases but nonzero net expenditure. The net share is therefore a quantity proxy for price exposure, not a bound on it.

Central calibration. FERC Form 1 Schedule 326 classifies purchases by duration and firmness. Short-term firm and other, including non-firm, service account for 0.625 of classified MWh and

[13] Generation and technology splits: US generation by fuel in TWh, EIA data retrieved through Ember's yearly full release, which is exhaustive and additive across fuels and therefore allows the component contributions to be checked against the total. Generation PPIs: BLS `PCU221110221110` (composite), `PCU2211102211104` (utilities), `PCU2211102211105` (non-utilities); monthly, annualized, rebased 2017=100.

0.571 of settlement value (**??**).[14] These codes do not disclose whether a contract price is fixed or indexed. For regulated bundled service I set

$$\varphi_{\text{regulated}} = \underbrace{\frac{\text{net purchased MWh}}{\text{supply}}}_{\text{quantity share}} \times \underbrace{\text{short/non-firm share of purchases}}_{\text{exposure proxy}},$$

with $\varphi_c = 0$ on own generation. This is a calibration, not an accounting identity: it also requires that exposed purchase costs move one for one with the capture-value change and reach tariffs within the year. Competitive supply is assigned $\varphi_c = 1$, the pass-through-maximising value; fixed retail margins, hedges, and capacity charges would lower it. The resulting value-weighted average is $\varphi = 0.386$ on the sectoral-output concept, comprising 0.172 from regulated bundled service and 0.214 from competitive supply.

High-pass-through scenario. Treating every net purchase as exposed and taking the generation share of a bundled bill from Parts B and C gives a value-weighted mean of $\varphi = 0.492$ over 2013–2024. This is a conditional scenario, not an identified upper bound on pass-through. The annual values rise from 0.458 in 2013 to 0.523 in 2024, so the window mean is exceeded in recent years. The scenario also rests on the restrictions $0 \leq \varphi_c \leq 1$, on own generation entering recorded output through cost-of-service revenue rather than wholesale transfer prices, and on the net-MWh exposure proxy; counting gross rather than net purchases as exposed raises the mean to 0.596. Form 861 identifies none of these, so the channel data do not furnish an unconditional upper bound on φ.

Omitted channels. Three simplifications all raise the calibrated φ. Corporate power purchase agreements and direct-access contracts are out-of-sector sales at a fixed strike, so $\varphi_c = 0$; they are omitted. Cooperatives and municipal utilities buy mainly on requirements contracts from generation-and-transmission cooperatives and the federal power marketing administrations, yet the investor-owned purchase split from FERC Form 1 is applied to them. Community choice aggregators procure on long-dated contracts and set rates administratively, yet are held at $\varphi_c = 1$.

Adding the positive gap between BEA gross output and EIA retail revenue at $\varphi_c = 1$ raises the gross-output analogue of the central value to 0.429. It does not change the surviving wedge. Because w_G measures retail-embedded generation revenue over gross output, G/Y, adding a fully exposed layer W requires the weight $(G + W)/Y$, and

$$\frac{G + W}{Y}(1 - \varphi_{\text{GO}}) = \frac{G}{Y}(1 - \varphi).$$

Only the sectoral-concept φ is therefore combined with w_G below.

The 9.7 percent of customers enrolled in any time-based rate in 2024 is a useful diagnostic of direct hourly indexation, but it does not bound φ: ordinary tariffs can transmit annual wholesale-cost changes through fuel clauses, default-service resets, or rate cases. It is therefore excluded from the calculation.

Implication. Under the physical-quantity construction described in Section E.3, no additional settlement-pass-through adjustment applies to the within-class timing margin. Under the deflated-value alternative, maintaining $\psi = 0$, the sector wedge is $(1 - \varphi)\, w_G B_V^{\text{gen}}$. The central calibration

[14] PUDL's post-2020 XBRL records include blank-seller respondent totals, identified by record identifiers ending in `_total`. I remove them before aggregation and exclude the remaining unclassified detail from the denominator of the proxy.

Table E.3: Settlement pass-through φ: measured channel weights and the sensitivity of the bound. Every weight is read off EIA Form 861 or FERC Form 1; only the channel pass-throughs φ_c are assumed, and each is assumed in the direction that raises φ and so shrinks the residual wedge $(1-\varphi)B_V^{\text{gen}}$. Window 2013–2024.

Specification	$\bar{\varphi}$	φ_{2024}	$1-\bar{\varphi}$
Baseline, sectoral output concept	0.305	0.308	0.695
BEA gross-output concept	0.347	0.308	0.653
Generation share from Part B/C	0.309	0.321	0.691
Purchase split, dollar-weighted	0.297	0.305	0.703
All utility purchases market-priced	0.360	0.369	0.640
Contract purchases 25% indexed	0.319	0.323	0.681
Competitive retail $\varphi_c = 0.75$	0.252	0.251	0.748
Competitive retail $\varphi_c = 0.50$	0.199	0.194	0.801
Aggregator load at $\varphi_c = 0$	0.294	0.280	0.706
Intra-industry layer at the EIA-861 total	0.545	0.519	0.455
Adversarial: all three together	0.398	0.369	0.602

Notes: The regulated channel is exposed only through fuel and purchased-power adjustment clauses, so its pass-through is the product of the net purchased share of supply (0.182 in 2024) and the market-priced share of purchases from the FERC Form 1 statistical classification (0.623, megawatt-hour weighted, 2013–2024 mean). Competitive supply is held at $\varphi = 1$ throughout, which is the correct value under full pass-through, not merely a conservative one. Delivery-only revenue carries no generation value and is excluded from the base. The gross-output rows add the intra-industry resale layer, which BEA includes and the BLS sectoral-output concept excludes, at $\varphi = 1$.

reduces the effective generation weight from 0.585 to 0.359, and the high-pass-through scenario to 0.297. Greater pass-through shrinks the residual output wedge and so delays the point at which the output wedge overtakes the capital wedge; it does not reverse the sign of the residual.

E.8 The European framework: one concept, two A-methods, national implementation

The European framework is examined at three levels: the accounting standard, the harmonised methodology, and the national instruments that implement it.

The standard endorses the concept. ESA 2010 classifies time-of-day variation — "electricity supplied at night compared with supply during the day" — as a quality margin to be recorded as a change in products rather than in price (ESA 2010, §10.16). An output account implementing that provision would carry each hour of delivery as a distinct product and would register the shift of generation toward low-value hours directly, exactly as Q^{val} does. The concept in this paper and the concept in the standard therefore coincide: the bias is a granularity gap in the instruments implementing the standard, not a conceptual error in the standard itself. The treatment is, however, permissive rather than required: the handbook's best-practice condition for a quantity-based volume measure is only that quantities be identified by tariff and customer type, in which hour of delivery does not appear.

The harmonised methodology permits two A-methods. For electricity (CPA 35), the Eurostat handbook on price and volume measures identifies as A-methods both the deflation of output values by a suitable producer price index and the use of quantity indicators *sufficiently detailed*

by tariff and customer type, which it regards as equivalent (Eurostat, 2016, CPA 35); value added is in principle obtained by double deflation, with separate volume measurement of output and intermediate consumption (ESA 2010, §10.32). The two branches map directly onto the families of Section E.1: under the direct branch there is no numerator and no deflator, and the premise follows from the quantity cells; under the deflated branch it follows from the accounting condition of Section E.2. In neither branch does the harmonised requirement introduce an hour or location stratum: the prescribed detail is tariff and customer type.

The price instruments are harmonised in concept, national in construction. The legal definition of the industrial producer-price variable requires an actual transaction price at basic prices, specified by its price-determining characteristics and adjusted for quality change, compiled monthly for NACE Section D (European Parliament and Council of the European Union, 2019; European Commission, 2020); the mandatory harmonised breakdown is generally at the electricity-gas-steam division level rather than electricity generation specifically. Within that frame, the elementary price concept, sampling, and aggregation are set nationally and differ materially: Finland's metadata state that "the price data collected from enterprises are quantity weighted average prices," resorting to "the price of the delivery/billing/payment time on the 15th day of the month" only where an average cannot be determined (Statistics Finland, 2026); Spain's describe prices observed around the 15th of the month with geometric-mean elementary aggregation; Denmark's describe geometric means of matched observations within detailed commodity aggregates, following the 2012 PPI handbook (Eurostat, 2026b, 2012). None of the examined metadata documents an hour-by-bidding-zone stratification. Two conclusions follow, one firm and one deliberately bounded. Firmly: no published harmonised requirement or examined national metadata stratifies the electricity PPI by hour and location of delivery, so no European deflator is documented to price the capture margin as a product dimension. Bounded: the compilation manual asks not for a price on a nominated day but for a period average, and singles out volatile products of macroeconomic importance.[15] Because some national constructions admit quantity-weighted or unit-value-type elementary prices, part of the within-month settlement variation may enter some national indexes, so $\psi \simeq 0$ is a country-by-country hypothesis for Europe rather than a proven uniform property. This bound does not weaken the premise: by Table E.2, $\psi > 0$ with low pass-through φ *raises* the residual bias above B_V^{gen}.

Translation into value-added productivity. European and UK multifactor productivity is computed on real value added obtained by double deflation, not on gross electricity output (Eurostat, 2016; Office for National Statistics, 2025). Let Y be gross output, M intermediate consumption, and $GVA = Y - M$. If only the output volume is mismeasured,

$$d\log GVA^{\text{off}} - d\log GVA^{*} \simeq \frac{Y}{GVA}\left(d\log Y^{\text{off}} - d\log Y^{*}\right),$$

[15] Eurostat (2021, pp. 34–35): "The index should in principle reflect the average price during the reference period. In practice the information actually collected may refer to a particular day in the middle of the reference period ... For products with a significant impact on the national economy that are known to have, at least occasionally, a volatile price development, it is important that the index does indeed reflect average prices." Electricity satisfies both conditions. If a respondent computes that average as revenue divided by quantity over the month, the observation is a within-month unit value and absorbs the intraday component of the capture-price decline, which is the larger component since the margin is primarily diurnal. If it is the average price of a fixed specification, hour-composition is not a specification characteristic and does not enter. The manual does not resolve which practice applies to wholesale electricity, and respondent-level collection procedures are not public.

so a gross-output volume bias is amplified in value-added growth by the ratio Y/GVA, holding intermediate-input volumes fixed. Measured on the same NAICS 2211 accounts used in Table 1, that ratio runs from 1.38 to 1.78 over 2010–2024 with a mean of 1.58, so the amplification is roughly one and a half rather than the two-to-two-and-a-half typical of more fuel-intensive definitions of the industry.[16] Holding intermediate-input volumes fixed. Purchased electricity, fuel, and balancing services on the input side may partially offset this, and the net value-added bias is $B_{GVA} \simeq (Y/GVA)B_Y - (M/GVA)B_M$. The gross-output basis used in the main text is therefore the conservative one for comparison with value-added productivity statistics.

E.9 Two direct-quantity systems: the United Kingdom and Norway

The direct branch of the European methodology is not hypothetical. Two systems with unusually rich underlying energy data implement it, and both illustrate that data granularity does not translate into index granularity.

United Kingdom. For the Index of Production, the ONS reports that Section D — electricity, gas, steam and air-conditioning supply — is measured from energy-supply volume data supplied by the Department for Energy Security and Net Zero, classified as direct-volume sources rather than Monthly Business Survey turnover deflated by a price index (Office for National Statistics, 2024). There is accordingly no current-value numerator and no deflator in the short-term electricity volume: the incidence parameters of Section E.2 do not arise, and the operative question is the stratification of the supplied volume data. The published methodology identifies energy-supply categories, not hour-location cells, so within each supplied stratum the index carries equal weight per unit of energy; the public documentation does not disclose the low-level aggregation cells, so the statement supported is that the construction does not mechanically incorporate capture prices, not that its residual bias is exactly B_V^{gen}. ONS multifactor productivity then uses quarterly-national-accounts value added as its output concept (Office for National Statistics, 2025), so the translation of the preceding subsection applies.

Norway. Statistics Norway extrapolates electricity-production volume from monthly electricity statistics — essentially physical generation — and inflates it with weighted price indexes for major recipient groups (households, ordinary industries, energy-intensive manufacturing), drawn from a quarterly survey of average prices and quantities by contract type and consumer group (Bougroug, 2021; Statistics Norway, 2026a). The construction is thus a quantity extrapolation on the output side: a wind MWh generated in a low-price, wind-abundant hour receives the same volume weight as a hydro MWh generated in a scarce evening hour. Norway is instructive because the raw data to do otherwise exist and are partially used: since October 2025 SSB values cross-border electricity flows at 15-minute Nord Pool prices in the external-trade statistics (Statistics Norway, 2026b). High-frequency *valuation* of the nominal numerator, however, is not high-frequency *volume aggregation*: if the same interval prices are used to deflate the interval-valued revenue, real output is again physical MWh, with the interval variation assigned entirely to the price component. Eliminating the bias requires carrying the interval-zone cells into the volume index as products — the treatment ESA 2010 §10.16 already endorses — and no documented Norwegian national-accounts series does so. On the capital side, SSB deflates electricity-industry investment by asset and product with

[16] Nominal gross output over nominal value added, BEA industry accounts for NAICS 2211; see `Results/provenance_gap_closure.md`.

separate domestic and import price components, which mitigates B_K without a renewable capacity-cost index; the Norwegian system therefore partially addresses the capital bias while leaving the output-value bias uncorrected, consistent with the timing asymmetry of Section 2.

E.10 What would overturn the premise

The premise fails under any of the following, and each is checkable.

1. *Hour-differentiated product classes.* If BLS or a European statistical institute were to stratify electricity output by time and location of delivery — whether as product classes for deflation or as products in a quantity aggregation — the resulting index would be value-weighted on the relevant margin and the bias would vanish by construction. This is the treatment ESA 2010 §10.16 already endorses in principle.

2. *A price cell settled on generation-weighted revenue.* If a PPI cell (in the US, the merchant generation quote; in Europe, a national elementary aggregate) were collected as revenue divided by output over a settlement period, with weights updated within the period, the cell would move with the capture price and ψ would rise toward one. The US published price concept — a specified transaction on a specified day — indicates otherwise, but the sampling detail for wholesale transactions is not public, and several European national metadata admit quantity-weighted monthly averages; this is the single most consequential unverified fact in this appendix. Note that by Table E.2, $\psi > 0$ with low φ enlarges rather than removes the residual bias; only $\psi \simeq \varphi \simeq 1$ would eliminate it through the price side.

3. *High settlement pass-through under deflated-value construction.* If output is built by deflating revenues and φ is near one while $\psi \simeq 0$, then by Equation (S.48) the published series already registers the capture-value loss, so the underlying wedge does not generate bias in official statistics. This combination is the substantive form of the objection, and it is why Section E.7 calibrates φ by settlement channel rather than asserting it: the central calibration is 0.386, and none of the reported scenarios approaches one. It cannot arise in the direct-quantity and extrapolation constructions of Sections E.3, E.5, and E.9, where recorded revenue does not enter the volume index.

E.11 Summary

Five constructions of official real electricity output are documented above, and the premise of Proposition 2 holds under each, for different reasons and with different strength.

In no documented case does the volume aggregation resolve the hour or location of delivery. That, and not any property of an output deflator taken in isolation, is what Proposition 2 requires. The capture-price decline is therefore not recoverable from published annual output aggregates and must be measured directly from hourly generation and prices, as in Section 4.2.

F Measuring the output-value bias in market data

This appendix documents the market-data measurement behind Section 4.2 and Figure 2. Section F.1 defines the price concept and the value factors and describes the European panel; Section F.2 describes the US market construction. Section F.3 reports value-factor schedules estimated from own-zone variation, as a diagnostic on the cross-border schedules adopted in the calibration. Section F.4 converts the level gap into annual growth-bias averages and adjusts for common price

Table E.4: Constructions of official real electricity output and the status of the equal-value-per-MWh premise.

Construction	Aggregation strata	Residual bias
BLS physical-quantity index (kWh sold to ultimate customers)	class-of-service prices; no hour or location	B_V^{gen}, unconditional
BLS deflated value (product-class deflation)	product classes = customer classes; no hour or location	$(1-\varphi)B_V^{\text{gen}}$; φ calibrated, not identified
BEA PCE quantity extrapolation (EIA kWh carried by CPI)	none; real output ≡ physical kWh	B_V^{gen}, exactly, on the household component
European deflated value (PPI A-method)	national elementary cells; no documented hour-location strata	$(1+\psi-\varphi)B_V^{\text{gen}}$; $\psi \simeq 0$ a national hypothesis
UK / Norway direct volume and extrapolation	energy categories, recipient groups; no hour or location	B_V^{gen} up to undisclosed low-level cells

shocks. Section F.5 shows that the European level evidence is robust to the treatment of negative prices and curtailment. Sections F.6 and F.7 construct a narrower chained-volume benchmark under alternative delivery-time stratifications and examine its sensitivity to the 2021–2023 price shock.

F.1 Price concept, value factors, and the European panel

Price concept. The price entering the capture ratio is the zonal day-ahead price, not the generator's total remuneration. Feed-in tariffs, contracts for difference, power-purchase agreements, and other contractual transfers may insulate receipts from the hourly market price; using those receipts would therefore combine the value of delivered energy with the support or risk-sharing arrangement attached to the plant. The day-ahead price instead measures the market value of an additional MWh in a given hour and zone under the prevailing network constraints and market rules. Negative prices are retained. This is a market-value benchmark, not a measure of generator profitability or social welfare.

Value factor. Following Hirth (2013), I define *value factor* as the generation-weighted average price a technology earns, divided by the time-weighted average price of the zone. For technology j in zone z and year t, over the settlement hours h of that year,

$$v_{j,z,t} = \frac{\bar{p}^{\text{cap}}_{j,z,t}}{\bar{p}_{z,t}} = \frac{\sum_h p_{z,t,h}\, g_{j,z,t,h} / \sum_h g_{j,z,t,h}}{\sum_h p_{z,t,h} / H_{z,t}}, \tag{S.51}$$

where $p_{z,t,h}$ is the day-ahead price, $g_{j,z,t,h}$ the generation of technology j, and $H_{z,t}$ the number of settlement hours. The numerator $\bar{p}^{\text{cap}}_{j,z,t}$ is the technology's *capture price* and the denominator $\bar{p}_{z,t}$ the *baseload price*; $v_{j,z,t}$ is a unit-free ratio, so no deflation enters. A value factor of one means the technology earns the average price, and $v_{j,z,t} < 1$ is the cannibalization discount that arises when output is concentrated in low-price hours. I compute Equation (S.51) directly from ENTSO-E hourly day-ahead prices and generation for 15 European bidding zones over 2015–2024; own-technology penetration $\theta_{j,z,t}$ is generation over load on the same panel.

F.2 US market construction

The underlying observations are published by CAISO and ERCOT; the gridstatus.io hosted API is used as a programmatic access and harmonization layer. CAISO day-ahead prices are drawn from the ISO's Open Access Same-Time Information System (OASIS), while generation and load are based on CAISO's official production and system-demand series (California Independent System Operator, 2026a,b,c). ERCOT day-ahead prices are drawn from the DAM Settlement Point Prices product, generation from the Fuel Mix reports, and load from the system-load archives (Electric Reliability Council of Texas, 2026a,b,c). The series are retrieved through the gridstatus.io hosted API (Grid Status, 2026).

I apply the same hourly accounting construction as in the ENTSO-E panel. For each technology j, the capture price is the generation-weighted day-ahead price, the baseload price is the time-weighted day-ahead price, and penetration is generation divided by load. Negative prices are retained. Negative generation observations, where present, are set to zero before constructing the generation weights. The resulting sample covers CAISO from 2020 to 2025 and ERCOT from 2018 to 2025, corresponding to the years for which the required price, generation, and load series are jointly available.

Because CAISO does not publish a single day-ahead system price, I construct the reference price as a load-share-weighted average of the PG&E, SCE, and SDG&E default load-aggregation-point prices, using approximate weights of 0.45, 0.43, and 0.12. These weights are calculated from their average shares of system load over the sample. The resulting benchmark matches the system-wide aggregation of generation and load used in the calculation. Pricing all CAISO generation at a single node instead would change the result substantially, and in either direction: in 2024 the solar value factor is 0.40 against the SCE load-aggregation point and 0.74 against PG&E, a spread of 0.34, compared with 0.58 on the load-weighted average.[17] The dispersion is a solar phenomenon and it supports the system-wide benchmark rather than treating any one regional node as the price of all CAISO generation.

For ERCOT, I use the ERCOT-wide hub average `HB_HUBAVG`, reported among the market's hub settlement points (Electric Reliability Council of Texas, 2026a,d). Generation and load are system-wide balancing-area totals. For the categories shown in Figure 2, a market-year is classified as solar-dominated when solar accounts for at least 60% of wind-plus-solar generation,

$$\frac{\theta_s}{\theta_s + \theta_w} \geq 0.6,$$

wind-dominated when this share is below 0.4, and mixed otherwise. Marker shape identifies the market-data source: circles denote ENTSO-E bidding zones, triangles CAISO, and squares ERCOT.

In 2024, CAISO records an output-value gap of 11.7% at 38.4% wind-and-solar penetration (6.5% in whole-sector units); ERCOT records 5.3% (2.9% in sector units). The elevated ERCOT observation in 2021 reflects the February winter-storm price spike.

F.3 Within-variation in the value-factor schedules

This subsection asks whether own-zone variation produces the same qualitative relationship between penetration and value factors. It is not the source of the schedule used in the calibration. The domestic specification includes own-zone penetration but not neighbouring penetration, whereas

[17] Computed from the hourly day-ahead prices and hourly solar generation in `Data/us_hourly_cache`; see `Results/provenance_gap_closure.md`. The corresponding wind value factors are tightly clustered, 0.96 at every node, because wind output is not concentrated in the hours when nodal prices diverge.

cannibalization in interconnected markets may operate through generation elsewhere in the market. It therefore does not identify the full cross-border effect estimated by Stiewe et al. (2025).

For each technology $j \in \{s, w\}$ I estimate the value-factor schedule on the panel of European bidding-zone-years (ENTSO-E, 2015–2024; zone-years with own-technology penetration above 0.5 percent),

$$v_{j,z,t} = v_{0,j} - a_j\, \theta_{j,z,t} + \mu_z + \lambda_t + \varepsilon_{j,z,t}, \qquad j \in \{s, w\}, \tag{S.52}$$

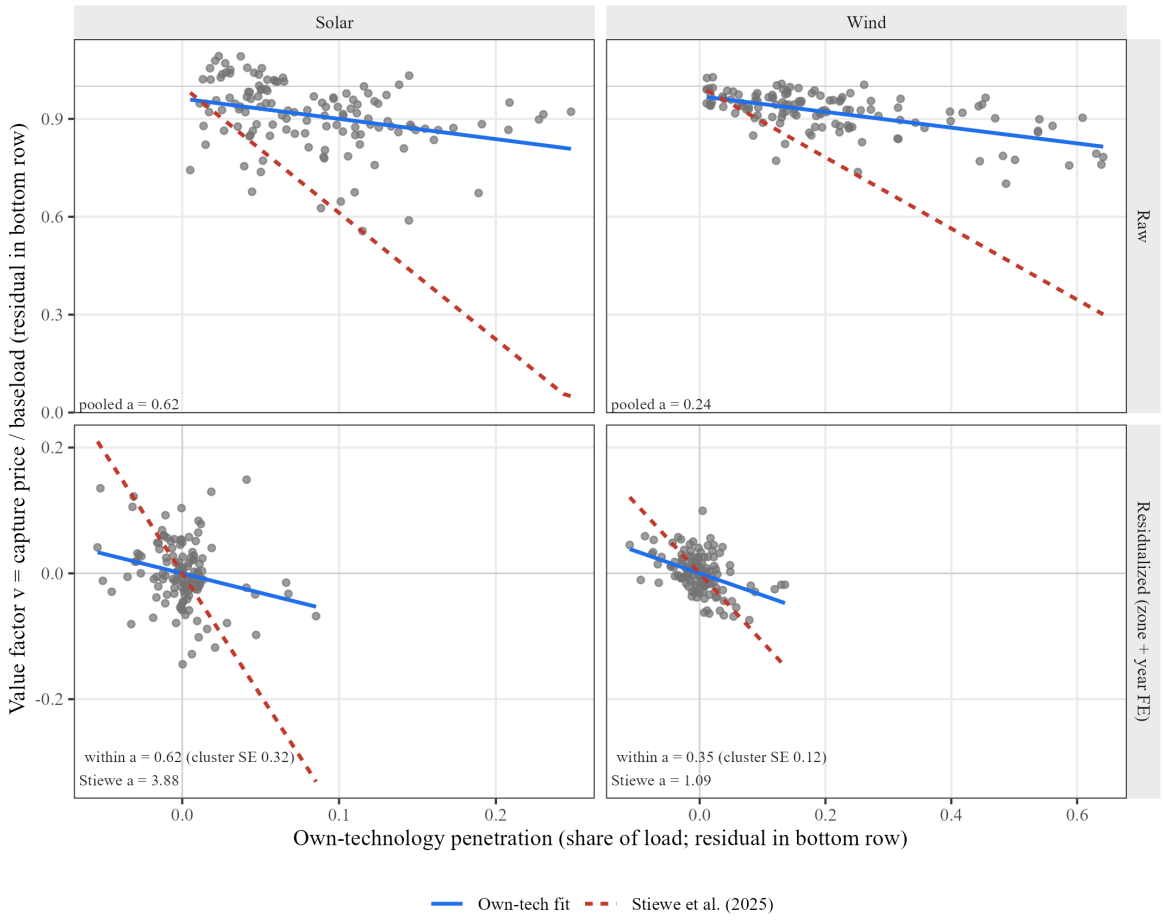
estimated separately by technology, where $v_{j,z,t}$ is the generation-weighted capture price over the time-weighted baseload price, $\theta_{j,z,t}$ is own-technology generation over load, and μ_z, λ_t are bidding-zone and year fixed effects. The schedule slope $a_j = -\partial v_j / \partial \theta_j$ is the cannibalization rate, signed so that a positive a_j means the value factor falls as own penetration rises, matching the $v = v_0 - a\theta$ parameterization of Hirth (2013) and Stiewe et al. (2025). Standard errors are clustered by bidding zone (CR1). Table F.1 reports the estimates.

Table F.1: Value-factor schedule slopes, European bidding-zone panel (ENTSO-E, 2015–2024). The slope $a_j = -\partial v_j / \partial \theta_j$ is the cannibalization rate (positive: value factor falls with own penetration). "Pooled" omits fixed effects; "Within" is the two-way (zone + year) fixed-effects estimator of Equation (S.52). Clustered standard errors (by bidding zone, CR1) in parentheses; with G clusters, t-statistics should be assessed against $G-1$ degrees of freedom. The Stiewe et al. (2025) row is the cross-border spatial estimate (domestic + neighbor) adopted in the calibration.

	Solar	Wind
Pooled slope a_j (no FE)	0.62	0.24
Within slope a_j (two-way FE)	0.62	0.35
(clustered SE)	(0.32)	(0.12)
t-statistic	1.96	2.90
Zone-years N	129	128
Bidding zones	15	14
Stiewe et al. (2025) spatial a_j	3.88	1.09

Figure F.1 displays the same estimates. Its two columns are the two technologies (solar, wind) and its two rows are two views of the identical data. The *top row* plots the raw value factor against own-technology penetration; the solid line is the pooled slope, which confounds the within-system relationship the mechanism concerns with fixed cross-system differences (a structurally sunny, high-solar-value zone need not resemble a windy one). The *bottom row* is the Frisch–Waugh–Lovell partialling-out plot: the value factor and penetration are each first regressed on the zone and year dummies, and the two sets of residuals are plotted against one another. This removes everything attributable to *which zone* an observation is and *which year* it falls in — fixed locational differences and common annual shocks such as gas-price cycles and weather years — leaving only the within-zone, within-year covariation that identifies a_j. By the Frisch–Waugh–Lovell theorem the slope of this residual scatter is exactly the two-way fixed-effects coefficient of Equation (S.52), so the bottom row is a direct picture of the identifying variation rather than the confounded pooled cloud above it. The red dashed line in each panel is the Stiewe et al. (2025) cross-border slope, drawn on the same axes for comparison.

The within slopes are positive and far below the spatial estimates (solar 0.62 against 3.88; wind 0.35 against 1.09). The wind slope is significantly positive ($t = 2.90$ against a 5% critical value of 2.160 at $G - 1 = 13$ degrees of freedom); the solar slope is marginal ($t = 1.96$ against 2.145



Bidding-zone-years, ENTSO-E 2015-2024. Top row: raw value factor vs own-tech penetration (confounded pooled slope). Bottom row: both variables residualized on zone + year fixed effects (Frisch-Waugh-Lovell) -- the within-variation the mechanism needs; slope = two-way FE estimator, standard error clustered by bidding zone (CR1). Red dashed: Stiewe et al. (2025) cross-border own-tech slope.

Figure F.1: Value factors against own-technology penetration, raw (top row) and residualized on zone and year fixed effects (bottom row), for solar (left) and wind (right). Solid line: own-technology fit (pooled slope in the top row, two-way within slope in the bottom); red dashed line: Stiewe et al. (2025) cross-border spatial slope. Standard errors clustered by bidding zone. Specification and point estimates in Equation (S.52) and Table F.1.

at $G - 1 = 14$), as one would expect with fifteen clusters. This is the pattern the design predicts rather than a contradiction: differencing out zone and year effects removes the neighbor-penetration spillovers that the cross-border spatial design exploits, and cannibalization operates through the interconnected market price, not the domestic generation share alone. The domestic panel therefore recovers only the own-market component of the schedule, which is why the calibration adopts the spatial estimates. What the exhibit establishes is narrower and sufficient for the argument here: within the same European data the sign is right and, for wind, statistically distinguishable from zero, and the pooled and within slopes are of the same modest order—so the object placed inside productivity accounting is a schedule these data are consistent with, not one re-estimated from within-country variation.

F.4 Annual growth-bias averages and common price shocks

The generation growth bias, $B_V^{\text{gen}} = -\Delta \log(1-D)$, averages 0.53 percentage points per year across the European panel. I compute the bias from the observed year-to-year change in $D = \theta(1-v)$ for each zone and then average across zone-year observations. The average is concentrated after 2019, when common price shocks—COVID, the 2022 gas-price spike, and the subsequent normalisation—moved value factors in all zones simultaneously. Removing the common annual movement, by subtracting each year's cross-zone mean deviation from every zone's value factor, lowers the average to 0.12 percentage points per year, with the pre-2020 bias values close to zero. To express the bias as a contribution to whole-sector TFP growth, I weight it by generation's share of sector output, $B_{V,t} = w_G B_{V,t}^{\text{gen}}$ with $w_G \simeq 0.55$ the maintained generation share for NAICS 2211, which also includes transmission and distribution. This exercise is descriptive: it is not used to calibrate the forward value-factor schedules of Appendix G. Unlike the capital-side correction, which also reinterprets past measured input growth, the calibrated output-side bias arises mainly as renewable penetration increases going forward.

F.5 Negative prices and curtailment

Two checks address the concern that the measured gap reflects market-design artifacts rather than value. Flooring day-ahead prices at zero, which deletes the contribution of negative-price hours, leaves the West Denmark 2024 generation-basis gap essentially unchanged (18.5 to 18.4 percent) and lowers the mean over the four-zone robustness subset (DK1, DK2, DE-LU, IE, 2019–2024) from 8.0 to 7.9 percent: the level gap is not a negative-price artifact, because negative hours in these zones are shallow relative to the mean price. Including curtailed energy, which the transparency-platform generation series excludes, works in the opposite direction from the concern: adding all-island Irish wind dispatch-down back into penetration raises the Irish gap by 0.3–0.5 percentage points (2020: θ 0.449 to 0.501). German curtailment is not consistently measurable across the October 2021 Redispatch 2.0 reform and is left out; Danish curtailment is negligible.

F.6 A narrower chained-volume benchmark

One issue concerns Q^{val} rather than official practice. Equation (4) evaluates output at contemporaneous relative values, whereas a chained volume index excludes pure relative-price movements. The two channels of Equation (5) therefore have different index-number interpretations. For each technology j, I decompose in a price and quantity terms,

$$\Delta[\theta_j(1-v_j)] = (1-\bar{v}_j)\,\Delta\theta_j - \bar{\theta}_j\,\Delta v_j.$$

where bars denote arithmetic means across the two periods. Let $\mathcal{L}_t$ denote the logarithmic mean of $1-D_{t-1}$ and $1-D_t$. Then

$$B_{V,t}^{\text{gen}} = \frac{\sum_j \left[(1-\bar{v}_{j,t})\Delta\theta_{j,t} - \bar{\theta}_{j,t}\Delta v_{j,t}\right]}{\mathcal{L}_t},$$

with $\mathcal{L}$ the logarithmic mean,

$$\mathcal{L}_t \equiv \frac{(1-D_t)-(1-D_{t-1})}{\log(1-D_t)-\log(1-D_{t-1})}.$$

which gives an exact additive decomposition of the log bias. The first term is a quantity-composition effect; the second captures changes in relative value, including pure relative-price movements that

a volume index records as prices rather than quantities. The composition channel accounts for 34% of the cumulative European bias and 41% of the US bias. Retaining that channel alone moves the central crossover from $\theta^\star = 0.34$ to 0.56; under the wind-heavy schedule, no crossover occurs within the modelled range.

Alternative delivery-time stratifications. Table F.2 constructs the narrower quantity-composition measure directly. Each specification applies the same chained Törnqvist formula to the same hourly prices and generation, varying only the aggregation strata. With month-by-hour cells, the cumulative gap is 6.2 log points in solar-dominated CAISO but slightly negative in wind-dominated ERCOT. With residual-load deciles, it is 6.9 log points in CAISO and 7.2 in ERCOT. Across Europe, the median rises from 0.25 log points under clock-time cells to 0.76 under residual-load deciles, with positive residual-load estimates in 13 of 15 zones.

The contrast is consistent with the technologies' temporal profiles. Solar output is strongly patterned by month and hour, so clock-time cells capture much of its reallocation across delivery periods. Wind varies substantially across days within the same clock-time cell. Residual-load strata distinguish those system states and can therefore reveal a quantity-composition margin that clock-time aggregation misses.

The levels are less robust than the ordering across stratifications. Dropping the 2021–2022 and 2022–2023 transitions, and then also 2020–2021, leaves the European residual-load median positive at 0.89 and 0.66 log points and positive in 12 and 10 of 15 zones. Clock-time estimates remain below residual-load estimates on the European median and in ERCOT, but the European clock-time median and pooled estimate turn negative. Individual zones can also be sensitive to short samples and exceptional price movements: Sicily, for example, changes by 4.5 log points when six earlier transitions are added, largely because of one year in which its baseload price nearly tripled.

These results therefore establish that stratification materially affects whether the quantity-composition gap is recorded; they do not identify a stable European magnitude. Residual-load cells are themselves a diagnostic construction, not uniquely correct output products, and the two system-state variants disagree in sign for IT-SICI and IT-NORD. IT-SUD is negative under every specification and is not explained here. The magnitudes used in the main analysis consequently come from the full current-value measure and its exact composition decomposition, while Table F.2 is interpreted as a test of the aggregation mechanism.

The exercise also clarifies the premise of Section 4.1. A sufficiently granular volume index is capable of registering delivery-time composition. The constraint is not index-number method but the aggregation strata used in official accounts, which distinguish customer classes, tariffs, or destinations rather than hour or system state.

The sensitivity of these comparisons to the 2021–2023 price shock is examined in Section F.7.

F.7 Sensitivity to the 2021–2023 price shock

Because T"ornqvist weights depend on cell revenue shares, exceptional changes in the hourly price distribution may affect the estimated gaps. Table F.3 therefore repeats the European comparison after excluding the 2021–2022 and 2022–2023 transitions, and then also the 2020–2021 transition.

The ordering across stratifications survives the exclusions. The European residual-load median remains positive and above the clock-time median in every sample, while ERCOT continues to show a substantially larger gap under residual-load strata. This supports the mechanism that clock-time cells can miss wind variation across days within the same hour-of-delivery cell.

The level estimates are not robust. Once the shock transitions are removed, the European clock-time median and pooled index turn negative, and the number of positive clock-time estimates falls

sharply. The chained-index exercise therefore shows that stratification materially affects whether the quantity-composition margin is recorded; it does not identify a stable European magnitude. The restricted samples are also shorter and exclude years of rapid renewable deployment, so the full panel remains the principal specification.

Taken together, the results show that a sufficiently granular volume index can register delivery-time composition, although the estimated magnitude is sensitive to the chosen strata and sample period. The constraint in official measurement is therefore not index-number method itself, but aggregation strata that distinguish customer classes, tariffs, or destinations rather than hour or system state.

G Calibration of the crossover

This appendix documents the calibration behind Figure 3. The exercise compares the capital-price bias and the output-value bias over a common path of wind-and-solar penetration. Both biases are expressed as annual contributions to whole-sector TFP. The object of interest is the penetration $\theta^\star$ satisfying

$$B_K(\theta^\star) = B_V(\theta^\star). \tag{S.53}$$

The exercise identifies a penetration threshold under alternative calibrations; it does not forecast the date at which any system reaches that threshold.

Mapping annual biases into penetration. Both objects in the crossover exercise are annual growth-rate biases:

$$B_{K,t} = s_K\, \omega_{R,t}\, \gamma_{K,t}, \qquad B^{\text{gen}}_{V,t} = -\Delta_t \log(1 - D).$$

To express them against a common penetration axis, I evaluate both along the logistic path

$$\theta_t = \frac{\theta_{\max}}{1 + \exp[-r(t - t_0)]}.$$

Conditional on this ceiling, the growth rate r and midpoint t_0 are chosen so that the path passes exactly through the EU wind-and-solar shares of 0.05 in 2010 and 0.29 in 2024. This gives $r = 0.149$ per year and $t_0 = 2029.0$. The path supplies the penetration level associated with each annual capital bias and, for the output-value bias, the annual change from θ_{t-1} to θ_t. The reported crossover $\theta^\star$ is therefore conditional on the assumed saturation ceiling and the resulting EU deployment path.

The forward capital bias uses the contemporaneous form, with the price gap $\gamma_{K,t}$ in place of the accumulated-stock gap $\kappa_{R,t}$ of Section A. Over 2011–2024 the two differ substantially: the price gaps sum to 119.3 log points and the stock gaps to 70.3. Along the forward path, earlier divergence can keep $\kappa_{R,t}$ elevated as $\gamma_{K,t}$ narrows, but whether it exceeds $\gamma_{K,t}$ by the crossover depends on the investment path and stock turnover. A vintage-consistent forward calculation is not attempted here.

I impose a saturation ceiling of $\theta_{\max} = 0.90$ following Jakhmola et al. (2026), it assumes an early and late acceleration scenario, where the ceiling L is respectively 92% (45% for wind, and 47% for solar), and 94% assumes L = 45% (47% for wind, and 47% for solar). For sensitivity analysis, I test whether a 80% ceiling for renewables would alter the result and find that the cross-over penetration rate θ^* is robust to the ceiling assumption (Table G.1).

But the crossover asks what happens at penetrations not yet observed, hence it requires a relationship between each technology's value factor and its penetration. To calculate D_t, I need an

assumption about how each value factor changes as penetration rises:

$$v_j(\theta_j) = v_{0,j} - a_j\theta_j,$$

where a_j is the cannibalization rate of the technology j, that how much technology j's value factor falls as its penetration increases,

$$a_j \equiv -\frac{\partial v_j}{\partial \theta_j} > 0.$$

The cannibalization rates are not used to construct the observed gap in Figure 2; they are needed only to project the output-value bias and locate the forward crossover. Namely, with the value of a_j, I have

$$D(\theta) = \theta(1 - v_0 + a\theta),$$

and

$$D'(\theta) = 1 - v_0 + 2a\theta.$$

A larger a makes the gap accelerate more sharply with penetration and therefore makes B_V rise sooner. That is why a solar-heavy installation schedule will cross the capital curve earlier.

Figure 2 cannot help to assess the cannibalization rate, as it only documents the realized level of the output-value gap. Because penetration enters D directly, the upward relationship in the figure combines the effect of greater renewable penetration θ_j with any decline in value factors v_j as penetration rises and doesn't allow to distinguish.

I calibrate a_j using the cross-border estimates of Stiewe et al. (2025), which incorporate both domestic and neighbouring penetration (Table G.2). As a diagnostic, Supplemental Appendix F.3 reports estimates using own-zone variation in the ENTSO-E panel. They have the expected sign but are imprecise and substantially smaller; because the domestic specification omits neighbouring penetration, it does not identify the same cross-border effect.

The three scenarios in the last row of Table G.2 differ only in the assumed wind–solar composition of future deployment, which is what the band is meant to bracket. Because the blend slope weights each technology by the square of its share and solar cannibalises its own value roughly three-and-a-half times faster than wind ($a_{\text{solar}} = 3.88$ against $a_{\text{wind}} = 1.09$), the aggregate slope is sensitive to the mix. The lower edge, $a \approx 0.99$, takes the current European mix (2024 Ember: about 61% wind and 39% solar of wind-plus-solar generation), which is wind-tilted and therefore flatter; the upper edge, $a \approx 1.57$, takes a solar-dominated future (60% solar), which cannibalises faster, thus leading the output-value bias to dominate the capital-deflator bias at a lower penetration. The central line is not a Stiewe blend but the independent Hirth (2013) estimate, $a = 1.30$, which falls inside the band; I use it as the reference schedule so that the central crossover does not rest on a single extrapolated mix assumption.

Capital-price bias. The forward capital curve is

$$B_K(\theta_t) = s_K\, \widetilde{\omega}_R(\theta_t)\, \gamma_{K,t}, \tag{S.54}$$

where $s_K = 0.57$ is the measured 2024 gross-output capital share of NAICS 2211, the basis on which the output-value bias in Equation (S.58) is measured, $\widetilde{\omega}_R(\theta_t)$ is current-cost renewable-capital exposure, and $\gamma_{K,t}$ is the annual gap between the broad electric-infrastructure deflator and the renewable capacity-cost deflator.

The path for $\gamma K, t$ is anchored to the IRENA installed-cost descent: it starts at the measured solar rate of 15.0% a year, is pinned to the recent three-year blend of 9.8% at $\theta = 0.29$, and decays

toward an asymptote $\gamma\infty$. The asymptote is not a free parameter. Learning does not terminate at high penetration; a mature fleet still replaces itself, so cumulative experience keeps growing at roughly $1/L$ with L the asset life, and the price decline settles at $\gamma\infty = b/L = 0.319/30 \simeq 0.011$ per year, using the solar-PV exponent and thirty-year lifespan of Way et al. (2022, Table S24). The crossover is close to invariant to this choice: $\theta^\star$ moves from 0.339 at $\gamma\infty = 0$ to 0.342 at $\gamma_\infty = 0.05$, against a span of 0.06 for the generation weight and 0.15 for the deployment mix. The blend weight behind the anchors, sixty percent solar, matters as little: over weights from 0.40 to a pure-solar 1.00, $\theta^\star$ moves by less than half a point, because the schedule takes the steeper of the solar and blended rates and solar is steeper throughout.

Renewable-capital exposure is observed only over the relatively low-penetration US range. The measured series ends in 2024 at $\theta = 0.172$, whereas the crossover exercise extends beyond $\theta = 0.5$. I therefore close the unobserved part of the capital curve by setting

$$\widetilde{\omega}_R(\theta) = \theta. \tag{S.55}$$

This is a parameter-free forward assumption, not an estimate of the exposure relationship.

The assumption is conservative for the estimated crossover penetration. In the 2024 US data, the repriced renewable assets account for 11.2% of total sector capital ($\widetilde{\omega}_R = 0.112$), compared with wind-and-solar penetration of 17.2%. Setting $\widetilde{\omega}_R = \theta$ therefore raises the forward capital curve relative to the observed US exposure-to-penetration ratio and delays the point at which the output-value bias overtakes it. Conditional on future exposure not exceeding penetration, the resulting $\theta^\star$ is an upper calibration on the exposure margin.

The measured US capital biases shown as open circles in Figure 3 use observed total-sector current-cost exposure $\widetilde{\omega}_R$ rather than Equation (S.55). The exposure closure is used only to extend the forward curve and does not enter the historical US reclassification reported in Table 1.

Output-value bias. For each technology $j \in \{s, w\}$, let $\theta_{j,t}$ denote generation penetration and $v_{j,t}$ its value factor. The generation-level output-value gap is

$$D_t = \theta_{s,t}(1 - v_{s,t}) + \theta_{w,t}(1 - v_{w,t}). \tag{S.56}$$

Its annual growth counterpart is

$$B^{\text{gen}}_{V,t} = -\Delta \log(1 - D_t). \tag{S.57}$$

Because the electricity sector also contains transmission and distribution, the object comparable with the sector capital bias is

$$B_{V,t} = w_G B^{\text{gen}}_{V,t}, \qquad w_G = 0.585. \tag{S.58}$$

All crossover calculations use Equation (S.58); no intersection between B_K and the unscaled generation bias is reported.

The forward value-factor paths use the technology-specific responses estimated by Stiewe et al. (2025), with $a_w = 1.09$ for wind and $a_s = 3.88$ for solar. I consider three calibrations. The wind-heavy calibration applies these responses to the current EU wind-and-solar mix. The solar-heavy calibration raises solar's share of wind-and-solar generation to 60 percent. The central calibration uses the Hirth (2013) blend, with an aggregate response of $a = 1.30$. In each case the calibrated value factors enter Equation (S.56), and the resulting generation bias is converted to whole-sector units using Equation (S.58).

Crossover results. For each value-factor calibration, I define the crossover as the first penetration at which

$$w_G B_V^{\text{gen}}(\theta) \geq B_K(\theta).$$

Under the central Hirth (2013) calibration, the two biases cross at approximately

$$\theta^\star = 0.34.$$

The solar-heavy calibration crosses earlier because solar has the steeper value-factor response. The wind-heavy calibration crosses later because its output-value bias rises more gradually. The sensitivity results are reported in Table G.1.

Adaptation. Storage, transmission and flexible demand flatten the net-load value profile and therefore reduce the effective cannibalisation slope. Rather than model those technologies, Table G.1 attenuates the central slope and asks what adaptation would have to deliver: a quarter off a moves the crossover from $\theta^\star = 0.34$ to 0.45, and a half off it defers the crossover beyond the modelled penetration range. The threshold is thus considerably more sensitive to the cannibalisation slope than to the saturation ceiling, which moves it by about one point; as with the composition-only column, the reason is the flatness of B_K at high penetration, where a small downward shift in B_V travels a long way along the penetration axis before the curves meet. A quarter off the central slope also happens to land almost exactly on the wind-heavy mix ($a = 0.98$ against 0.99), so those two rows describe nearly the same schedule reached by different routes, one compositional and one behavioural. Two things follow. The forward statements in the main text should be read as conditional on value-factor schedules estimated on a near-storage-free system, and, on the main-text calibration, the omission of flexibility places $\theta^\star = 0.34$ toward the early end of the plausible range. But adaptation on this scale is a postponement, not a removal: storage arbitrage value itself compresses with storage penetration—the same mechanism one step downstream—so the flexibility that flattens the renewable schedule is subject to it too (Supplemental Appendix I).

The generation weight. Because $B_V = w_G B_V^{\text{gen}}$, the generation share of sector gross output scales the entire output-value bias one-for-one, and the crossover inherits that scaling. Neither BEA nor Eurostat splits the electricity industry into generation and delivery—NAICS 2211 and NACE D35 are single industries—so w_G has to be measured from the published price and revenue decompositions. Table G.3 reports three.

The Annual Energy Outlook decomposes the retail electricity price into generation, transmission and distribution components, and the base year of each vintage is observed rather than projected (U.S. Energy Information Administration, 2026b); stacking the base years of AEO2015 through the current release gives a generation share of 0.612 on average over 2012–2025, in a range of 0.559 to 0.655. That estimator divides by retail revenue. Dividing the same EIA generation revenue instead by the BEA gross output for NAICS 2211—the denominator the sector productivity accounts actually use, and the larger of the two—gives 0.585 on average, in a range of 0.523 to 0.648. The European counterpart, the energy-and-supply component of the Eurostat price decomposition net of taxes (Eurostat, 2026a), averages 0.582 before 2021; it is an upper bound, because the energy component bundles the retail supply margin with generation, and after 2021 the gas-price spike inflates it against regulated network charges.

The three estimators share a numerator and differ in the denominator, so they are not three readings of one quantity to be averaged: only the second divides by the gross output for NAICS 2211 on which the sector productivity accounts, and therefore B_K's counterpart in this comparison, are built. The calibration accordingly maintains $w_G = 0.585$ and takes its band from that estimator's

own range, $[0.52, 0.65]$. Three checks support it. The price share rescaled by the measured ratio of retail revenue to BEA gross output, $0.612 \times 0.947 = 0.585$, reproduces it, as the arithmetic requires. Pooling the years rather than averaging nine vintage ratios gives 0.581. And the estimator is centred on its own range, $(0.523 + 0.648)/2 = 0.586$, so the base is not a point chosen inside a span. The estimator is, if anything, a mild lower bound: its numerator counts generation value once, embedded in retail revenue, while BEA gross output includes intra-industry sales, so the ratio of retail revenue to gross output falls from 0.97–1.00 in 2012–2017 to 0.85 in 2022 as wholesale trading value rose. Understatement is the safe direction, since a larger generation weight brings the crossover forward. Over the measured range the central crossover moves from 0.37 to 0.31, against 0.34 at $w_G = 0.585$; the window-to-window variation within the estimator (a mean of 0.607 over 2012–2019 against 0.540 over 2019–2022) sits inside that band. This remains a smaller lever than either the deployment mix or adaptation. One limitation: the deployment path is European while w_G is American, and no concept-matched European estimator exists, because NACE D35 carries no generation/delivery split — which is why the European row is a price decomposition and enters as a bound.

H Cross-country harmonization: US-anchored exposure and the convention gap

The unobserved exposure problem. The capital-price correction depends on renewable-capital exposure,

$$\widetilde{\omega}_R \equiv \frac{K_R}{K_D},$$

where K_R is the current-cost stock of the renewable assets affected by the deflator change and K_D is total sector capital. European national accounts do not separately identify renewable and conventional electric-power capital within NACE Section D. European exposure must therefore either be constructed indirectly or calibrated from a system in which the relevant capital split is observed.

Diagnostic European constructions. I construct a renewable-capital numerator by valuing Ember wind and solar capacity at IRENA installed-cost benchmarks and accumulating investment under a perpetual-inventory rule with depreciation $\delta = 0.04$. The resulting exposure levels depend materially on the valuation basis used for the total-sector denominator.

On this vintage the penetration shortcut overstates input-side exposure for every system reported, most sharply for Denmark, whose high penetration coexists with low constructed exposure. But the construction is not stable. Repeating it on the matched 2010–2017 window against Eurostat net book values—rather than current replacement cost—inverts the picture: the ratio $\widetilde{\omega}_R/\theta$ then ranges from 0.58 in Denmark to 6.49 in France, with Germany at 4.74 against 0.56 on the current-cost vintage. The French and German estimates are not credible as measures of renewable assets' share of total sector capital, particularly given France's low wind-and-solar generation share over that window. They diagnose a valuation mismatch: the renewable numerator is measured at estimated current replacement cost, whereas the Eurostat net book denominator carries a long-lived capital stock at depreciated historical cost, so the ratio is inflated by a basis difference rather than by any feature of European capital. An order-of-magnitude swing in the same country across two defensible constructions is why the harmonization below is anchored to the one system in which the split is measured, and why constructed European levels are not inserted directly.

The US anchor. The United States is the only system with a measured renewable/conventional capital split. Over 2011–2017 the BEA data give a mean exposure-to-penetration ratio of $\widetilde{\omega}_R/\theta = 1.00$; by 2024, with penetration at 0.172 and exposure at 0.112, the ratio is 0.65. I therefore write European exposure as $\widetilde{\omega}_R = \phi\theta$ and report a band $\phi \in [1, 2]$: $\phi = 1$ imposes the measured US relationship, and $\phi = 2$ is a deliberately conservative stress in the upward direction suggested by the matched-window construction for Germany, although that construction is not treated as a credible level estimate..[18]

The annual rule. At the measured US deflator gap, the bias per unit of exposure follows directly from Equation (2). With $s_K = 0.70$ on a value-added basis, ten percentage points of renewable-capital exposure imply

$$0.70 \times 0.10 \times \gamma_K = \begin{cases} 0.60 \text{ pp/yr}, & \gamma_K = 8.5\% \text{ (smoothed 2010–2024 bias)}, \\ 0.77 \text{ pp/yr}, & \gamma_K = 11.0\% \text{ (recent annual value)}, \end{cases}$$

that is, approximately 0.6–0.8 percentage points of measured sectoral TFP growth per year. This is the rule quoted in Section 5.

The EU–US convention gap. A system still deflating renewable investment with broad construction indexes books renewable-capital learning as TFP; the post-2023 United States books approximately none of it. The published gap therefore contains a pure convention component equal to the B_K the European accounts have not yet removed. Cumulating $B_{K,t} = s_K\phi\theta_t\gamma_{K,t}$ along the EU deployment path over the fourteen growth years 2011–2024 gives approximately 14 cumulative log points at the measured anchor $\phi = 1$ and approximately 28 at the $\phi = 2$ stress, on the value-added capital share $s_K = 0.70$ used above; the LCOE-based upper share $s_K = 0.85$ raises these to 17 and 34.[19] Both figures are sectoral. The Domar weight is irrelevant here, because the comparison being corrected is itself a sectoral one.

The Germany–US Section-D gap. The same correction applied to the matched EU KLEMS Section-D window is reported in Table H.2. Published KLEMS TFP growth is −1.31% per year for Germany against −1.59 for the United States, so Germany leads by 0.28 points. The US-anchored correction removes between 0.47 and 1.21 points per year of that lead, turning it into a deficit of 0.19 points at $\phi = 1$ and 0.93 at $\phi = 2$. Over the seven growth years 2011–2017—equivalently, the 2010–2017 endpoints—this cumulates to 3.3–8.5 log points of the published Germany–US difference attributable solely to the deflator convention.

The point estimates reverse the Germany–US ordering, but the corrected gap of roughly 0.2 percentage points per year at the measured anchor is smaller than the uncertainty in the imported γ_K the capital shares, and ϕ. The correction therefore makes the published ranking fragile but does not establish that the underlying ranking is reversed. The claim the paper makes is that published rankings are sensitive to the renewable-capital correction, not that they reverse.

I Storage as a second member of the self-correlated class

The paper's class claim is that any technology whose output is self-correlated pushes measurement bias from input prices toward output value as it scales. Grid-scale storage provides a second

[18] This ϕ is unrelated to the $\kappa \equiv s_K\omega_R$ of Supplemental Appendix C.

[19] Computed from the Ember EU wind-and-solar share and the measured annual deflator gaps; see `Results/provenance_gap_closure.md`. The bias is linear in s_K, so the two share conventions scale it proportionally.

member, observable in US data. A battery's service is intertemporal arbitrage; its per-unit value is the price spread it straddles; and, like a solar MWh, that value is eroded by neighbours doing the same thing, because charging in the trough and discharging at the peak compresses the trough-peak spread itself.

Storage also has an input-side analogue. Battery costs have followed a steep learning curve (Way et al., 2022), so a broad capital deflator could misclassify embodied progress in storage investment in the same way as for renewable generating equipment. The paper does not quantify this additional capital-price bias; the point here is only that storage can be subject to both sides of the measurement problem.

CAISO provides the test. Operating battery power in the CAISO balancing area grew from 0.5 GW in 2020 to 15.4 GW in 2025. Regressing the monthly mean intraday arbitrage spread (top four hours minus bottom four hours of day-ahead prices, April 2023 to 2025) on log storage penetration with month-of-year fixed effects gives an elasticity of -0.73 (SE 0.15, $t = -4.8$): conditional on seasonality, doubling the fleet compresses the per-unit arbitrage value by about half. Because the spread is a fuel-cost differential across hours, the natural-gas price scales it mechanically and could in principle drive this slope; gas was also volatile over the window. It does not: adding the log monthly Henry Hub gas price to the regression leaves the storage elasticity essentially unchanged at -0.74 (SE 0.19, $t = -3.8$), with the gas coefficient itself insignificant ($t = 0.1$) once month-of-year effects absorb the seasonal fuel cycle, and battery penetration and gas only moderately collinear (correlation 0.57). Indeed gas rose through 2025 while the spread kept falling, so a falling fuel price cannot account for the compression. The same distinction between a conditional value object and the fuel-driven price level separates the value factor from the fuel-driven PPI in Supplemental Appendix E.

The mapping into the accounting is identical in form to Proposition 2, with θ the storage share and v the realized-over-initial spread ratio. Storage is still small in output terms, so this is a class illustration rather than a quantitatively material bias; its role in the argument is twofold. It converts the class from an assertion into a two-member falsifiable pattern, and it bounds the adaptation escape valve of Section 5: the flexibility that flattens the renewable value-factor schedule is itself subject to the same self-cannibalization, so adaptation defers the crossover rather than removing it. The estimate rests on thirty monthly observations from one ISO and should be read as preliminary.

J Official sources

This appendix records the document-specific methodological details drawn from official sources, so that they remain visible without appearing as annotations in the reference list. The bibliography entries for these sources contain publication metadata only.

J.1 Bureau of Economic Analysis

Private fixed investment and the electric-power deflators. The general framework for private fixed investment is described in the NIPA Handbook, including the source-and-method summary in Table 6.A (Bureau of Economic Analysis, 2024b, Table 6.A). For electric-power structures, three sources establish the relevant chronology. First, the June 2023 preview of the comprehensive update announced a separate "alternative electric" price index for wind and solar structures, constructed as a weighted composite of EIA average construction costs in dollars per kilowatt and applied retrospectively from 2010. The same document states that traditional and alternative electric structures had previously been deflated together with a composite of Handy–Whitman electric-power and utility-building construction-cost indexes (McCulla et al., 2023). Second, the

November 2024 methodology summary identifies the source indexes used for alternative and other electric-power structures (Bureau of Economic Analysis, 2024c). Third, the 2025 annual update records a separate change for "other electric power structures": beginning in 2024:Q1, the source moved from Handy–Whitman to Bureau of Reclamation construction-cost data (Mataloni and McCulla, 2025). In BEA's June 2023 SCB preview (McCulla et al., 2023):

> ***Improved prices for private fixed investment in wind and solar power structures within electric power structures*** *Beginning with 2010, BEA will introduce a new price index to deflate wind and solar power structures within its estimates of real private fixed investment in electric power structures. Additionally, measures of wind and solar power structures will be presented separately in NIPA underlying table group 5.4.xU. Currently, traditional and alternative power structures are measured and presented as a single series; they are deflated using the same price index, which is based on a composite of Handy Whitman's electric power and building construction cost indexes. However, Handy Whitman's cost indexes do not cover construction cost data for alternative electric power structures; as a result, the indexes do not reflect the significant price decreases in wind and solar structures over recent years. As part of this year's update, wind and solar power structures will be deflated using a new composite price index based on a weighted composite of the Energy Information Administration's average annual costs of construction for wind and solar generators (measured in dollars per kilowatt of installed nameplate, or maximum, capacity).*

The published alternative-electric series is a Fisher price index in NIPA underlying Table 5.4.4U, series `IA001174-A`. The values used for the latest-vintage check are 78.104 in 2023 and 70.870 in 2024 (Bureau of Economic Analysis, 2026). Handy–Whitman is a proprietary trade-source family rather than a published formula (Whitman, Requardt and Associates, LLP, 2026). Reclamation's Construction Cost Trends are composite construction-cost indexes whose current movements are driven chiefly by BLS producer price indexes, supplemented by other price sources and engineering judgment (U.S. Bureau of Reclamation, 2026). This later source change remains within the broad construction-cost family and does not alter the paper's identified comparison, which reprices renewable investment while holding the conventional-capital series fixed.

Depreciation parameters. The source for the perpetual-inventory parameters is BEA's table "BEA Rates of Depreciation, Service Lives, Declining-Balance Rates, and Hulten–Wykoff Categories" (Bureau of Economic Analysis, nda). The parameter choices used in the capital-services calculation are described where they enter the analysis in Supplemental Appendix B.4.

Industry and household electricity output. For the industry accounts, BEA constructs chain-type quantity indexes for gross output by deflating the current-dollar commodity composition of an industry's output with corresponding commodity price indexes and combining the resulting commodity quantity indexes; real value added is obtained by double deflation and Fisher aggregation (Bureau of Economic Analysis, ndb). For household electricity, the 2024 NIPA Handbook describes current-dollar personal-consumption estimates based on EIA residential revenue, with quantity and price information from EIA kilowatt-hours and cents per kilowatt-hour, adjusted by BEA from a billing to a usage basis, together with the BLS electricity CPI (Bureau of Economic Analysis, 2024a, Ch. 5). Because that handbook reflects the 2024 annual update, the current source-data description is the 2025 methodology summary: annual EIA electricity quantities are combined with BLS CPIs for personal consumption, while monthly EIA physical quantities

combined with BLS PPIs are used for the electric-utility components of government consumption and private inventories (Bureau of Economic Analysis, 2025).

Access dates. The Table 5.4.4U series was accessed on May 11, 2026. The 2023 preview and 2025 annual-update pages were accessed on July 21, 2026. The output-methodology sources were accessed on July 28, 2026.

J.2 Bureau of Labor Statistics

The bibliography entries for the BLS sources contain publication metadata only. This subsection records the document-specific methodological details used in the paper.

Industry productivity output. The BLS industry productivity program defines sectoral output as the goods and services produced by an industry for sale to consumers or to businesses outside that industry. Real output is generally obtained by deflating nominal sales or values of production with BLS price indexes. For some industries, including parts of the utilities sector, it is instead measured directly from physical quantities. Where the necessary product detail is available, BLS constructs physical-quantity output indexes as Törnqvist aggregations of individual products. Under the deflated-value method, revenues are deflated at the detailed product-class level to account for changes in product mix; BLS describes the resulting indexes as conceptually equivalent to indexes constructed from physical quantities (Bureau of Labor Statistics, 2017).

Electric-utility productivity output. The documented BLS method for electric utilities bases output on kilowatt-hours sold to ultimate consumers rather than kilowatt-hours generated. Sales to each class of service are weighted by the price associated with that service class. This weighting is intended to prevent shifts in the distribution of sales among customer classes from appearing spuriously as productivity changes. Inter-utility sales are excluded from both output and input. The technical note states its coverage using the former SIC classification, including SIC 491 and the electric operations of combination utilities in SIC 492–493, and identifies the Energy Information Administration as the source of the underlying sales data (Bureau of Labor Statistics, nd).

Producer price indexes. The BLS Producer Price Index measures the net revenue received by a specified producing establishment from a specified buyer for a specified product or service under specified transaction terms. BLS seeks actual transaction prices, including applicable discounts, premiums, rebates, and allowances, rather than list or book prices. Its methodology also provides for quality adjustment when the sampled product, transaction terms, or unit of measure changes (Bureau of Labor Statistics, 2026b).

For electric power generation, NAICS 221110, the industry PPI measures price changes for the initial commercial transaction received by generating establishments and distinguishes generation by utilities from generation by non-utilities. For bulk transmission and control, NAICS 221121, the index measures changes in transmission and control fees and charges. Separate distribution indexes within NAICS 221122 cover residential, commercial, and industrial electricity; these series extend back to December 1990 (Bureau of Labor Statistics, 2005).

Household electricity prices. In the CPI, BLS prices electricity under separate rate-schedule specifications, including general residential, heating, non-heating, and water-heating service. The

consumption amounts attached to the sampled rate schedules are selected using probability procedures based on consumption averages supplied by the Department of Energy. BLS also publishes average electricity prices per kilowatt-hour. Those series are dollar levels and should not be confused with the CPI electricity index (Bureau of Labor Statistics, 2026a).

Access dates and document vintage. The BLS Handbook of Methods pages and the two electricity price fact sheets were accessed on July 28, 2026. The electric-utility productivity technical note is entered under 2007, corresponding to its stated revision date; a claim that depends on a particular document vintage should refer to the date shown in the technical note itself.

J.3 Energy Information Administration

Form EIA-861 reports electricity sales to ultimate customers, associated revenue, and customer counts by state and customer class. The data distinguish bundled service, energy-only service, and delivery-only service. EIA's reported average retail price of electricity is calculated as reported revenue divided by reported kilowatt-hours sold. It is therefore an average-revenue or unit-value statistic, rather than a fixed-specification price index (U.S. Energy Information Administration, 2026a).

The precise field definitions and service-type classifications should be read together with the instructions for the applicable Form EIA-861 reporting vintage. The source was accessed on July 28, 2026.

J.4 Eurostat

Eurostat, Handbook on prices and volume measures in national accounts, 2016 edition (Publications Office of the EU; ISBN 978-92-79-37859-1), §4.4 "CPA D — Electricity, gas, steam and air conditioning," pp. 79–80. The recommended "A method" for production: (Eurostat, 2016)

> "The A method here is to deflate value of output by the available PPIs... A fully equivalent approach, which would also be an A method, is to directly use the quantity data available on products, providing the quantity data is sufficiently detailed."

ESA 2010 §10.32 (ESA 2010) (quoted in the Handbook, ch. on value added): "the theoretically correct method to calculate value added in volume terms is by double deflation, i.e. deflating separately the two flows of the production account (output and intermediate consumption)..."

The Handbook explicitly names the intraday-value problem and instructs statisticians to treat it as a quality difference, i.e. as different products, not a value change (§4.4, p. 79, citing ESA 2010 §10.16): (ESA 2010)

> "a clear distinction is drawn between price differences based on quality (for example electricity supplied at night compared with supply during the day... or there may be a different tariff for 'green' electricity) and price discrimination."

J.5 Federal Reserve Board

The physical-output benchmark for NAICS 2211 is the Federal Reserve's industrial-production index for electric power, series `IPG2211A` of the G.17 statistical release, annual and indexed to 2017=100 (Board of Governors of the Federal Reserve System, 2026). The generation-only and technology splits used in Figure E.1 are `IPG22111A` (generation, NAICS 22111), `IPN221114T8A` (renewable and other) and `IPN221112A` (fossil). This is a quantity index built from physical energy

data rather than a deflated value, which is why it serves here as the physical counterpart to the BLS deflated-output series. Two differences from the BLS sectoral-output concept should be kept in view: the Federal Reserve index measures generation rather than sales to ultimate customers, and it is partly weather- and workday-adjusted. Neither series resolves the hour of delivery, which is the margin at issue. The series were retrieved through FRED, which mirrors the release, on July 17, 2026.

J.6 Lawrence Berkeley National Laboratory

The independent US installed-cost series in Figure 1 and Figure B.4 are taken from a single 2025 vintage of two Berkeley Lab market reports, so that both technologies rest on one consistent basis. Utility-scale solar is the capacity-weighted average installed cost of ground-mounted photovoltaic plants above 5 MW-AC, from the tab "CapEx Trend (PV-only)" of the data file accompanying Seel et al. (2025); land-based wind is the capacity-weighted average installed cost of land-based projects, from the tab "CapEx Over Time" of the data file accompanying Wiser et al. (2025). Both are published in real dollars and are converted to a nominal basis with the BEA implicit GDP price deflator wherever the comparison is against the BEA deflator, which is nominal. The 2024 observation is EIA-preliminary for both technologies, and the wind sample is thin in 2013, 2022 and 2024; this is why the settled 2010–2023 window is the reported comparison in Section B.6. The capacity weights for the Törnqvist composite are annual gross additions from the same two reports, solar from Berkeley Lab and EIA and wind from ACP, a provider difference inherent to the two publications. The equipment-import and turbine-price evidence in Section C instead uses the 2024 edition of the wind report (Wiser et al., 2024), whose data file carries the import series.

J.7 Ember

The observed penetration path $\theta(t)$ for the European Union, Germany and the United States is the "Wind and Solar" share of total electricity generation, in percent, from the yearly full release of Ember's *Yearly Electricity Data* (Ember, 2025). The 2024 values anchoring the system positions in Figure 3 are 0.288 for the European Union, 0.435 for Germany and 0.172 for the United States. The same release supplies the capacity series used for the sourcing shares in Section C and the physical generation in Figure E.1. Ember compiles from EIA, Eurostat, ENTSO-E and national statistical sources rather than collecting primary data, so it is used here as a harmonised cross-country access layer; as a level check, its global solar capacity at end-2024 is within one percent of the IRENA figure. Accessed June 22, 2026.

J.8 ENTSO-E

The European bidding-zone panel is built from the ENTSO-E Transparency Platform (ENTSO-E, 2026): hourly day-ahead prices (Article 12.1.D), actual generation per production type (Article 16.1.B&C) and actual total load (Article 6.1.A), retrieved through the Web API for 15 bidding zones over 2015–2024. Penetration on this panel is generation over total generation, not generation over load; the two differ for net-trading zones, of which West Denmark (bidding zone DK1) is the clearest case, and the generation basis is the one under which the gap $D = \theta_s(1 - v_s) + \theta_w(1 - v_w)$ equals the proportional shortfall of value-weighted relative to physical generation exactly. On this basis DK1 reaches $\theta = 0.751$ in 2024, which is the position reported for West Denmark in the main text; the whole-country Danish figure of 0.692 in Table H.1 is a different aggregate and is not comparable. Negative prices are retained. Accessed July 16, 2026.

Table F.2: Quantity-composition gaps under alternative delivery-time stratifications, 2015–2024 (Europe) and 2018–2025 (United States). $B_V^{\text{cell}} = \Delta \log Q^{\text{phys}} - \Delta \log Q^T$, where Q^T is a chained Törnqvist volume index. The first three columns apply that index to the same hourly prices and generation, varying only the aggregation strata: residual-load deciles, residual-load deciles interacted with quarter, and month-by-hour-of-delivery cells. Residual-load breakpoints are fixed within each zone using the pooled distribution, so that cells are comparable across years. The final column reports the current-value bias $B_V^{\text{gen}} = -\Delta \log(1 - D)$ over the same transitions, computed from the same hourly observations. n is the number of annual transitions, which differs across zones, so cumulative figures are not comparable across rows without it. Rows are ordered by B_V^{gen}. Zone-specific cumulative log points, generation basis. The ordering clock-time < system-state survives excluding the transitions spanning the 2021–2023 gas-price shock; the levels do not (Supplemental Note F.7).

		System state		Clock time	Current value
Zone	n	Deciles	Deciles × quarter	Month × hour	B_V^{gen}
Europe (ENTSO-E), 123 transitions					
DK-1	9	+4.28	+5.48	+0.08	+11.53
DK-2	9	+2.77	+1.63	−0.30	+10.48
ES	9	+1.76	+1.03	+0.05	+8.16
DE-LU	5	+2.21	+5.03	+2.44	+7.99
BE	9	+3.48	+2.95	+3.71	+7.55
AT	5	+0.37	+2.11	+3.06	+4.01
IT-SICI	9	+0.37	−2.74	−3.29	+3.11
IT-SARD	9	−0.16	−0.40	+0.25	+2.99
NL	9	+1.39	+1.83	+0.25	+2.62
FR	9	+1.33	+0.27	+1.84	+2.24
IT-CSUD	9	+0.02	+1.20	+1.00	+1.98
IT-CNOR	9	+0.72	+0.08	+0.22	+1.16
GB	5	+0.66	+1.45	+0.01	+1.11
IT-NORD	9	+0.76	−0.51	+0.61	+0.93
IT-SUD	9	−4.50	−4.77	−3.40	−0.80
Median		+0.76	+1.20	+0.25	
Positive		13/15	11/15	12/15	
United States, 12 transitions					
CAISO	5	+6.87	+5.33	+6.23	+8.20
ERCOT	7	+7.21	+9.38	−1.65	+2.89

Table F.3: Sensitivity of the European stratification results to the exclusion of price-shock transitions. Entries are cross-zone medians of cumulative log points, counts of positive zone estimates, and the pooled chained index over zone-by-month-by-hour cells.

	Full panel	Excl. 2021–23	Excl. 2020–23
Transitions	123	95	81
Median: residual-load deciles	+0.76	+0.89	+0.66
Median: deciles × quarter	+1.20	+0.83	+0.48
Median: month × hour	+0.25	−0.30	−0.39
Positive: residual-load deciles	13/15	12/15	10/15
Positive: month × hour	12/15	6/15	5/15
Pooled hour-location index (Europe)	+1.34	−0.14	−0.84
ERCOT: residual load / clock time	$+7.21/-1.65$	$+7.92/+1.22$	$+6.81/+2.15$

Table G.1: Sensitivity of the whole-sector crossover $\theta^\star$. Both wedges are in whole-sector gross-output units, $B_V = w_G B_V^{\text{gen}}$ against B_K built with the gross-output capital share $s_K = 0.57$, with $w_G = 0.585$ unless the row says otherwise. The first block varies the assumed wind–solar composition of future deployment (Table G.2). The second is a reduced-form adaptation check: it attenuates the central cannibalisation slope by a quarter and by a half, asking what storage, transmission and demand response would have to deliver to defer the handoff, without modelling any of them. The third varies the generation share of sector gross output, which scales the whole output-value wedge one-for-one. Columns: the baseline logistic ceiling $\theta_{\max} = 0.90$; the $\theta_{\max} = 0.80$ robustness check; and a specification retaining only the $(1-v)\Delta\theta$ composition term of Equation (5), discarding the $-\theta\Delta v$ value-decline term, at the baseline ceiling. On the calibrated path the two channels of Equation (5) are of equal size by construction, because a linear schedule $v = v_0 - a\theta$ with $v_0 = 1$ implies $1-\bar{v} = a\bar{\theta}$; the final column is therefore a halved wedge and a bounding specification, not an empirical decomposition. The third block spans the measured range of w_G documented in Table G.3. The fourth varies the gross-output capital share, which scales B_K one-for-one: from its 2010–2017 mean to 0.70, the level it approaches as fuel inputs leave the sector; the central value is the measured 2024 share. "> range" denotes a handoff deferred beyond the modelled penetration range.

		Full wedge		Composition channel
Sensitivity	Assumption	$\theta_{\max} = 0.90$	$\theta_{\max} = 0.80$	$\theta_{\max} = 0.90$
Central	$a = 1.30$, $w_G = 0.585$, $s_K = 0.57$	0.34	0.35	0.56
Wind-heavy	EU 2024 mix, $a = 0.99$	0.44	0.49	> range
Solar-heavy	60% solar, $a = 1.57$	0.29	0.29	0.47
Adaptation	a cut 25%: $a = 0.98$	0.45	0.49	> range
Strong adaptation	a cut 50%: $a = 0.65$	> range	> range	> range
Low generation share	$a = 1.30$, $w_G = 0.52$	0.37	0.39	0.60
High generation share	$a = 1.30$, $w_G = 0.65$	0.31	0.32	0.52
Low capital share	$a = 1.30$, $s_K = 0.44$	0.28	0.28	0.47
High capital share	$a = 1.30$, $s_K = 0.70$	0.40	0.42	0.63

Table G.2: Value-factor slopes used to calibrate the output-value bias. The technology-specific slopes a_j in $v_j(\theta_j) = v_0 - a_j\theta_j$ are the sum of the own-market and neighbouring-market coefficients from the cross-border spatial estimates of Stiewe et al. (2025), in percentage-point value-factor loss per percentage point of own-technology generation share. Mix blends apply $a_{\text{blend}} = \sum_j f_j^2\, a_j$, with f_j the technology's share of wind-plus-solar generation; the Hirth (2013) blend is the central calibration and the Stiewe-based mixes define the band.

	Own-market	Neighbouring	Slope a_j
Wind	0.61	0.48	1.09
Solar	1.49	2.39	3.88
Blended calibrations (slope a)			
EU wind-heavy mix (2024)			0.99
60% solar mix			1.57
Hirth (2013) blend (central)			1.30

Table G.3: The generation weight w_G, measured. Each row is the generation share of electricity-sector output implied by one published decomposition. The two American estimators share a numerator and differ in the denominator: retail revenue in the first, BEA gross output for NAICS 2211 in the second. The European row excludes taxes and is an upper bound, since the energy-and-supply component includes the retail supply margin.

Estimator	Window	Mean	Min	Max
US, AEO price share: generation/(generation+T+D)	2012–2025	0.612	0.559	0.655
US, EIA generation revenue / BEA 2211 gross output	2012–2024	0.585	0.523	0.648
EU27, Eurostat energy-and-supply share, ex-tax	2017–2020	0.582	0.515	0.665
Maintained in the calibration		0.585	0.520	0.650

Sources: EIA, *Annual Energy Outlook*, Table 8, "Prices by Service Category", base year of vintages AEO2015 to the current release; BEA, Underlying GDP-by-Industry, gross output, NAICS 2211; Eurostat `nrg_pc_204_c` and `nrg_pc_205_c`, EU27, all consumption bands, components `NRG_SUP` and `NETC`.

Table H.1: Constructed renewable-capital exposure outside the United States, latest available vintage. Numerator: Ember wind and solar capacity at IRENA replacement cost with a perpetual-inventory net-stock adjustment. Denominator: Eurostat `nama_10_nfa_st`, NACE Section D, total fixed assets at current replacement cost.

System	Vintage	θ	$\widetilde{\omega}_R$	$\widetilde{\omega}_R/\theta$
Germany	2024	0.435	0.243	0.56
European Union	2023	0.270	0.178	0.66
Spain	2023	0.404	0.208	0.51
Denmark	2024	0.692	0.115	0.17

Notes: Eurostat NACE D is broader than electric power, offshore wind is not separately identified, and the ratios are sensitive to the valuation basis of the denominator. The table is a diagnostic, not an exposure measurement.

Table H.2: Matched 2010–2017 Section-D ranking under the US-anchored correction. This is a cross-country EU KLEMS diagnostic, not the paper's preferred US NAICS 2211 benchmark.

System	KLEMS TFP (%/yr)	$\bar{\theta}$	$\widetilde{\omega}_R/\theta$ (constructed)	B_K band (pp/yr)	Corrected TFP (%/yr)
Denmark	+1.15	0.406	0.58	2.31–4.61	−3.47 to −1.16
Netherlands	−0.17	0.071	1.71	0.37–0.74	−0.90 to −0.54
France	−0.59	0.043	6.49	0.20–0.40	−0.99 to −0.79
United States	−1.59	0.053	1.00 (measured)	0.26	−1.85
Germany	−1.31	0.160	4.74	0.74–1.47	−2.78 to −2.05
Italy	−3.19	0.123	1.73	0.68–1.36	−4.54 to −3.87
Finland	−2.22	0.027	2.34	0.15–0.30	−2.52 to −2.37

Notes: B_K bands apply the US-anchored $\widetilde{\omega}_R = \phi\theta$, $\phi \in [1, 2]$, not the constructed ratios in the fourth column, which are shown only to document their scatter. France's ratio of 6.49 is the clearest evidence that the raw construction is unreliable and is retained for that reason.

Table J.1: BEA source documents cited for the electric-power capital-deflator comparison.

Citation	Document	Used for
Bureau of Economic Analysis (2024b)	NIPA Handbook, Ch. 6 (Private Fixed Investment)	General PFI methodology
Bureau of Economic Analysis (2024c)	Updated Summary of NIPA Methodologies (Nov. 2024)	Deflation-methods table
McCulla et al. (2023)	Preview of 2023 Comprehensive Update (Jun. 2023)	Handy–Whitman → \$/kW switch, applied back to 2010
Mataloni and McCulla (2025)	2025 Annual Update (Nov. 2025)	Conventional line → Bureau of Reclamation, from 2024:Q1
Bureau of Economic Analysis (2026)	Table 5.4.4U data series (via DB-nomics)	Actual deflator numbers

Table J.2: BEA deflation of private fixed investment in electric-power structures.

Structure type	Price index, all years except most recent	Price index, most recent year
Alternative electric (wind and solar)	BEA weighted composite of EIA average annual construction costs for wind and solar generators, measured in dollars per kilowatt of installed nameplate (maximum) capacity (Bureau of Economic Analysis, 2024c)	
All other electric (conventional)	Weighted average of trade-source cost indexes for electric light-and-power plants and utility buildings (Handy–Whitman) (Whitman, Requardt and Associates, LLP, 2026; Bureau of Economic Analysis, 2024c)	Moving average of the Bureau of Reclamation composite construction-cost index (U.S. Bureau of Reclamation, 2026; Bureau of Economic Analysis, 2024c)

Note: The alternative-electric index was introduced in the 2023 comprehensive update and applied back to 2010 (McCulla et al., 2023). The source for the conventional line changed from Handy–Whitman to Bureau of Reclamation data beginning in 2024:Q1 (Mataloni and McCulla, 2025). These are separate changes: the first is the renewable-capital deflator studied in this paper; the second changes only the source used for other electric-power structures.