Australian National University
Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

Declining Fossil-Plant Utilisation and Falling Substitutability with Wind & Solar

CAMA Working Paper 91/2025
October 2026

Anthony Wiskich
Centre for Applied Macroeconomic Analysis, ANU
Swinburne University of Technology

Abstract

The elasticity of substitution between wind and solar (W&S) and fossil-fuel electricity generation is pivotal for the cost of decarbonisation, yet unsettled. It is not a constant: as intermittent W&S penetration deepens, integration costs rise and substitutability falls. We estimate this profile from the capacity margin—a reduced-form, supply-side measure of how fast utilisation of the retained fossil fleet falls as the W&S share rises. In a country panel through 2025 the elasticity is high—about nine at today's 15% share of generation—and falls to about four as W&S penetration deepens; an independent, price-based integration-cost literature finds a similar decline. Because substitutability falls as abatement deepens, non-uniform carbon prices across regions grow more inefficient the deeper the abatement, strengthening the case for harmonising carbon prices internationally as the transition advances.

Keywords

elasticity of substitution, clean transition, electricity, variable renewable energy, production function, carbon pricing

JEL Classification

O33, Q41, Q42, Q54, Q58

Address for correspondence:

(E) cama.admin@anu.edu.au

ISSN 2206-0332

The Centre for Applied Macroeconomic Analysis in the Crawford School of Public Policy has been established to build strong links between professional macroeconomists. It provides a forum for quality macroeconomic research and discussion of policy issues between academia, government and the private sector.

The Crawford School of Public Policy is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

Declining fossil-plant utilisation and falling substitutability with wind & solar

Anthony Wiskich*

This draft: August 2026

Abstract

The elasticity of substitution between wind and solar (W&S) and fossil-fuel electricity generation is pivotal for the cost of decarbonisation, yet unsettled. It is not a constant: as intermittent W&S penetration deepens, integration costs rise and substitutability falls. We estimate this profile from the capacity margin—a reduced-form, supply-side measure of how fast utilisation of the retained fossil fleet falls as the W&S share rises. In a country panel through 2025 the elasticity is high—about nine at today's 15% share of generation—and falls to about four as W&S penetration deepens; an independent, price-based integration-cost literature finds a similar decline. Because substitutability falls as abatement deepens, non-uniform carbon prices across regions grow more inefficient the deeper the abatement, strengthening the case for harmonising carbon prices internationally as the transition advances.

Keywords: elasticity of substitution; clean transition; electricity; variable renewable energy; production function; carbon pricing.
JEL: O33, Q41, Q42, Q54, Q58.

*Centre for Applied Macroeconomic Analysis, Crawford School of Public Policy, ANU College of Asia & the Pacific, Australian National University, Australia. Email: tony.wiskich@anu.edu.au. School of Business, Law & Entrepreneurship, Swinburne University of Technology.

1 Introduction

Responding to climate change requires transitioning the global economy from dirty to clean energy, and the substitutability between these inputs governs the cost of that transition. In the electricity sector—a major source of emissions and the enabler of electrification elsewhere—wind and solar (W&S) are projected to drive decarbonisation, rising from around 15% of global generation in 2024 [Ember, 2025] to dominate by mid-century in the IEA's Net Zero Emissions scenario [IEA, 2024]. How readily W&S displaces fossil generation is therefore central to the cost and feasibility of the transition. We show that this substitutability is not a constant: it is high while W&S is a small share of generation and falls sharply as penetration deepens.

Macroeconomic models typically parameterise this substitutability with a constant elasticity of substitution (CES) between clean and dirty aggregates. The elasticity is a focal point in the debate between techno-optimists and techno-pessimists: a high value implies cheap, rapid decarbonisation, a low value implies a costly and slow one. There is, however, a disconnect between the values used and the values estimated. Integrated-assessment and directed-technical-change models routinely adopt elasticities above two [Acemoglu et al., 2012]. Empirical estimates, by contrast, are typically low, between about 0.5 and 2 [Pelli, 2012, Bello & Solarin, 2022, Papageorgiou et al., 2017, Lanzi & Sue Wing, 2011]—but these come either from price-based methods [Pelli, 2012, Bello & Solarin, 2022], which the supply approach here avoids, or from a clean aggregate that mixes W&S with hydro and nuclear [Papageorgiou et al., 2017], and a non-parametric study likewise questions how strong substitutability really is [Malikov et al., 2018]. Recent macroeconomic evidence adds a short-run/long-run dimension: from aggregate energy data, Lebbe & Heylen [2026] and Casey et al. [2026] estimate a short-run elasticity well below one that rises only in the long run, so that estimates conflating the two overstate substitutability. These are economy-wide *energy* elasticities, not electricity-specific, and their long-run margin runs *opposite* to the within-electricity intermittency margin we isolate below. Relatedly, Jo & Miftakhova [2024] obtain an elasticity that *rises* with the clean share through directed technical change [see also Jo, 2025]—an offsetting force to ours, as innovation and cheaper storage raise substitutability over time while intermittency lowers it at high instantaneous shares.

That the marginal cost of integrating intermittent W&S rises with its share is well understood in the electricity-systems literature. Although electricity is physically identical whatever its source, the intermittency of W&S makes clean and dirty imperfect substitutes; that literature summarises the extra cost of accommodating intermittent generation as *integration costs*—close to zero at low shares and rising with penetration [Gowrisankaran et al.,

2016, Ueckerdt et al., 2013, Hirth et al., 2015, Elliston et al., 2016]—whose largest long-run component is a fall in the utilisation of dispatchable (dirty) capacity. Rising integration costs imply falling substitutability, and bottom-up dispatch models make this quantitative, finding it high at low shares and falling as penetration deepens [Stöckl & Zerrahn, 2023, Wiskich, 2019].

The current paper shares the supply-side, quantity-based orientation of Papageorgiou et al. [2017]—it works from generation and capacity, not prices—but departs from it in both method and conclusion. It also revisits our own earlier output-based finding that W&S and fossil generation look like near-perfect substitutes [Wiskich, 2023], which we now show is an artefact of slow capacity adjustment rather than a true elasticity. Rather than fitting a production function to *output*, which we show does not identify the elasticity here (Appendix A), we recover the elasticity from a reduced-form, first-order relationship on the capacity margin: the rate at which utilisation of the retained fossil fleet falls as W&S penetration rises.

Our paper makes two contributions, bridging two literatures that rarely meet: the electricity-systems literature on integration costs, in value ($/MWh) terms, and the macroeconomic literature that needs a clean–dirty substitution elasticity and often assumes a single constant. First, we recover the declining elasticity profile (Section 4) from our own quantity panel, and find that the integration-cost literature arrives at the same decline from the opposite side of the market—working from hourly market prices, it measures how fast the marginal value of W&S falls with penetration [Hirth, 2013, López Prol et al., 2020, Stiewe et al., 2025], independent corroboration of what we read off quantities. Because the transition is a near-global time trend, the panel identifies the *level* of the decline only partially—we bracket the slope rather than pin it, and its high-share magnitude leans on the handful of systems furthest through the transition—but its *sign and shape*, a finite and *decreasing* elasticity, are robust across every specification.

Second, a decreasing elasticity raises the cost of adopting different carbon prices across regions, because deep abatement concentrated in one region pushes it into a low-elasticity, high-marginal-cost range (Section 5). Considering two identical regions, skewed abatement that drives only one region to an 80% clean share costs more than the same global abatement shared equally (40% in each); such skewed costs are about a quarter higher using our estimated elasticity than under a *constant* elasticity carrying the same total cost of a transition to the 80% cap. So it is the decreasing *shape* of the elasticity, not merely its average level, that can raise the inefficiency of skewed pricing, strengthening the case for harmonised carbon pricing.

2 Data

We use Ember's yearly electricity data, which reports both generation (by source) and capacity (including a fossil aggregate) for around 200 countries over 2000–2025 and is freely available (Table 1). Fossil-plant vintages used to anchor the capacity-adjustment horizon (Section 3) come from the WRI Global Power Plant Database.

Output is total generation Y (GWh). The clean input is W&S generation C; on the dirty side we use fossil generation G and the fossil *capacity* stock K (GW), whose ratio is the fleet's capacity factor—the object the primary estimate reads. Everything else—hydro, nuclear and a small bioenergy/geothermal residual—is collected in O, so that $Y = C+G+O$. The estimation indexes W&S penetration by the *fossil-only* share $s = C/(C+G)$ (Section 3); the descriptive total-generation share $C/Y = C/(C + G + O)$ is also reported. Generation is overwhelmingly fossil, W&S, hydro or nuclear; the residual proper (chiefly bioenergy and geothermal) is small—a generation-weighted 1.9% of output, median 0.4% per country-year. Regions with high hydro, nuclear, or trade shares are dropped—hydro above 30% or nuclear above 35% of the country's generation, or absolute net imports above 60% of generation—so that the sample is systems where the integration costs of intermittency are borne by fossil utilisation rather than absorbed by hydro, nuclear, or imports. We exclude a single negative-generation data anomaly and five economies whose generation is dominated by documented war or supply crises rather than the energy transition (Lebanon, Yemen, Syria, Libya, Iraq); their series are the most volatile in the panel (four to five times the median year-on-year variation). Table 1 describes the resulting sample. The fossil fleet mixes coal and gas—whose capacity Ember reports separately—and its composition shifts over the transition; we show in Appendix C that both fuels are throttled as W&S rises and that the utilisation fall we identify is a within-fuel effect, not an artefact of the changing coal/gas mix.

Curtailment. Ember's generation series are metered, delivered output, so W&S is measured *net* of curtailment—arguably the economically relevant quantity, since curtailed energy has no market value, displaces no fossil generation, and is itself one of the integration costs the marginal-value measure of Section 3 is built to capture. It enters our estimate *only* through the penetration measure—a curtailing system appears at a lower share than its build-out implies, which if anything *steepens* the measured slope among the leaders, where curtailment is largest [IEA, 2023]. But the capacity build-out robustness of Section 4.2, built from installed W&S *capacity* rather than net-of-curtailment generation, is free of this channel and returns a similar slope, so the finite elasticity is not a curtailment artefact.

The identification mechanism is already visible descriptively. The capacity factor of the fossil fleet falls sharply and monotonically with the W&S share—from about 0.54 where

Table 1: Estimation sample (Ember, 2000–2025)

	Ember panel (filters applied)	Value-factor sample (crisis-excluded)
Country-years	1,694	1,560
Countries	90	85
Years covered	2000–2025	2000–2025
Median W&S share	0.4%	0.6%
Mean W&S share	4.8%	5.1%
Country-years with W&S share $\geq 10\%$	281	271
Country-years with W&S share $\geq 30\%$	56	55
Country-years with W&S share $\geq 40\%$	29	29
Maximum W&S share	72% (Denmark, 2025)	

Ember panel: fossil capacity > 0.5 GW, positive generation, and the hydro ($< 30\%$ of generation), nuclear ($< 35\%$) and net-import ($|\text{net imports}| < 60\%$ of generation) filters above. Value-factor sample additionally excludes five war/supply-crisis economies and requires a well-defined fossil capacity factor ($0 < \text{cf} < 1.2$; the upper bound admits mild overshoot above one—an artefact of within-year capacity additions/retirements and of the generation and capacity series being drawn from different underlying sources—while excluding gross inconsistencies). In this table and the descriptive facts, the W&S share is the total-generation share $C/Y = C/(C + G + O)$, with O all non-W&S, non-fossil generation (hydro, nuclear and the small other residual); the estimation index is the fossil-only share $C/(C + G)$ (Section 3).

W&S is below 10% of generation to 0.25 where it exceeds 40% (Table 2, Figure 1a)—and high shares are no longer rare: 55 country-years exceed a 30% W&S share and shares reach about 72% (Denmark) by 2025 (Table 1). A first look at the systems furthest through the transition previews the mechanism. Among the early adopters—the systems that crossed a 15% W&S share by 2016—fossil utilisation falls and stays low as their shares climb to 40–70% (Denmark $0.58 \to 0.11$, Portugal $0.62 \to 0.24$, Spain $0.46 \to 0.26$, Germany $0.52 \to 0.34$), with no return toward the ≈ 0.5 low-share baseline (Figure 1b). Two features stand out and are used later. First, the drop is large. Second, it does not reverse even where systems have had the most time to adjust: within the leaders (seven systems) the recovery is statistically present but tiny—about 0.4 capacity-factor points per year against a 20–40 point drop, an order of magnitude too small for full re-adjustment. And across the panel, in the thirty-six country-spells where the share moved less than two points over three years, the capacity factor barely changes (median +0.4 point). This descriptive decline is consistent with the value-factor slope that Section 4 estimates—with country and year effects—and determines the elasticity.

Table 2: Fossil-fleet utilisation falls with the W&S share (Ember, 2000–2025)

W&S share of generation	0–10%	10–20%	20–30%	30–40%	>40%
Mean fossil capacity factor	0.54	0.44	0.39	0.34	0.25
Country-years	1,289	154	62	26	29

Value-factor sample (crisis-excluded, $n = 1{,}560$).

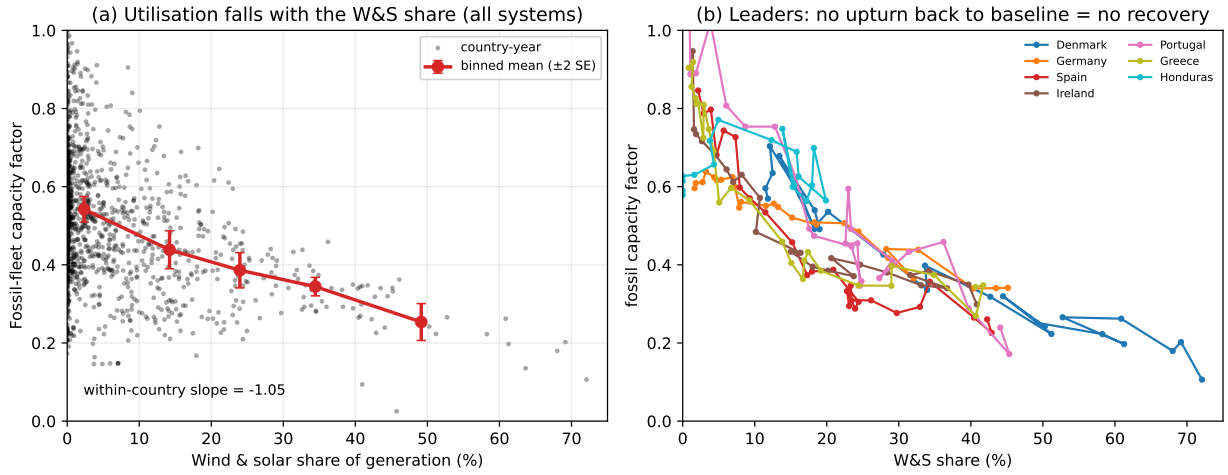


Figure 1: Fossil-fleet utilisation falls as the wind & solar share rises. *Panel (a):* every country-year (grey) with binned means (±2 SE, red); the within-country slope is −1.05 (a raw country-fixed-effects reading; the primary two-way estimate is smaller, Section 4.1). *Panel (b):* the same fall traced within the early adopters (systems that crossed a 15% share by 2016); utilisation stays low as penetration rises with no upturn back toward the low-share baseline—the signature of a *permanent* component, not merely a transitional adjustment lag (Appendix A). From Ember data.

3 Empirical strategy

We work on the supply side, treating W&S generation and the fossil *fleet* (its capacity stock) as inputs to electricity output, which avoids prices and the regional policies that drive W&S investment. The natural implementation—fitting a CES production function (or the variable-elasticity VES form of Antony, 2009) to output following Papageorgiou et al. [2017]—does not, however, identify the elasticity here. Recovering σ from a production function is delicate in general, as the elasticity is entangled with biased technical change [León-Ledesma et al., 2010] and output responds to the substitution curvature only at second order; here that second-order response is swamped, and a realistic multi-year capacity-adjustment lag makes a genuinely finite elasticity read as "perfect" (Appendix A; Monte Carlo, Appendix D). Faced with this, we build the estimate on a *first-order* signal from the

capacity margin rather than the *second-order* signal in output.

Definitions. Output Y_{it} is total generation. The clean input is W&S generation, C_{it}. The dirty input is the fossil *capacity* stock K_{it} (GW); fossil generation is G_{it}, and fossil utilisation is the capacity factor

$$\mathrm{cf}_{it} = \frac{G_{it}}{8760\, K_{it}}, \tag{1}$$

where annual generation G_{it} (in energy units) is divided by the fleet's potential annual output, capacity K_{it} times the 8760 hours in a year; the ratio is the dimensionless capacity factor. Pairing clean generation with dirty *capacity* is deliberate: the object of interest is how W&S output displaces the *stock* of dispatchable plant and its utilisation (as in Papageorgiou et al., 2017; we do not use the fossil *fuel*/generation flow as the input because it merely tracks fossil output and, being the dispatchable balancing supply, moves with demand shocks, whereas capacity is predetermined and is the stock whose utilisation the value factor traces). Penetration is measured by the *fossil-only* W&S share $s_{it} = C_{it}/(C_{it} + G_{it})$, clean generation relative to the fossil generation it displaces. This is the exact index for the clean–fossil substitution margin the value factor traces: hydro, nuclear and other non-W&S generation O enter output *separably* from the clean–fossil pair (the nest $F(H(C, G), O)$ below), so the marginal rate of substitution between C and G—and hence the value factor and the elasticity—depends only on their own mix $C/(C + G)$, not on O.[1] It is also the margin the outcome reads, since fossil is the marginal, load-following plant whose utilisation traces the value factor while must-run or energy-limited O does not. Pairing fossil utilisation on the outcome side with the fossil-only share is moreover the internally consistent choice: the conventional total-generation share $C/Y = C/(C + G + O)$ would lump O in with fossil as a single firm aggregate, which to be coherent would require *total* system utilisation—not fossil utilisation—on the outcome side, a different object from the one we measure.

The value factor and the elasticity. The *value factor* $v(s)$ is the market value of a marginal W&S unit relative to a unit of firm (dispatchable) generation—the object of the market-value, or "cannibalisation", literature [Hirth, 2013, López Prol & Schill, 2021, López Prol et al., 2020]. Under a constant-returns production function in which the clean and fossil inputs form a separable nest, $Y_V = F\big(H(C, G), O\big)$—value-adjusted (economic) electricity output, distinct from physical generation $Y = C + G + O$, to which a marginal W&S unit adds less than a firm fossil unit does—that value equals the marginal rate of substitution

[1] This separability is testable, and holds over the other-generation range we observe (an O/Y tenth-to-ninetieth percentile of 0–0.34). Letting the permanent value slope vary with the mean-centred other share, $v = 1 - (k_{\mathrm{perm}} + k_1 \widetilde{O/Y})\, s - k_{\mathrm{tr}}(s - s_{i,t-L})$, gives $k_1 = 0.02$ (SE 0.31, 90% CI $[-0.23, 0.58]$)—indistinguishable from zero—and splitting the panel by countries' median other share leaves the slope flat across terciles ($k_{\mathrm{cf}} = 0.63, 0.68, 0.64$ from the lowest to the highest O/Y). The clean–fossil margin is thus insensitive to the level of O, as the nest requires.

within the nest,

$$v(s) = \frac{\partial Y_V/\partial C}{\partial Y_V/\partial G} = \frac{\mathrm{MP}_C}{\mathrm{MP}_G} = \frac{H_C}{H_G}, \tag{2}$$

which by separability depends only on the fossil-only mix $s = C/(C+G)$. Integration costs are the gap $1 - v(s)$, near zero at low shares and rising with penetration. The elasticity of substitution is tied to the value factor by

$$\sigma(s) = -\frac{v(s)}{s(1-s)\,v'(s)}, \tag{3}$$

derived in Appendix B from the definition $\sigma = \mathrm{d}\ln(G/C)/\mathrm{d}\ln v$ along an isoquant. The integration-cost literature finds the decline is approximately linear, $v(s) = 1 - ks$; a widely-cited decomposition [Hirth, 2013] puts the slope at $k \approx 0.67$ for the observed W&S mix, with illustrative values of order 1 for wind and 2 for solar. With a linear v,

$$\sigma(s) = \frac{1-ks}{k\,s(1-s)}, \tag{4}$$

and the single slope k fixes the entire elasticity profile.

The estimating equation. The regression reads a *utilisation slope* off the capacity margin—the rate at which firm fossil utilisation falls with the W&S share,

$$\mathrm{cf}(s)/\mathrm{cf}(0) = 1 - k_{\mathrm{cf}}\, s. \tag{5}$$

This k_{cf} is a *quantity* object and must be kept distinct from the *value-factor slope* k in $v(s) = 1 - ks$, the marginal object (2) that drives the elasticity. Our identification is the first-order relationship $k_{\mathrm{cf}} = k$: once the fossil stock has adjusted, an *average* object (utilisation) traces the *marginal* one (the value factor $\mathrm{MP}_C/\mathrm{MP}_G$), $\mathrm{cf}(s)/\mathrm{cf}(0) = v(s)$. Free entry on both sides of the market delivers this equality (derived in full, with its two first-order caveats and its factor-demand interpretation, in Appendix B). The clean-side leg anchors the value of a clean unit at its policy-inclusive cost (subsidy or carbon price included), which varies over time and across regions; we treat this level as common within a year, so the year effects γ_t absorb it, and only country-specific, time-varying policy that co-moves with a country's own rollout—the within-country confounder of Section 4.1—can tilt the bridge.

Two penetration measures. We estimate two co-primary specifications that differ only in how W&S penetration is measured. The *realised fossil-only share*

$$s_{it} = \frac{C_{it}}{C_{it} + G_{it}}$$

is direct and observed, but with realised fossil *generation* G_{it} in its denominator, the same demand-responsive quantity that enters the outcome $\text{cf} = G/(8760K)$, so a country-specific demand innovation could move regressor and outcome together (W&S generation C is relatively demand inelastic, and fossil capacity K does not respond to annual demand shocks). The *demand-clean share*

$$z_{C,it} = \frac{C_{it}}{C_{it} + G_{i,t-1}} \tag{6}$$

lags realised fossil generation G by one year, so the *contemporaneous* demand shock—which moves G and the outcome cf together—is predetermined, while the persistent, W&S-driven decline in fossil utilisation is kept in the measure. It is *demand-clean* under the paper's assumption that the demand innovation is transitory (Appendix D), and tracks s closely. A capacity-based build-out share (Section 4.2), which keeps *contemporaneous* capacity and so registers fossil retirements in the same year, returns a similar slope. In the estimating equation below the penetration regressor is written generically as x_{it}, standing for either the realised share s_{it} or the demand-clean share $z_{C,it}$; the two co-primary specifications differ only in this substitution.

Our *primary* specification takes the structural ratio $\text{cf}(s)/\text{cf}(0) = v(s) = 1 - k_{\text{cf}} s$ at face value and lets the country and year effects, and the error, act *proportionally* to utilisation—the natural form, since $\text{cf}(s)/\text{cf}(0)$ is a ratio, so each country's baseline $\text{cf}_i(0)$ should *scale* the fall rather than sit beside it. In logs this is the two-way (country *and* year) fixed-effects regression

$$\ln \text{cf}_{it} = a_i + \gamma_t + \ln\big(1 - k_{\text{perm}}\, x_{it} - k_{\text{tr}}\,(x_{it} - x_{i,t-L})\big) + u_{it}, \tag{7}$$

estimated by nonlinear least squares, whose year effects γ_t remove anything common across countries within a year, so that the slope is identified only from cross-country differences in the *timing* of adoption. Here k_{perm} is the *permanent* (integration-cost) value slope—the object of interest—and $k_{\text{tr}}\,(x_{it} - x_{i,t-L})$ the *transitory* drag from firm capacity displaced but not yet retired over the horizon L. The multiplicative country effects a_i (additive in logs) absorb each country's own log-baseline $\ln \text{cf}_i(0)$, so the common low-share level differences out and *no* $\text{cf}(0)$ normalisation is needed; the value factor v itself stays *linear*, $v(s) = 1 - k_{\text{cf}} s$, exactly as in (8). The impact (contemporaneous) reading is $k_{\text{cf}} = k_{\text{perm}} + k_{\text{tr}}$ and the long-run reading is $k_{\text{cf}} = k_{\text{perm}}$; under the maintained first-order relationship this long-run k_{perm} is our estimate of the value slope k, and the elasticity follows from (3). This split resolves an apparent tension: the bridge $\text{cf}(s)/\text{cf}(0) = v(s)$ (Appendix B) is a *long-run*, free-entry relation, so it is the permanent (long-lag) slope—not the impact one—that identifies k, the transitory term being exactly the capacity disequilibrium the free-entry condition nets out. The impact reading used below is the case $L = 0$, where the transitory term vanishes.

Dropping γ_t leaves the country-fixed-effects variant, which reads higher (Section 4.1).

A single constant-returns clean–fossil nest generates exactly this linear value factor. With marginal products $\mathrm{MP}_C = \partial H/\partial C$ and $\mathrm{MP}_G = \partial H/\partial G$, and $A > 0$ a productivity constant,

$$H(C,G) = A\sqrt{(1-k)C^2 + 2CG + G^2}, \qquad \frac{\mathrm{MP}_C}{\mathrm{MP}_G} = \frac{(1-k)C+G}{C+G} = 1 - k\,s = v(s), \quad (8)$$

the $-kC^2$ term (the radicand is $(C+G)^2 - kC^2$) an integration penalty growing with the square of the clean input; further discussed in Appendix B. The nest indexes penetration by the fossil-only share $s = C/(C+G)$, our primary index; it enters value-adjusted output through $Y_V = F(H,O)$, so the total-generation share C/Y figures only via the separable O.

The adjustment horizon. The impact-versus-permanent split in (7) depends on the delay L over which capacity adjusts. Electricity capital is long-lived and costly, so the dirty stock adjusts to its long-run (post-W&S) level only slowly—plants stay in service at reduced utilisation as long as they cover operating costs. Estimating L jointly with k_{cf} pins it down only weakly—on the capacity margin the objective is nearly flat in the lag, as it is for the output-based fit [2] —so rather than commit to one value we bracket $L = 5$–9 years and take its midpoint, $L = 7$, as our central horizon: the joint-estimation fit criterion is nearly flat across lags (footnote above), so no lag is sharply selected, and the plant-vintage evidence independently points to a long adjustment horizon, consistent with though not uniquely pinning seven years. From WRI plant vintages the fossil fleet's capacity-weighted mean age is about 20 years (coal 19, gas 21) and only about 14% is past its economic life (coal 10%, gas 20%), so a displaced plant idles for many years rather than retiring at once.

4 Results

4.1 Main

We estimate the value slope with the two co-primary penetration measures of Section 3—the realised fossil-only share $s = C/(C+G)$ and the demand-clean share z_C, which lags fossil generation by a year and so brackets the concern that cf and s share the fossil residual G.

[2] The joint-estimation fit criterion is nearly flat in the lag under the primary proportional spec—the sum of squares moves by about 3% across $L = 0$–8 with no sharp minimum (its nominal low falls at $L = 2$, but $L = 7$ lies within a fraction of a percent of it), and a geometric (Koyck) distributed lag estimated jointly with k collapses to a zero decay—so it does not select a horizon and we anchor L on plant vintages instead; reduced-form capacity-adjustment regressions likewise imply mean lags of roughly five to eleven years—the within and pooled dynamic-panel estimators bracket the truth from opposite sides, as Nickell bias pulls the within estimate down and the pooled one up—but we do not rely on them, as they share a common capacity-on-generation specification and so are not independent evidence.

In the primary two-way (country + year) FE form (7) both give the same impact slope, $k_{\text{cf}} = 0.87$ (the country-fixed-effects reading is higher, Table 3); lengthening the lag L then separates the permanent integration-cost component from the transitory one (the fossil stock still adjusting), and at our vintage-anchored horizon $L = 7$ the long-run reading—which under the first-order identification estimates the value slope k—is our headline $k \approx 0.63$ (SE 0.11), with z_C giving a near-identical 0.66 (SE 0.14). That the two agree shows the mechanical simultaneity is not manufacturing the decline. The slope sits close to the $k \approx 0.67$ that Hirth [2013]'s decomposition implies for the observed W&S mix—an independent, price-based route to nearly the same number—and rules out both a *constant* elasticity ($k = 0$) and one that *rises* with the share ($k < 0$); the implied elasticity is of order thirty at a 5% share, about nine at 20% (roughly where the world sits today, some 15% of *total* generation), around six by 30%, and about four at the 50–70% fossil-only shares now reached by the leaders (Figure 2).

Why two-way FE is primary. Country FE attributes the near-global common trend (Section 3) to W&S and reads high; netting out year effects strips it, lowering both the impact and long-run readings (Table 3). Part of the common decline is surely genuine, system-wide displacement by W&S; part could be confounders correlated with the transition

Table 3: The value factor estimated from fossil utilisation, and its permanent (integration-cost) component

	realised share s		demand-clean share z_C	
	k_{cf} (impact)	k_{cf} (long-run)	k_{cf} (impact)	k_{cf} (long-run)
Two-way (country + year) fixed effects — primary				
$L = 0$ (no delay)	0.87 (0.07)	0.87 (0.07)	0.87 (0.06)	0.87 (0.06)
$L = 3$ years	0.87 (0.07)	0.79 (0.07)	0.87 (0.06)	0.80 (0.07)
$L = 5$ years	0.87 (0.07)	0.74 (0.08)	0.87 (0.06)	0.71 (0.10)
$L = 7$ **years**	**0.87 (0.07)**	**0.63 (0.11)**	**0.87 (0.06)**	**0.66 (0.14)**
$L = 9$ years	0.87 (0.07)	0.62 (0.13)	0.87 (0.06)	0.70 (0.17)
Country fixed effects only (reads higher; absorbs the common trend)				
$L = 0$ (no delay)	0.96 (0.05)	0.96 (0.05)	0.96 (0.05)	0.96 (0.05)
$L = 7$ years	0.96 (0.05)	0.70 (0.04)	0.96 (0.05)	0.77 (0.04)

Utilisation slope read as the normalised fossil capacity factor $\text{cf}(s)/\text{cf}(0) = 1 - k_{\text{cf}} s$, estimated in the primary multiplicative (log-cf) form (7) by NLS; under the first-order identification the long-run k_{cf} estimates the value-factor slope k (Section 3). Primary estimates are two-way (country + year) fixed effects; the country-fixed-effects-only rows absorb the common global W&S/utilisation trend and so read higher (Section 4.1). Impact $k_{\text{cf}} = k_{\text{perm}} + k_{\text{tr}}$; long-run $k_{\text{cf}} = k_{\text{perm}}$ (7). The right-hand pair is the demand-clean co-primary z_C (6) (Section 3), which tracks the realised-share profile closely. Country block-bootstrap standard errors (400 reps) in parentheses; long-run 90% confidence intervals (±1.645 SE) are reported in the text.

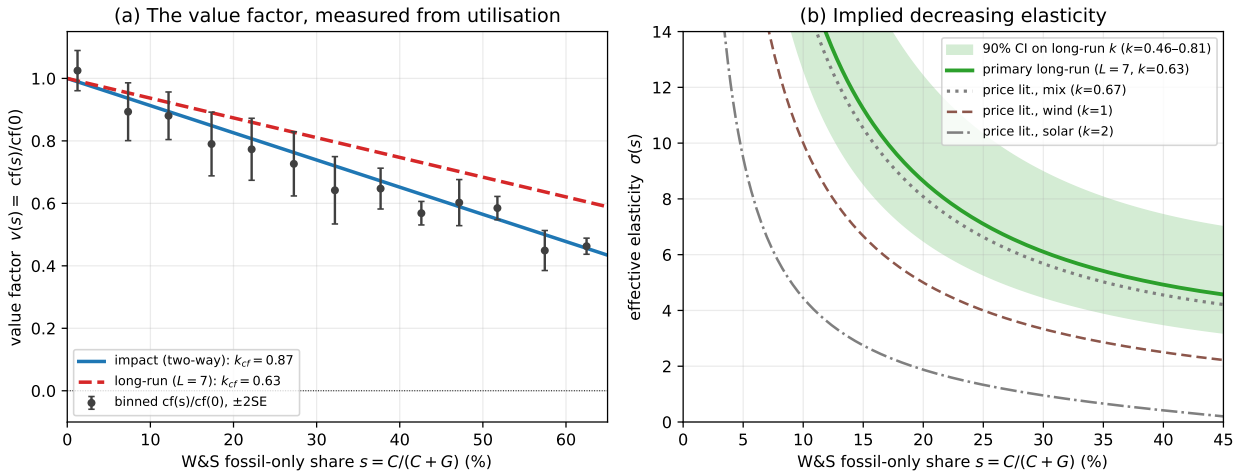


Figure 2: The value factor and elasticity estimated from the capacity margin (primary two-way FE). *Left:* the fossil capacity factor normalised to its low-share level, $\mathrm{cf}(s)/\mathrm{cf}(0)$, against the W&S share (binned, ±2 SE), with the fitted utilisation slope $\mathrm{cf}(s)/\mathrm{cf}(0) = 1 - k_{\mathrm{cf}} s$ for the two-way impact ($k_{\mathrm{cf}} = 0.87$) and long-run ($L = 7$, $k_{\mathrm{cf}} = 0.63$) readings; the data show mild convexity, so the linear form is a first-order approximation. *Right:* the implied elasticity $\sigma(s)$ from the bridge (3). The shaded band is the 90% confidence interval on the primary long-run value slope ($L = 7$, $k_{\mathrm{cf}} = 0.63$, CI $[0.45, 0.81]$), with the point estimate drawn through it; illustrative reference slopes from Hirth [2013] are shown for comparison—the observed W&S mix ($k = 0.67$) falls within the band, while the wind ($k = 1$) and solar ($k = 2$) ends lie below it (lower elasticities). The point estimate ($k_{\mathrm{cf}} = 0.63$) sits essentially at the price-mix value; the lagged demand-clean share gives a similar slope ($k = 0.66$).

(secular gas-price and demand trends, plant ageing); two-way FE cannot separate them, so in this respect it *understates* the slope—it also strips out the part of the true W&S effect that is common across countries.

The lag and the plateau. Lengthening the lag L credits more capacity inertia, but the proportional long-run reading falls only modestly and then *plateaus*, staying near 0.62 at $L = 9$ (Table 3). This fall-then-flatten shape is the signature of a slow capacity rebound: a constant true slope, run through the estimator on the real share paths, over-reads at short lags (capacity still frozen) and settles at the truth once the retirement horizon is reached. The capacity-inertia Monte Carlo of Appendix D makes this quantitative—a constant permanent slope of 0.63 reproduces the observed lag profile almost exactly, while a no-inertia placebo stays flat across every lag, so the decline is the inertia signature and not an artefact of the lagged term. It is also not an artefact of the sample shrinking as the lag lengthens (the L-lagged share is unavailable for the earliest years, so n falls from 1,560 at $L = 0$ to 895 at $L = 9$): re-estimating the whole profile on a *common* sample—the country-years present at the longest lag, holding n fixed—reproduces the same shape, the long-run reading

falling from a 0.79 impact to about 0.62 by $L = 7$–9 and plateauing there. The plateau also pins the permanent slope: because the reading flattens rather than keeps falling, it bounds k from both sides, narrowing the profile-consistent range to about 0.60–0.70 around the headline, with the plant-vintage horizon $m \approx L = 7$ (Section 3) selecting 0.63 within the band (Appendix D).

Within-country confounders. Two-way FE removes only what is common across countries in a given year; it does not remove a *country-specific*, time-varying driver that co-moves with a country's own W&S rollout. A rising domestic carbon price is the clearest candidate, but the accounting limits how much it can contaminate the estimate. Because supply meets demand and fossil generation is the balancing residual, a carbon price mostly acts *through* the regressor—by raising the W&S share s itself—which is the variation we want, not a confound. Its only routes to fossil utilisation *at a given share* are a demand effect (a lower Y) or faster fossil-capacity retirement (a lower denominator K), and once high-hydro, -nuclear and -trade systems are excluded little else can absorb the balance. Those channels are moreover of ambiguous sign—capacity retirement, for one, *raises* the utilisation of the surviving fleet—so they blur the *magnitude* of k rather than manufacture a decline.

4.2 Robustness

Fleet composition and capital cost. The dirty input is a coal–gas aggregate, and as W&S deepens the fleet shifts toward gas (gas rises about six points of fossil capacity per ten-point rise in the W&S share, coal falling correspondingly; Appendix C). Two concerns follow. First, the utilisation fall might be a *composition* artefact—a shift toward lower-utilisation gas—rather than genuine throttling; it is not, as a shift-share decomposition attributes almost all of the aggregate fall to within-fuel throttling rather than to the changing mix (Appendix C). Second, the value-factor bridge $\mathrm{cf}(s)/\mathrm{cf}(0) = v(s)$ ties utilisation to the firm plant's capacity cost r_K (Appendix B); a shift toward cheaper gas capital lowers $r_K(s)$, which biases the measured slope *up* (an over-statement of σ). Re-valuing the fleet in coal-equivalent GW at the market-implied capital ratio $\theta = \mathrm{cf}_{\mathrm{gas}}(0)/\mathrm{cf}_{\mathrm{coal}}(0) = 0.89$ (Appendix C, which under free entry equals $r_{\mathrm{gas}}/r_{\mathrm{coal}}$) purges this channel: on the both-fuels subsample—country-years reporting coal and gas capacity separately—it leaves the slope at $k \approx 0.70$ (realised) and 0.50 (demand-clean), in the same range as the headline (Table 4). The correction is small, so the changing coal/gas capital mix neither manufactures nor materially moves the decline.

Sample stability. Holding the FE structure at the primary two-way specification, the estimate is stable across the sample (Table 4): leaving out any single country, applying a volatility filter, dropping the lowest shares, restricting to a post-2010 start, tightening

or fully relaxing the trade/hydro/nuclear exclusions—the latter readmitting interconnected high-share systems otherwise dropped (chiefly Lithuania, whose low fossil utilisation partly reflects neighbours' W&S through cross-border spillovers, Stiewe et al., 2025)—and size-weighting by generation all leave the long-run k_{cf} clearly positive and near the headline, the low end (about 0.5) reached by re-admitting the crisis economies or filtering on generation volatility; only its level moves across these choices.

Mechanical simultaneity, further corroborated. The demand-clean co-primary of Section 4.1 already meets head-on the concern that cf and s share fossil generation G; two further checks agree that the simultaneity does not drive the estimate. First, a *build-out share* takes W&S penetration from installed *capacity*, $z = (\bar{f}_W K_W + \bar{f}_S K_S)/(\bar{f}_W K_W + \bar{f}_S K_S + \text{cf}_{F,t-1} K_{F,t})$—W&S capacity at fixed representative capacity factors $\bar{f}$, over that same clean capacity plus the fossil fleet valued at its *lagged* utilisation $\text{cf}_{F,t-1}$ times *contemporaneous* capacity $K_{F,t}$. Its clean numerator carries no generation (so it is free of W&S curtailment, Section 2) and its fossil denominator carries only *lagged* generation (so the contemporaneous demand shock in cf does not enter it), while contemporaneous fossil capacity keeps it sensitive to retirements. It tracks the realised share almost exactly (correlation 0.98) and, used as the penetration measure in (7), gives a long-run slope of 0.68 (SE 0.16)—essentially the headline (Table 4); building the clean side from generation rather than capacity leaves this unchanged (0.67, correlated 0.99 with the capacity version), so the capacity/generation choice is immaterial. Freezing the fossil side instead at a *fixed* representative capacity factor—the pure-capacity build-out—reads far higher (0.98), an artefact of holding constant the very fossil utilisation the estimate is about rather than letting it fall. Second, the Monte Carlo of Appendix D injects a country-specific transitory demand shock at the volatility actually observed in the panel (a within-country log-output innovation of SD ≈ 0.06): it biases the reading *upward* but loads far more on the contemporaneous impact slope than on the long-run reading we headline—a further reason to prefer the long-run one—and, because the demand-clean and build-out purges read at or above the headline rather than below, this potential inflation does not bind in the data. The mechanical simultaneity thus blurs the impact slope, not the long-run k.

The 2021–2023 gas-price crisis. A related worry is that the European gas-price crisis of 2021–2023 depressed fossil utilisation in exactly the high-share leaders that identify the deep end of the profile, co-moving with their W&S shares and so steepening the measured slope. Two checks address it. Replacing the common year effects with *Europe-specific* year effects—a full set of year dummies interacted with a European-market indicator, so that any Europe-wide shock in any year is netted out and identification rests on within-Europe adoption timing plus the rest of the world—*raises* the long-run slope rather than lowering

it; dropping the crisis years outright raises it similarly; and the demand-clean measure z_C agrees under both checks (Table 4)—so if anything the crisis attenuated rather than inflated the measured decline.

Table 4: Value-factor elasticity: robustness across sample and estimation choices, for both co-primary penetration measures

	k_{cf} (long-run, $L = 7$)		
Specification	realised s	demand-clean z_C	n
Main	0.63 (0.11)	0.66 (0.14)	1031
Capital-cost adjusted (coal-equiv., $\theta = 0.89$)	0.70 (0.20)	0.50 (0.17)	634
Leave-one-region-out (range)	0.53–0.65	0.51–0.69	1012–1031
Volatility filter (SD $\Delta \ln G_f < 0.2$)	0.52 (0.12)	0.50 (0.13)	978
Drop W&S share $< 1\%$	0.69 (0.12)	0.73 (0.21)	347
Stricter trade/hydro/nuclear exclusions	0.67 (0.15)	0.75 (0.19)	795
Include interconnected systems (no trade cap)	0.67 (0.23)	0.48 (0.20)	1050
Size-weighted (by generation)	0.73 (0.27)	0.73 (0.32)	1031
Include crisis economies	0.48 (0.22)	0.48 (0.24)	1121
Start year 2010	0.72 (0.13)	0.71 (0.15)	505
Europe-by-year FE (gas crisis)	0.71 (0.11)	0.74 (0.14)	1031
Drop 2021–2023 (gas crisis)	0.71 (0.11)	0.75 (0.16)	845
Build-out share z (W&S capacity; lagged fossil util)	0.68 (0.16)		963

*Primary multiplicative-FE (log-*cf*), two-way (country + year) FE reading*: long-run $k_{\text{cf}} = k_{\text{perm}}$ of (7) at $L = 7$, for each of the two co-primary penetration measures—the realised share s and the demand-clean share z_C (6) (Section 3)—estimated on the same base sample per row (n is the realised-share count; the demand-clean reading drops each country's first year). The FE-choice sensitivity (two-way versus country FE) is in Table 3. The sample is 2000–2025; the L-year-lagged share is available from $2000 + L$, so the long-run reading is identified off 2007 onward at $L = 7$ (the "start year 2010" row restricts it further). Standard errors (country block-bootstrap, 400 replications) in parentheses; the leave-one-region-out row leaves out each of the 85 countries in turn and reports the range of k; its n column gives the corresponding range of sample sizes (1012–1031 country-year observations). The *capital-cost* row re-values fossil capacity in coal-equivalent GW, $K_{\text{val}} = K_{\text{coal}} + \theta K_{\text{gas}} + \theta_o K_{\text{other}}$, at the market-implied $\theta = \text{cf}_{\text{gas}}(0)/\text{cf}_{\text{coal}}(0) = 0.89$ (Appendix C), netting out gas's cheaper capital and leaving k near the price-mix value (0.67); its lower n reflects the both-fuels subsample, which Ember covers thinly. The final row reports a further demand-clean measure, the build-out share z (Section 4.2; reported once, as its single denominator has nothing to vary between the two columns). Both z_C and z are reduced forms (no instrument) addressing the mechanical simultaneity between cf and s.

5 Implications for carbon pricing

A high elasticity is good news: the clean transition in electricity generation can proceed quickly and cheaply. But a *decreasing* elasticity means the transition becomes more costly as it deepens, and—the point of this section—that it matters more whether regions adopt the *same* carbon price.

It is well known that a uniform carbon price is efficient because it equalises marginal abatement costs across regions. The novel point is that a *decreasing* elasticity can amplify the welfare loss from non-uniform pricing at deeper abatement. The intuition is simple. Under a skewed policy—a carbon price in one region only—that region must undertake all the abatement, so it is pushed to a deep abatement level where, with a decreasing elasticity, it is in a low-elasticity, high-marginal-cost regime. Under a uniform price each region abates only shallowly and stays in the high-elasticity, low-cost regime. Because marginal abatement cost is convex, splitting abatement is cheaper than concentrating it; a decreasing elasticity makes the marginal cost curve *more* convex at high abatement, so the penalty for concentrating abatement is larger.

Figure 3 makes this concrete for two identical regions, using our estimated elasticity. We take the value-factor elasticity $\sigma(s) = (1 - ks)/[k\,s(1 - s)]$ with $k = 0.65$ (midway between the two co-primary slopes) over a clean share s rising from 0 to 80%, and cost each region's abatement with a constant-returns CES in which clean and dirty are equally valuable at the margin at the low-share baseline—the same normalisation as our value factor, which is measured *relative* to its low-share level and so equals one there ($v(0) = 1$). This fixes the marginal cost of the *first* unit of abatement at one for every elasticity, so the decreasing and constant cases start from a common point and any difference in total cost reflects the *shape* of $\sigma(s)$, not a different starting level. [3] Panel (a) plots the cost of the transition against the clean share reached, for the estimated elasticity and for the constant elasticity carrying the same total cost to the 80% cap ($\bar{\sigma} = 4.40$). The estimated σ is high at low shares—early abatement is cheap—and falls into the low single digits (about four) by mid-transition, reaching a minimum at $s^* = [1 - \sqrt{1 - k}]/k = 0.628$ before the linear value factor turns it up again. That turn-up beyond s^*—which lies inside the range spanned by the 80% cap—is an artefact of the linear $v(s) = 1 - ks$ rather than a measured feature; were the value factor to keep declining past s^*, deep abatement would be costlier still, so treating σ as rising there *understates* the penalty in the policy comparison below. The cost-matched constant therefore lies near these high-share values, not near the (far larger) average of $\sigma(s)$:

[3] The cost of raising a region's clean share to a while holding output fixed is $c(a;\sigma) = \left[\left(1 - \frac{1}{2}(1 - a)^{\rho}\right)/\frac{1}{2}\right]^{1/\rho} - 1$ with $\rho = (\sigma - 1)/\sigma$; the decreasing case integrates the local marginal cost at the local $\sigma(s)$, and the constant case fixes σ at the value giving the same total cost to an 80% share, $\bar{\sigma} = 4.40$.

the total cost of reaching 80% is dominated by the costly deep-abatement end, so the single constant that reproduces it is pulled down into the low-elasticity range. Panel (b) plots the extra cost of a skewed carbon price (one region taken to the 80% cap, the other left at zero) over a uniform price (each region abates equally), against global abatement. The penalty rises steeply as abatement deepens: at 40% global abatement—where the skewed region reaches the 80% cap—it is about 18% under the estimated elasticity against about 14% under the cost-matched constant, so the decreasing *shape*, not merely its average level, adds roughly a quarter to the cost of skewed pricing. At shallow abatement the ranking reverses, since a lightly abating region stays in the cheap high-elasticity range; the shape bites only once abatement is deep. This is an *illustrative* calibration on two identical regions, not an estimate: the specific magnitude ("about a quarter") depends on the assumed 80% cap and the symmetric two-region setup, and should be read as an order of magnitude rather than a measured number. What is robust is the *shape* argument, which relies on the decline of $\sigma(s)$ rather than on a precise k or a particular cap: any decreasing profile in our estimated range (k from about 0.6 to 0.7) makes the marginal-cost curve more convex at deep abatement, so the qualitative conclusion—that a decreasing elasticity amplifies the cost of skewed pricing—holds for any deep cap (a 70% cap, say, gives a smaller but same-signed penalty) and does not ride on the exact value of k.

Two considerations qualify the conclusion. First, externalities from leading regions—

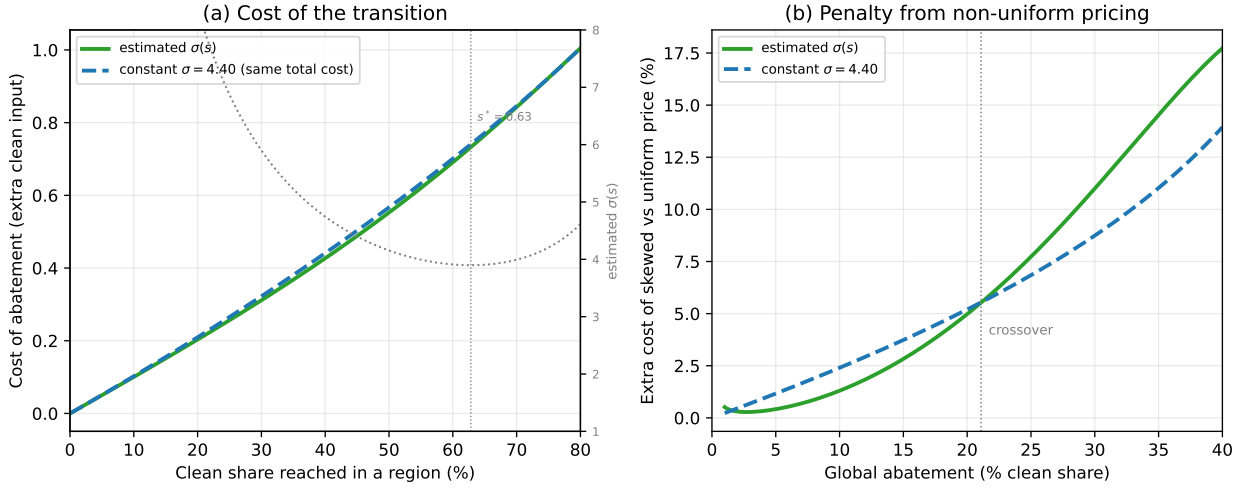


Figure 3: A decreasing elasticity raises the cost of non-uniform carbon pricing (two identical regions), using our estimated elasticity. Panel (a): the cost of raising a region's clean share, for the estimated $\sigma(s)$ (grey dotted, right axis; $k = 0.65$) and for the constant elasticity carrying the same total cost to an 80% share ($\bar{\sigma} = 4.40$). Panel (b): the extra cost of a skewed carbon price (one region to the 80% cap, the other at zero) relative to a uniform price (both abate equally), against global abatement (Section 5).

technology spillovers, regulatory learning, social acceptance, reduced perceived risk—can favour a sequential transition. Second, the extent of the elasticity's decline depends on prices, and batteries—which firm intermittent output—are falling in cost over time. This too favours uniform pricing. Cheaper batteries help most where abatement is deep, so it pays to leave the deepest cuts late. Under a uniform price both regions abate gradually and reach their deepest cuts late; under a sequential policy the lead region is pushed into deep abatement early, before battery prices fall, so only the follower captures the cheap-battery benefit.

6 Conclusion

The substitutability between wind and solar and fossil-fuel generation is summarised by a single parameter—the value-factor slope k—but that slope implies an *elasticity* of substitution that is not constant: high at low penetration and declining as the W&S share rises. That elasticity is hard to read off generation *output*, because output responds to the substitution elasticity only at second order (Appendix A). This is why we read the elasticity off the first-order signal in the capacity margin and, since a young fossil fleet idles rather than retires, bracket the adjustment horizon rather than fixing it. The resulting decline gives the object a transition model needs: an elasticity high at today's shares that falls into the low single digits as penetration deepens. A policy corollary follows (Section 5): because substitutability falls as abatement deepens, non-uniform carbon prices grow more costly the deeper the transition.

Several extensions follow. The long-run shape—here imported from the integration-cost literature—could be measured directly by estimating the W&S value factor from hourly day-ahead and capacity-market prices, which reveal the declining marginal value that generation quantities cannot; this would replace the calibrated $v(s)$ with a data-based one and pin the decline empirically. The capacity-adjustment delay could be modelled as a continuously estimated partial-adjustment process and allowed to differ between coal and gas, using unit-level retirement data—the fuel split of Appendix C shows coal and gas throttle at different rates, and the fleet's shift toward flexible gas is itself part of the firming response. Finally, incorporating electricity storage—increasingly important as batteries scale—would firm intermittent output, raise substitutability at high shares, and sharpen both the long-run elasticity and the policy conclusion.

References

Acemoglu, D., Aghion, P., Bursztyn, L., & Hemous, D. (2012). The environment and directed technical change. *American Economic Review*, 102(1), 131–166.

Antony, J. (2009). A dual elasticity of substitution production function with an application to cross-country inequality. *Economics Letters*, 102(1), 10–12.

Bello, M.O., & Solarin, S.A. (2022). The elasticity of substitution between clean and dirty energy sources. (Working paper.)

Casey, G., Gao, Y., & Kruse-Andersen, P.K. (2026). Some macroeconomic evidence on substitution between clean and dirty energy. *AEA Papers and Proceedings* (forthcoming).

Elliston, B., Riesz, J., & MacGill, I. (2016). What cost for more renewables? The incremental cost of renewable generation—an Australian National Electricity Market case study. *Renewable Energy*, 95, 127–139.

Ember (2025). *Yearly Electricity Data*. Ember, London. Available at `https://ember-energy.org/data/`.

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

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

Hirth, L., Ueckerdt, F., & Edenhofer, O. (2015). Integration costs revisited—An economic framework for wind and solar variability. *Renewable Energy*, 74, 925–939.

International Energy Agency (2023). Will more wind and solar PV capacity lead to more generation curtailment? In *Renewable Energy Market Update—June 2023*. IEA, Paris.

International Energy Agency (2024). *World Energy Outlook 2024*. IEA, Paris.

Jo, A. (2025). Substitution between clean and dirty energy with biased technical change. *International Economic Review* (forthcoming). doi:10.1111/iere.12743.

Jo, A., & Miftakhova, A. (2024). How constant is constant elasticity of substitution? Endogenous substitution between clean and dirty energy. *Journal of Environmental Economics and Management*, 125, 102982.

Lanzi, E., & Sue Wing, I. (2011). Directed technical change in the energy sector: An empirical test of induced directed innovation. WCERE 2010 Conference, mimeo.

Lebbe, T., & Heylen, F. (2026). On the elasticity of substitution between clean and dirty energy: Reconciling empirical estimates and their implications for model calibration. Ghent University, Faculty of Economics and Business Administration, Working Paper 2026/1140.

Malikov, E., Sun, K., & Kumbhakar, S.C. (2018). Nonparametric estimates of the clean and dirty energy substitutability. *Economics Letters*, 168, 118–122.

León-Ledesma, M.A., McAdam, P., & Willman, A. (2010). Identifying the elasticity of substitution with biased technical change. *American Economic Review*, 100(4), 1330–1357.

López Prol, J., Steininger, K.W., & Zilberman, D. (2020). The cannibalization effect of wind and solar in the California wholesale electricity market. *Energy Economics*, 85, 104552.

López Prol, J., & Schill, W.-P. (2021). The economics of variable renewable energy and electricity storage. *Annual Review of Resource Economics*, 13, 443–467.

Papageorgiou, C., Saam, M., & Schulte, P. (2017). Substitution between clean and dirty energy inputs: A macroeconomic perspective. *Review of Economics and Statistics*, 99(2), 281–290.

Pelli, M. (2012). The elasticity of substitution between clean and dirty inputs in the production of electricity. SURED 2012 Conference, Ascona.

Stiewe, C., Xu, A.L., Eicke, A., & Hirth, L. (2025). Cross-border cannibalization: Spillover effects of wind and solar energy on interconnected European electricity markets. *Energy Economics*, 143, 108251.

Stöckl, F., & Zerrahn, A. (2023). Substituting clean for dirty energy: A bottom-up analysis. *Journal of the Association of Environmental and Resource Economists*, 10(3), 819–863.

Ueckerdt, F., Hirth, L., Luderer, G., & Edenhofer, O. (2013). System LCOE: What are the costs of variable renewables? *Energy*, 63, 61–75.

Wiskich, A. (2019). Decreasing long-run substitutability between solar, wind and fossil fuels in electricity generation. CAMA Working Paper 70/2019.

Wiskich, A. (2023). Perfect so far? Substitutability between wind & solar and dirty electricity generation. CAMA Working Paper 27/2023.

A Why we do not estimate the elasticity from output

The output-based specifications. The natural supply-side approach fits a production function to output Y_{it} (generation) with clean input C_{it} (W&S) and dirty input the fossil capacity K_{it} of Section 3, here used as a productive input rather than as the denominator of the capacity factor. The CES, following Papageorgiou et al. [2017], is

$$\ln Y_{it} = a_i + \tau\, t + \tfrac{1}{\psi} \ln\!\big(\omega\, C_{it}^{\psi} + (1-\omega)\, K_{it}^{\psi}\big) + \varepsilon_{it}, \tag{9}$$

with elasticity $\sigma = 1/(1-\psi)$ ($\psi \to 1$ perfect substitution, $\psi < 1$ finite), $\omega \in (0,1)$ the clean-share weight and τ a common output trend. To let the elasticity *fall* with the W&S share we also fit the variable-elasticity (VES) form of Antony [2009],

$$Y_{it} = K_{it}^{\alpha} \left(C_{it} + \alpha\beta K_{it}\right)^{1-\alpha}, \qquad \sigma_{it} = 1 + \beta\, \frac{K_{it}}{C_{it}}, \tag{10}$$

whose Hicks elasticity is infinite as the clean share $\to 0$, equals $1+\beta$ when $C = K$, and declines as the share rises. Both are estimated by NLS in first differences within country. Two findings emerge, explaining why we read the elasticity off the capacity margin instead: from output the elasticity is *not identified*, and no such production function can even reproduce the fall in fossil utilisation that Section 2 documents.

Failure 1: the elasticity is not identified from output. Output responds to the curvature of the production function only at second order, and here that response is swamped. The differenced sum of squares moves by under 2% across the entire range of ψ, and a single high-leverage system—Kenya, dominated by geothermal—flips the sign of the implied elasticity (moving ψ across 1). Allowing the elasticity to decline does not help: the CES and the VES fit essentially identically (Table 5), so output cannot even distinguish a share-declining elasticity from a constant one. The estimate of Papageorgiou et al. [2017] shows the same flatness directly in the data. They report a robust $\sigma \approx 2$ from production-function regressions of clean against dirty electricity, but that finite value turns on a single high-trade observation: excluding Luxembourg—which they note should be dropped for its very high electricity-trade share, yet retain—moves ψ from 0.46 to above 1 in levels and from 0.49 to above 1 in first differences, i.e. from a finite $\sigma \approx 2$ to a convex production function with $\psi > 1$, the same perfect-substitution reading the uncorrected CES gives here. That a single observation can move the estimate across the finite/perfect boundary is precisely the flatness of the output objective—not evidence that a finite elasticity is wrong, but that output cannot pin it down.

Failure 2: no CES or VES in dirty capacity reproduces the fall in utilisation.

Dirty utilisation is dirty generation per unit of dirty capacity; where dirty capacity K is a productive input, that generation is, under constant returns, its marginal product $\mathrm{MP}_K = \partial Y/\partial K$. For the CES (9),

$$\mathrm{MP}_K = (1-\omega)^{1/\psi}\Big[\,1+\tfrac{\omega}{1-\omega}\big(\tfrac{C}{K}\big)^{\psi}\Big]^{(1-\psi)/\psi}, \qquad \operatorname{sign}\tfrac{\partial \mathrm{MP}_K}{\partial C} = \operatorname{sign}(1-\psi), \tag{11}$$

so for any finite elasticity ($\psi < 1$) utilisation strictly *rises* with the clean input, from the very first unit. The VES behaves identically: writing $x = C/K$, its MP_K has $\partial \mathrm{MP}_K/\partial x = \alpha(1-\alpha)\,x\,(x+\alpha\beta)^{-\alpha-1} > 0$ at every positive share (flat only at $x = 0$, where its elasticity is infinite). *Neither* form, with dirty capacity as an input, can generate a fall in utilisation over the positive shares we observe: the sign of the effect is fixed by $1-\psi > 0$ (with $0 < \omega < 1$) alone, independent of the estimated parameter values, so no choice of coefficients can rescue it. The data show the opposite—fossil utilisation falls with the W&S share (Table 2, Figure 1).

If one nonetheless pushes the output route. A finite elasticity *can* be forced from output by modelling the capacity-adjustment lag directly. Because dirty capacity adjusts to its long-run (post-W&S) level only slowly, output at t is in effect produced with the capacity optimal for the clean level of period $t-\delta$; replacing observed capacity with its long-run-adjusted level gives, for the CES,

$$\tilde{K}_{it} = K_{it} - \frac{\omega}{1-\omega}\Big(\tfrac{C_{it}}{K_{it}}\Big)^{\psi-1}\big(C_{it} - C_{i,t-\delta}\big), \tag{12}$$

the increase in clean input over the delay window weighted by the CES marginal rate of technical substitution $\mathrm{MP}_C/\mathrm{MP}_K$ (the VES analogue follows from its own MRTS). The output-based fit identifies the horizon δ poorly—the per-observation objective is again nearly flat, with no sharply selected minimum—so we take $\delta \approx 7$ from the plant-vintage evidence of Section 3, the same horizon as the headline value-factor lag L. The correction then returns a finite $\sigma \approx 2.5$, insensitive to the assumed delay (Table 5). But this is not an identified elasticity: the Monte Carlo of Appendix D shows it is a mechanical compression of a wide range of true elasticities (including perfect substitution) into $\sigma \approx 1.6$–2.5, driven by the observed capacity inertia rather than the underlying substitutability.

Table 5: Output-based estimates: near-perfect uncorrected, an unidentified finite value once corrected

Specification	ψ (SE) / parameters	Implied σ
CES, base sample (Ember 2000–2025)		
Uncorrected	1.10 (0.24), $\omega = 0.47$	∞
Capacity-delay, $\delta = 5$	0.58 (0.06), $\omega = 0.48$	2.4
Capacity-delay, $\delta = 7$ (base case)	0.60 (0.06), $\omega = 0.48$	2.5
VES, capacity-delay ($\delta = 7$)		
Decreasing $\sigma_{it} = 1 + \beta K/C$	$\alpha = 0.59,\ \beta = 0.15$	median ≈ 2.5; 1.1 at 50% share

Country block-bootstrap standard errors (200 reps) in parentheses. The base sample excludes a negative-generation data error and five documented war/supply-crisis economies (Lebanon, Yemen, Syria, Libya, Iraq). ω is the estimated CES clean-share weight; the estimated output trend τ runs from about -0.01 (uncorrected) to -0.02 (corrected). The VES median σ is taken over the country-year observations with a W&S share of at least 2%; over the full panel it is higher (near 7), as the sample is dominated by near-zero-share observations where the VES elasticity is mechanically large. On this sample the CES and VES fit almost identically.

B The value-factor bridge and its closed-form production function

From utilisation to the value factor: the identity $\mathrm{cf}(s)/\mathrm{cf}(0) = v(s)$**.** The estimator reads a *quantity* slope—the normalised fossil capacity factor $\mathrm{cf}(s)/\mathrm{cf}(0)$, an *average* object—whereas the elasticity is driven by the value factor $v(s) = \mathrm{MP}_C/\mathrm{MP}_G$, a *marginal* object. The reading requires the two to coincide, $\mathrm{cf}(s)/\mathrm{cf}(0) = v(s)$ (hence $k_{\mathrm{cf}} = k$); we set out here the two equilibrium conditions that deliver it and the two margins along which it bends. This formalises the free-entry argument sketched in Section 3, the single step that carries the reading from an observable average to the marginal object the elasticity depends on.

Let $p_G(s)$ and $p_C(s)$ be the market values of a unit of firm and of clean energy. *(i) Competitive input pricing.* If value-adjusted output $Y_V = F(C, G)$ is produced competitively, each input is paid its marginal product, so the value ratio equals the marginal rate of substitution, $p_C/p_G = F_C/F_G = v(s)$. *(ii) Free entry, both sides.* A unit of firm capacity earns p_G only on the energy it sells—its capacity factor $\mathrm{cf}(s)$ times nameplate—so its annual quasi-rent is $\mathrm{cf}(s)\, p_G(s)$ per unit of nameplate (net of marginal fuel cost, folded into p_G). Long-run free entry drives this to the constant annual capacity cost r_K,

$$\mathrm{cf}(s)\, p_G(s) = r_K \qquad \Longrightarrow \qquad \mathrm{cf}(s) \propto 1/p_G(s). \tag{13}$$

Clean, in turn, is built until its marginal value equals its levelised cost, anchoring $p_C(s) = c_C$

(fixed at a point in time). Combining, $v(s) = c_C/p_G(s) \propto 1/p_G(s)$ moves with the *same* object as utilisation—the scarcity value p_G of firm energy—so

$$\frac{\mathrm{cf}(s)}{\mathrm{cf}(0)} = \frac{p_G(0)}{p_G(s)} = \frac{v(s)}{v(0)} = v(s), \tag{14}$$

using the normalisation $v(0) = 1$—the value factor is measured relative to its zero-penetration level, so only its normalised slope enters and nothing economic rides on the level (in particular this neither assumes zero integration cost at $s = 0$ nor conflicts with $K(s)$, which is simply $K(0)$ there). This is $k_{\mathrm{cf}} = k$ exactly, and it is the bridge from the observable average to the marginal v that (3) below then turns into $\sigma(s)$. The step is not an accounting identity: were firm capacity to track its own generation one-for-one, utilisation would be flat; (14) holds only because capacity sits on the retention schedule $K(s) = K(0)(1-s)/(1-ks)$—the fossil fleet is retained while its energy is competed away by W&S, rather than retired in lockstep with the fossil generation that W&S displaces.

Retained capacity as factor demand. The retention schedule has a factor-demand reading. Near $s = 0$, adding clean changes the fossil capacity the system wants at the rate $\mathrm{d}K/\mathrm{d}s = (k-1)K(0)$—shed when $k < 1$, held flat at $k = 1$, added at $k = 2$—but always by *less* than the generation it displaces, which is why utilisation falls rather than staying flat. While the stock is still adjusting, the measured slope departs from the value slope: since fossil generation is the residual $G = (1-s)(C+G)$, we have $\mathrm{cf}(s)/\mathrm{cf}(0) = (1-s)K(0)/K(s)$, which for a *frozen* stock ($K(s) = K(0)$) is just $1-s$—$k_{\mathrm{cf}} = 1$ whatever the true k, pure generation displacement—while a stock being *retired* along the schedule ($K(s) < K(0)$) gives $k_{\mathrm{cf}} < 1$, and one still *built out* on net ($K(s) > K(0)$) gives $k_{\mathrm{cf}} > 1$. So the measured k_{cf} can sit anywhere while capacity adjusts—in particular above 1—and only its *long-run* value identifies the value slope k (the Monte Carlo of Appendix D makes this quantitative).

The same two first-order caveats as in Section 3 attend (14), each bending v marginally off the linear $1-ks$—consistent with the mild convexity in the fitted $\mathrm{cf}(s)/\mathrm{cf}(0)$ (Figure 2). And (14) is a *long-run* statement: it holds once the fossil stock has adjusted onto the schedule, whereas before then utilisation overshoots its long-run fall and the measured slope carries the transitory component that the lagged term in (7) separates out.

Deriving the bridge (3)**.** For the constant-returns clean–fossil nest $H(C, G)$ of clean generation C and firm fossil generation G (the value-adjusted nest of Section 3), the elasticity of substitution is $\sigma = \mathrm{d}\ln(G/C)/\mathrm{d}\ln(\mathrm{MRS})$ along an isoquant, where the marginal rate of substitution $\mathrm{MRS} = H_C/H_G$ equals the value factor v. Write the clean share $s = C/(C+G)$, so $G/C = (1-s)/s$ and $\mathrm{d}\ln(G/C) = -\,\mathrm{d}s/[s(1-s)]$. With $v = v(s)$, $\mathrm{d}\ln v = v'(s)\,\mathrm{d}s/v(s)$,

hence

$$\sigma(s) = \frac{\mathrm{d}\ln(G/C)}{\mathrm{d}\ln v} = \frac{-\,\mathrm{d}s/[s(1-s)]}{v'(s)\,\mathrm{d}s/v(s)} = -\,\frac{v(s)}{s(1-s)\,v'(s)},$$

which is (3). A linear value factor $v(s) = 1 - ks$ then gives $\sigma(s) = (1-ks)/[\,k\,s(1-s)\,]$ directly.

The closed form. The production function (8), $Y_V = A\sqrt{(C+G)^2 - kC^2}$ with $s = C/(C+G)$, has marginal rate of substitution $\mathrm{MP}_C/\mathrm{MP}_G = 1 - ks = v(s)$ exactly, so it realises this $\sigma(s)$. The limiting cases are instructive (the wind/solar labels follow Hirth, 2013's single decomposition and are illustrative): $k = 0$ gives $Y_V = A(C+G)$, perfect substitutes ($\sigma = \infty$)—value-adjusted output then coincides with physical generation; $k = 1$ (the wind end) gives $Y_V = A\sqrt{G^2 + 2CG}$ and $\sigma(s) = 1/s$ exactly; $k = 2$ (the solar end) keeps a non-negative clean marginal product only for $s < 1/2$ (there $\mathrm{MP}_C \propto G - C$), beyond which marginal curtailment makes further clean input unproductive. The general quadratic $Y_V = A\sqrt{aC^2 + 2bCG + G^2}$ nests (8) and is linear in (a, b) in Y_V^2; estimating it freely on the panel is what the footnote in Section 4.1 refers to, and it is where the output-based objective is nearly flat—the flatness the capacity-margin reading (7) avoids.

C Fleet composition: coal versus gas

The dirty input is the fossil *fleet*, an aggregate of coal, gas and a small residual of other fossil. Ember reports capacity separately for coal and gas (on thinner coverage than the fossil aggregate, one reason why we keep the aggregate as primary). But is the utilisation fall that identifies k a genuine within-fuel effect or an artefact of the changing coal/gas mix, and does the value-factor bridge's constant-r_K assumption survive the shift toward gas?

Both fuels throttle. A shift-share decomposition of the aggregate capacity-factor fall across the early adopters (comparing their sub-10% and above-30% observations) attributes about 100% of it to within-fuel throttling and about zero to composition—coal and gas begin from similar baseline capacity factors (≈ 0.61 and 0.53), so reallocating capacity between them barely moves the weighted mean. The fall is thus within-fuel, not a reweighting. The fleet does nonetheless shift markedly toward gas: across the panel gas rises from under half of fossil capacity at low shares to about 60% by a 30–40% W&S share, coal falling correspondingly—itself a signature of the firming the integration-cost mechanism describes (systems retain flexible gas and shed inflexible coal). Figure 4 shows both margins across the whole sample—the fossil fleet's coal/gas capacity mix (panel a) and each fuel's utilisation (panel b): as the W&S share rises, both fuels' capacity factors fall together while the capacity mix tilts toward gas over the bulk of the range.

Valuing the dirty capital: coal versus gas. Our dirty input, like Papageorgiou et al. [2017]'s, is installed capacity; they aggregate the fossil technologies as a simple megawatt sum in their baseline and, as a robustness check, re-value each technology's capacity at its technology-specific investment cost. Both steps bear on us, because the bridge $\mathrm{cf}(s)/\mathrm{cf}(0) = v(s)$ of Appendix B rests on a firm capacity cost r_K constant along s (free entry, $\mathrm{cf}(s)\, p_G(s) = r_K$). If the marginal and average firm plant shifts from high-capital coal to cheaper-capital gas, the condition instead delivers

$$\frac{\mathrm{cf}(s)}{\mathrm{cf}(0)} = \frac{r_K(s)}{r_K(0)}\, v(s), \tag{15}$$

so the measured slope is $k_{\mathrm{cf}} = k + \phi$ with $\phi = -\,\mathrm{d}\ln[r_K(s)/r_K(0)]/\mathrm{d}s$. A shift to cheaper capital (r_K falling, $\phi > 0$) makes utilisation fall *faster* than the value factor—free entry tolerates lower utilisation when capacity is cheap to build—so the simple megawatt-sum reading *over*states k and *under*states σ. Note that gas's *higher* variable cost does not enter here—it is already netted out in the scarcity price p_G and shows up instead in *where* each fuel sits in the merit order; what the capital valuation corrects is the fixed-cost, capacity-value dimension.

Following Papageorgiou et al.'s robustness step, we re-value the fleet in coal-equivalent capacity, $K_{\mathrm{val}} = K_{\mathrm{coal}} + \theta\, K_{\mathrm{gas}} + \theta_o K_{\mathrm{other}}$ with $\theta = r_{\mathrm{gas}}/r_{\mathrm{coal}}$ (θ_o the analogous ratio for other fossil, taken near 0.4), and re-estimate k on the resulting capacity factor. At $\theta = 1$—the plain megawatt sum—$k = 0.75$ on this both-fuels subsample (above the full-panel 0.63, a sample-composition effect—the coal/gas-reporting systems carry little other generation, so their fossil-only share nearly equals the total-generation share); valuing cheaper gas capital below coal lowers k monotonically and raises σ. How far to down-weight is disciplined by the data: under free entry $\mathrm{cf}_j(0) \propto r_j$, so the ratio of baseline capacity factors gives a *market-implied* $\theta = \mathrm{cf}_{\mathrm{gas}}(0)/\mathrm{cf}_{\mathrm{coal}}(0) \approx 0.89$ (baseline capacity factors of coal 0.58 and gas 0.52 on the both-fuels subsample, a little below the full-sample 0.61 and 0.53 quoted above)—reading the capital-cost ratio off baseline utilisation under free entry, which holds insofar as coal and gas capture the same average energy price at low W&S penetration (a fair approximation there, where both still run as high-utilisation plant and the value gap between them has yet to open)—at which $k \approx 0.70$—near the 0.67 that Hirth [2013]'s decomposition implies for the observed mix. This market-implied θ is moreover the *internally consistent* choice: it is read off the same free-entry condition ($\mathrm{cf}_j(0) \propto r_j$) that underlies the value-factor bridge itself, so adopting the engineering ratio instead would be inconsistent with the very estimator that produces k. Crediting still cheaper gas capital (engineering costs put θ nearer 0.4, gas roughly half of coal per MW) lowers k further, to about 0.27 at $\theta = 0.4$; we report that case

only to *sign* the capital channel, not as a preferred estimate.

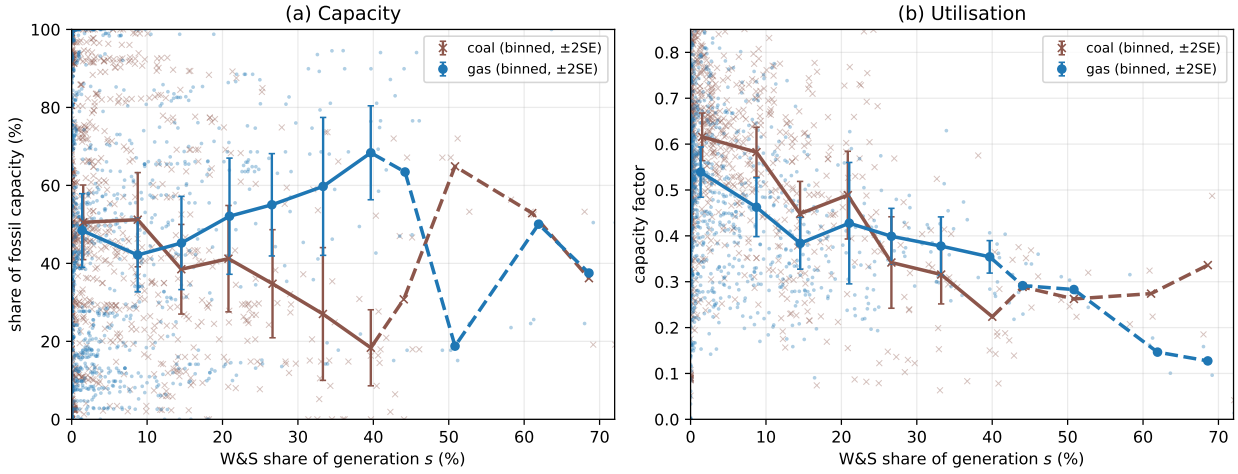


Figure 4: Coal versus gas across the whole value-factor sample (all country-years); faint points are individual country-years (coal crosses, gas dots), bold the binned means with country-clustered ±2 SE; means are drawn solid with error bars up to a 40% W&S share and dashed beyond, where too few systems are observed to bin reliably. *Panel (a):* each fuel's share of fossil capacity against the W&S share—the mix tilts toward gas over most of the range (noisily at the highest shares, where few fossil plants remain). *Panel (b):* each fuel's capacity factor—both coal and gas are throttled down as the W&S share rises. From Ember data.

D Monte Carlo evidence

We run one Monte Carlo experiment for our *primary* value-share (capacity-margin) estimator (7), and the output-based CES estimator of Appendix A. Both use the same template. Rather than a synthetic panel, we generate the outcome on the *actual* crisis-excluded input paths, so the share distribution, adoption timing and panel structure match the data by construction; we impose a structural capacity-inertia data-generating process with a *known* true elasticity; and we then apply the paper's own estimator over 400 replications, with country effects, common year effects and iid noise added and then removed as in the data. The real paths carry timing heterogeneity (staggered start years, plateaus, occasional share reversals) that identifies the capital lag, and spare us any arbitrary choice about how much of it to inject. Both experiments deliver the same lesson—capacity inertia biases the naive short-run reading toward perfect substitution, while crediting that inertia recovers the finite, decreasing long-run elasticity.

The value-share (capacity-margin) estimator

In the two-way panel the long-run utilisation slope k_{cf} falls as the assumed capacity-adjustment lag L in (7) lengthens and then *plateaus* (Table 3), and we show this fall-then-flatten shape is exactly the signature a genuine, slow capacity rebound leaves on that estimator, not an accident of the lag choice. On the real share paths we impose the structural capacity-inertia process

$$\frac{\text{cf}_{it}}{\text{cf}(0)} = 1 - k_{\text{perm}}\, s_{it} - k_{\text{tr}} \left(s_{it} - s_{i,t-m}\right)_{+}, \tag{16}$$

the same permanent-plus-transitory form and notation as the estimating equation (7) (with $(\cdot)_+$ the positive part), except that the process clips the transitory term at the fixed retirement horizon m whereas the estimator regresses on the linear change $s_{it} - s_{i,t-L}$; on the real paths the clip changes the estimated profile negligibly.

We feed in our *own* headline estimate as the truth: $k_{\text{perm}} = 0.63$ (the $L = 7$ value slope of Section 4.1), with transitory slope $k_{\text{tr}} = 0.50$ and retirement horizon $m = 7$ years—so the exercise asks whether our estimate, once put through capacity inertia, reproduces the readings we actually see. It does, almost exactly: estimating (7) on the simulated panel gives $k_{\text{cf}} = 0.85, 0.79, 0.75, 0.63, 0.62$ at $L = 0, 3, 5, 7, 9$ (mean of 400 replications) against the data's 0.87, 0.79, 0.74, 0.63, 0.62 (Figure 5a). Imposing $L = 0$ misreads the still-frozen transitory drop as permanent and inflates k_{cf} to 0.85 (above the permanent slope 0.63); lengthening L credits the drop to inertia and walks the estimate back to the truth, recovering $k_{\text{perm}} = 0.63$ at $L = 7 = m$ and holding there at $L = 9$—the fall-then-plateau

shape of the data. A placebo with $k_{\text{tr}} = 0$—a purely permanent value slope and no inertia—produces a *flat* profile ($k_{\text{cf}} \approx 0.63$ at every L), confirming that the observed decline is the capacity-inertia signature and not a mechanical consequence of adding a lagged regressor (Figure 5b)—which is why the paper headlines the $L = 7$ reading rather than the impact value.

What the profile does and does not identify. The experiment above feeds in our headline as the truth and confirms consistency; it is worth being explicit about the limits of that. Holding the retirement horizon m fixed, the lag profile identifies the permanent slope—but in principle m and k_{perm} can trade off. Searching over $(k_{\text{perm}}, k_{\text{tr}}, m)$ for the triples whose noise-free profile best reproduces the data readings $\{0.87, 0.79, 0.74, 0.63\}$ at $L = 0, 3, 5, 7$, a lower permanent slope paired with a longer horizon can match that fitted range almost as well: $k_{\text{perm}} = 0.63$ ($m = 7$) is the best (per-point RMSE 0.011), but $k_{\text{perm}} = 0.60$ ($m = 8$) lies within RMSE 0.014 (Table 6). What breaks this near-degeneracy is the *plateau*: the data flatten at 0.62 from $L = 7$ to $L = 9$, whereas the lower-slope fits reach the fitted range only by pushing the horizon out and so keep falling past it, predicting $L = 9$ readings of 0.55–0.58 below the observed plateau. The plateau thus bounds the permanent slope from *below*, while the in-sample fit bounds it from above—a higher slope overshoots at short lags ($k_{\text{perm}} = 0.90$ reads 1.04 at $L = 0$, RMSE 0.10). Together they pin k_{perm} to a narrow band of

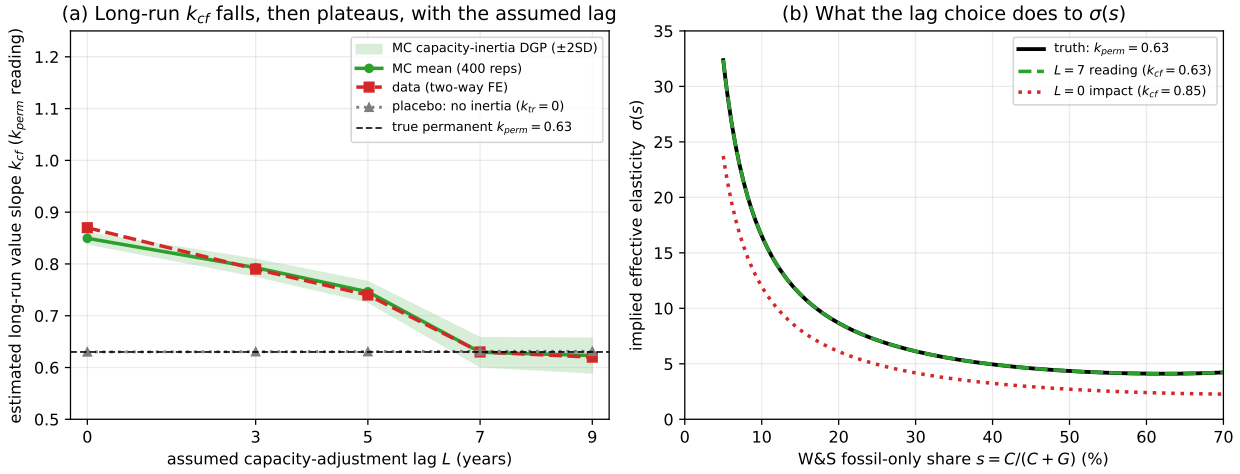


Figure 5: A slow capacity rebound reproduces the fall of the estimated value-share with the assumed lag. *Panel (a):* the long-run k_{cf} (the k_{perm} reading of the two-way FE regression (7)) against the assumed lag L, in the data (red) and in the capacity-inertia Monte Carlo (16) run on the real share paths (green, mean of 400 replications ±2SD; true permanent $k_{\text{perm}} = 0.63$, dashed). The no-inertia placebo ($k_{\text{tr}} = 0$) is flat at k_{perm}, so the decline is the inertia signature, not an artefact of the lag term. *Panel (b):* the implied elasticity $\sigma(s) = (1 - ks)/[k\, s(1 - s)]$ for the truth, the $L = 0$ impact reading (a k_{cf} that overstates the true value slope k and so understates σ), and the $L = 7$ reading (on the truth).

roughly 0.60–0.70 around the headline, tighter than the horizon trade-off alone would allow, with the vintage anchor $m \approx L = 7$ (Section 3) selecting 0.63 within it. Across the whole band the elasticity is finite and decreasing, with σ at a 70% share running from about 4.6 (at $k_{\text{perm}} = 0.60$) to about 3.5 (at 0.70); the qualitative reading is invariant, and only the high-share *level* of the elasticity moves within it.

A demand-shock placebo. The estimator's outcome cf $= G/(8760K)$ and the realised regressor $s = C/(C+G)$ (with fossil generation G in its denominator) share G, so a country-specific demand innovation moves both—utilisation down, share up—and could bias the slope. We add one to the DGP: total demand scales by $e^{d_{it}}$, $d_{it} \sim N(0, \sigma_d^2)$; W&S is must-run so fossil absorbs it, giving observed $s = s_0\, e^{-d}$ and cf $= \text{cf}_{\text{struct}}\,(e^d - s_0)/(1 - s_0)$, which reduces to the baseline at $d = 0$ (the control then recovers the truth 0.63 exactly). At the demand volatility observed in the panel ($\sigma_d \approx 0.06$, the within-country innovation SD of log output after country trend and year effects), the reading is biased *upward*, but unevenly: the contemporaneous impact reading rises from 0.85 to about 0.94, while the long-run reading we headline moves only from 0.63 to about 0.69; the transitory shock loads onto the lagged term, so the long-run k is far less exposed, and only at an implausible $\sigma_d = 0.15$ does it reach about 0.80. Crucially, this potential inflation does not bind in the data: the demand-clean co-primary of Section 4.1 and the capacity build-out share of Section 4.2, which purge the shared G directly, both read *at or above* the realised headline rather than below it—so the mechanical simultaneity is not inflating the estimate; where it bites is the impact slope, not the headline long-run k.

Table 6: The plateau pins the permanent slope from both sides: $(k_{\text{perm}}, k_{\text{tr}}, m)$ triples reproducing the observed two-way FE readings

True k_{perm}	best k_{tr}	m	RMSE/pt (L= 0–7)	fitted $L = 9$ (data 0.62)	implied σ $s = 50\%$	$s = 70\%$
0.50	0.60	10	0.017	0.55	6.0	6.2
0.60	0.60	8	0.014	0.58	4.7	4.6
0.63	**0.50**	**7**	**0.011**	**0.62**	**4.3**	**4.2**
0.75	0.50	5	0.044	0.71	3.3	3.0
0.90	1.00	8	0.102	0.55	2.4	2.0

Each row: the (k_{tr}, m) minimising the squared distance of the noise-free *primary proportional* (log-cf NLS) two-way FE lag profile from the data readings $\{0.87, 0.79, 0.74, 0.63\}$ at $L = 0, 3, 5, 7$, over the real crisis-excluded share paths. Lower permanent slopes match the fitted range only with longer horizons and then keep falling, undershooting the observed $L = 9$ plateau of 0.62 (the "fitted $L = 9$" column, an out-of-sample check); a higher slope $k_{\text{perm}} = 0.90$ instead overshoots the short-lag readings (RMSE 0.10). The plateau bounds the slope from below and the fit from above, leaving a narrow band around the headline. $m \approx L = 7$ is pinned externally by plant vintages (Section 3). Implied $\sigma(s) = (1 - ks)/[k\, s(1 - s)]$.

The output-based CES estimator

The same template applied to the output estimator makes the identical point. On the real crisis-excluded C_{it} (clean generation) and K_{it} (fossil capacity) paths we generate output from the CES (9) with a known finite elasticity ($\psi = 0.70$, $\sigma = 3.3$) and a genuine capacity delay—output at t is produced with the capacity optimal for the clean level δ years earlier, the delay-adjusted input (12)—and estimate ψ both uncorrected (feeding observed K_{it}) and with the capacity-delay correction. A genuine delay biases the uncorrected estimator upward, from the true $\sigma = 3.3$ to about 7–8 (occasionally to the perfect-substitution bound) at $\delta = 3$–5—exactly the near-perfect reading the raw CES returns on the data—while the correction recovers the truth ($\sigma \approx 3.3$–3.4) at every delay.

Imposing no delay, by contrast, leaves the same estimator unbiased across the whole range from $\sigma = 2$ to 20, confirming that the upward bias is the delay's signature rather than the estimator's.

A companion experiment probes the opposite risk—that the capacity-delay correction *manufactures* a finite elasticity even when the truth is perfect substitution. Suppose what truly enters production is fossil *generation* (the MWh dispatched) rather than capacity; as W&S rises, generation falls far below capacity as the fleet is throttled, so a capacity-based CES—as any capacity study must estimate—reads that widening gap as apparent inertia. Generating output on the real paths as a CES of clean and *actual* dirty generation, for true elasticities from $\sigma = 2$ up to perfect substitution, and then estimating from dirty *capacity*, only the perfect-substitution case reproduces the data (uncorrected estimate at the perfect bound, corrected $\sigma \approx 2.5$); any true finite elasticity leaves the uncorrected estimate below the bound, unlike the data. So the corrected $\sigma \approx 2.5$ is not an identified substitution elasticity but a property of the correction and the observed inertia—best read as the finite *capacity* elasticity of the maintained model, not as evidence that W&S is an imperfect substitute for fossil electricity.