Young, Dirty, and Stranded: Capital Vintage and the Carbon Risk Premium

Simon Gilchrist* Joseba Martinez† Natalie Rickard§

February 2026

Preliminary and Incomplete

Abstract

We argue theoretically and show empirically that carbon risk premiums compensate investors for stranding risk: the exposure of fossil-fuel-emitting capital to policy-induced value loss. We develop a putty-clay model in which machines are locked into fuel-to-labor ratios chosen at installation and newer vintages embody superior technology. An unanticipated carbon tax strands dirty capital through two competing forces: a growth wedge (disembodied productivity growth makes revenue age more slowly than fuel costs, concentrating losses on young machines) and operating leverage (which disproportionately compresses old machines' profit margins). When the growth wedge dominates, young dirty capital bears the largest stranding losses and therefore commands the highest risk premium. We find support for this prediction in two settings. In the utility sector, where generator-level data allow direct measurement of capital age and fuel type, event-study returns around cost-side regulatory shocks are significantly higher for firms with older fossil-fuel generators. In the broad cross-section, using Trucost emissions and Compustat-based capital age, the carbon risk premium is largest among capital-intensive firms in the youngest age quintile and indistinguishable from zero in the oldest quintile across all sectors, consistent with the model's prediction that old capital, having depreciated through most of its fuel exposure, carries negligible stranding risk. The emissions-age interaction is concentrated in the post-Paris period and strengthens further around 2020–2021, consistent with markets repricing stranding risk at moments when cost-side policy became more credible.

*Department of Economics New York University and NBER. Email: sg40@nyu.edu

†London Business School and CEPR. Email: jmartinez@london.edu

‡Bank of England. Email: natalie.rickard@bankofengland.co.uk

§The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees.

1 Introduction

The existence and structural origins of carbon risk premiums are of first-order importance for both finance and climate policy. If high-emission firms earn higher expected returns, this price of carbon risk governs the cost of capital for the dirtiest sectors of the economy, shapes the incentives to decarbonize, and determines who bears the financial burden of the energy transition. Yet whether such premiums reflect a coherent economic mechanism, and if so, which one, remains unresolved.

We argue that carbon risk premiums compensate investors for *stranding* risk: the exposure of fossil-fuel-emitting capital to policy-induced value loss. To make this claim precise, we develop a putty-clay (Johansen, 1959; Gilchrist and Williams, 2000) model in which machines are locked into fuel-to-labor ratios at installation and newer vintages embody superior technology. The model delivers a hierarchy of cross-sectional predictions. The first is standard: dirtier firms face higher stranding exposure and therefore earn higher expected returns. The second is new and more discriminating: the carbon premium varies with capital age. Young dirty capital has locked in large fuel commitments with long remaining lives; old dirty capital has depreciated through most of its fuel exposure and is partially hedged by competitive price pass-through. Under empirically relevant conditions, the model predicts that the emissions–age interaction is negative: the carbon premium is steepest for firms with young capital and attenuates toward zero for firms with old capital. Capital-light sectors, where irreversibility is weak, serve as a built-in placebo: the model predicts a flat, near-zero carbon premium at all ages.

We find support for market pricing of stranding risk in two settings. Our first empirical setting exploits the US electric utility sector, which offers two advantages. First, utilities have been at the front line of US climate policy, from the Clean Air Interstate Rule through the Clean Power Plan to the EPA's 2024 power-plant standards, providing a rich calendar of events that shift the probability or reality of cost-side regulation. Second, the EIA's generator-level data let us observe, for each publicly traded utility, the capacity-weighted age of its fossil-fuel generators and its renewable generators separately: precisely the characteristics the model identifies as driving stranding exposure. In event studies around 19 cost-side regulatory events over 2007–2024, firms with younger fossil fleets lose significantly more value when regulation tightens; sorting firms by fossil-fuel generator age, the signed Old−Young return spread is positive and significant. When firms are instead sorted by the age of their renewable generators, the spread is indistinguishable from zero. This is the model's clean-side placebo: renewable generators have no fuel-cost exposure, so their vintage should be irrelevant for stranding.

We then turn to the broad stock market to test whether the mechanism generalizes beyond utilities. We measure emissions using the Trucost database and proxy capital age with a recursive investment-based measure from Compustat (Lin et al. (2020)). The model predicts that the emissions-age interaction should be strongest in capital-intensive industries, where putty-clay irreversibility binds, and weakest in capital-light industries, where it does not. This is what we find. In pooled cross-sectional regressions, the carbon premium among capital-intensive firms is strongly decreasing in capital age: firms in the youngest age quintile

earn roughly 70 basis points per month more per standard deviation of emissions, declining monotonically to approximately zero in the oldest quintile. Capital-light firms display no significant carbon premium at any age. The model also predicts that stranding risk should be priced only when cost-side policy is credible. Consistent with this, the emissions-age interaction is concentrated in the post-Paris period and strengthens further around 2020-2021, coinciding with moments when the market's assessment of cost-side policy risk shifted discretely.

Our findings reconfirm the central result of Bolton and Kacperczyk (2021): firms with higher emissions earn higher expected returns, and this premium reflects risk compensation rather than mispricing. But our model and evidence qualify the finding in an important way. The carbon premium is not a uniform tax on dirty firms; it is a stranding premium whose magnitude depends on the vintage structure of the firm's capital. A high-emission firm with young capital, locked into fuel commitments with long remaining lives, commands a substantially higher premium than a high-emission firm with old capital whose fuel exposure has largely depreciated. In the oldest age quintile, the premium is indistinguishable from zero. The BK specification, which regresses returns on emissions levels without conditioning on capital age, estimates an average across these heterogeneous exposures. Our contribution is to show that this average masks economically large variation within the high-emission group, variation that is predicted by the theory and that identifies the structural source of the premium as cost-side policy risk operating through irreversible capital.

The finding that the stranding premium is concentrated in young capital has implications for the transition. Because the premium falls most heavily on recently installed dirty capital, it raises the cost of financing new emission-intensive projects relative to cleaner alternatives. Financial markets, in effect, penalize the flow of dirty investment, not merely the stock of existing dirty capital. Had the finding been reversed, with old dirty capital bearing the largest premium, the financial pressure would instead fall on holders of legacy assets, encouraging early retirement rather than cleaner installation at the margin. The distinction matters: one channel discourages new dirty investment, the other accelerates the phase-out of old dirty capital. Our evidence points to the former.

Importantly, our findings are robust to lagging the emissions measure by one year, addressing concerns about look-ahead bias in the timing of emissions data releases (Aswani et al., 2024; **?**), and to scaling emissions in a theory-consistent way by gross PP&E rather than using levels, which captures the fuel intensity of the installed capital stock rather than firm size. Both of these choices align with **?**, who show that the carbon premium is correctly identified using lagged emissions intensity and that its magnitude is sensitive to the inclusion of "super-emitters", notably electric utilities. Our model provides a structural account of why utilities appear anomalous: as the most capital-intensive sector with the oldest fleets, their capital sits in the region where the stranding premium attenuates toward zero.

This paper provides such a theory and the empirical evidence to support it. We argue that the carbon premium reflects the market pricing of *stranding risk*: the risk that climate policy will destroy the value of fossil-fuel-using capital that, once installed, cannot be repurposed. The finance literature lacks a model of how stranding risk varies across the cross-section of firms. We build one, grounded in the putty-clay tradition (Gilchrist and Williams, 2000):

machines are installed with a fixed fuel-to-labor ratio that cannot subsequently be adjusted. A carbon tax therefore falls on locked-in fuel commitments, and its incidence depends on both how dirty a firm's capital is and how *old* it is. The model's central prediction is an emissions × age interaction in returns: among firms with substantial physical capital, young dirty firms bear the largest stranding losses and therefore command the highest risk premia. Clean capital, which has zero fuel exposure by construction, exhibits no age gradient whatsoever, a built-in placebo that no alternative theory of "young capital is riskier" can replicate.

The mechanism turns on two competing forces within the model. A *growth wedge* arises because disembodied productivity growth causes revenue to decay more slowly with vintage than fuel costs do: young machines have locked in large fuel commitments with long remaining lives, while old machines' fuel exposure has largely run off. *Operating leverage* works against this because the wage bill does not decay with age, compressing old machines' profit margins and exposing them to tax shocks. In capital-intensive industries, where the labor share is low and general-equilibrium price adjustments reinforce the wedge, the growth effect dominates and young dirty capital suffers the most. We further develop a *taxonomy of stranding mechanisms* across policy instruments: the age × dirty interaction is specific to cost-side policies (carbon taxes, cap-and-trade, emissions regulations) that raise operating costs of installed capital. Green technology improvements and investment subsidies strand dirty capital through the output-price channel, producing weaker dependence on age. Strikingly, dirty investment taxes and ESG-driven cost-of-capital wedges can actually *benefit* existing dirty incumbents by reducing competitive entry.

We test these predictions in two empirical settings. In the utility sector, where EIA plant-level data allow direct measurement of capital vintage, we conduct event studies around 19 cost-side regulatory events over 2007–2024. The signed Old−Young cumulative abnormal return spread is significantly positive for fossil-fuel vintages: firms with older fossil fleets outperform those with younger fossil fleets when cost-side regulation tightens. The renewable-vintage spread, by contrast, is indistinguishable from zero, exactly the placebo the model predicts, since clean generators have no fuel-cost exposure. Cross-sectional regressions confirm that orthogonalized fossil fleet age is the dominant predictor of signed abnormal returns, robust to controls for firm size, leverage, and generation mix.

We extend the analysis to a broad Trucost–Compustat panel covering all industries. In pooled cross-sectional regressions of unconditional monthly returns, the emissions × age interaction is significantly negative in capital-intensive industries, with a coefficient on the emissions × age interaction of -0.0020 ($t = -3.34$). In capital-light industries, the interaction is indistinguishable from zero ($t < 1$), providing a second built-in falsification test. These results are robust across all emission scopes (Scope 1, 2, 3, and 1+2) and hold in both pooled OLS and Fama–MacBeth specifications. A pooled OLS event study following Gómez-Cram and Olbert (2023) yields a triple-interaction estimate (emissions × age × event day) of +24 basis points ($t = 2.70$) in capital-intensive industries, with no significant effect in capital-light industries or around subsidy events.

These findings matter for three reasons. First, they supply the structural mechanism that the carbon-premium literature has lacked: the premium arises from the irreversibility of

physical capital, and it varies predictably across the cross-section (by emissions, capital age, and capital intensity) in ways that pin down stranding risk as the operative channel. Second, the return pattern around policy events is diagnostic of the *type* of transition risk the market is pricing. Cost-side regulations produce the age × dirty signature our model predicts; subsidies and investment restrictions do not. This taxonomy can be used by investors and policymakers to read the cross-section of event-day returns as a signal of which stranding channel is operative. Third, the finding that stranding risk falls disproportionately on young dirty capital has important implications for the energy transition. On one hand, the stranding premium encourages reallocation toward green investment through the standard cost-of-capital channel. On the other hand, by penalizing recent dirty investment most heavily, it discourages replacement of older, dirtier capital with newer, cleaner capital *within* fossil-fuel-consuming industries, exactly the capital-replacement margin that Bertolotti et al. (2025) show is central to the green transition. More broadly, by raising the cost of capital at the dirty technological frontier, stranding risk can redirect innovation toward clean technologies (Acemoglu et al., 2012), acting as an implicit tax on dirty-frontier investment that complements explicit carbon pricing in steering the direction of technical change.

Related literature. This paper contributes to the literature on climate risk and asset pricing. Engle et al. (2020), Bolton and Kacperczyk (2021), Bolton and Kacperczyk (2023), and Hsu et al. (2023) document a positive return–emissions relation. The existence and interpretation of this carbon premium are debated: Aswani et al. (2024) argue that the association is driven by vendor-estimated (rather than firm-disclosed) emissions and by unscaled emissions correlated with firm size, while Zhang (2025) shows that after accounting for data-release lags, realized carbon returns turn negative in the U.S. and are concentrated in the post-Paris period. Crosignani et al. (2025) reconcile conflicting findings by showing that the premium is driven by expected emissions and by the post-Climate Accord regime. Our emissions × age interaction addresses these concerns directly: the interaction is a *within-emission-level* predictor that cannot be explained by size, data timing, or level effects, and the capital-light placebo rules out confounds that operate across all industries.

Gârleanu and Pedersen (2025) and Barnett et al. (2020) study how climate risk is priced in general equilibrium; Pástor et al. (2021) and Pedersen et al. (2021) model ESG preferences as affecting expected returns. Chittaro et al. (2025) show that in frictionless markets, pollution premia provide incentives equivalent to a carbon tax, and note that vintage capital is a key feature of decarbonization paths. Our contribution is to show that the premium has a *within-brown* structure, an emissions × age interaction, that these models do not generate, and to provide a structural mechanism (putty-clay irreversibility) that explains it. Acharya et al. (2025) build a GE model of the energy sector in which incumbents with developed reserves respond differently from entrants; we differ in scope (all industries, not just energy) and in the source of heterogeneity (capital *age* rather than reserve composition). Lin et al. (2020) document an unconditional old-capital-age premium driven by technology-frontier shocks. Our result is the opposite *conditional* sign: within dirty, capital-intensive firms, young capital earns higher returns. The two channels coexist: the unconditional old-age premium operates across all firms, while stranding risk reverses it for the dirtiest, and the capital-light placebo confirms our channel is distinct. Geelen et al. (2024) show that capital age drives leverage

and debt maturity; we show it drives risk premia through an entirely different mechanism. On the real side, the model builds on the putty-clay framework of Gilchrist and Williams (2000), its application to climate policy in Gilchrist et al. (2026), and Gourio (2011), who shows that putty-clay irreversibility generates endogenous stock-market volatility by creating a wedge between the value of installed and new capital. Lanteri and Rampini (2022) study irreversible investment under climate-policy uncertainty but do not derive cross-sectional asset-pricing implications. Bertolotti et al. (2025) study optimal replacement subsidies with vintage-dependent environmental externalities; our market-pricing results complement their normative analysis. Our taxonomy of stranding mechanisms connects to the literature on directed technical change and climate policy (Acemoglu et al., 2012; Aghion et al., 2016).

2 Model

We build a putty-clay model with embodied productivity growth in which firms differ in the fossil-fuel intensity of their capital. Machines are installed with fixed factor ratios and cannot be retrofitted; newer machines embody superior technology. An unanticipated carbon tax strands dirty capital, and the severity of stranding depends on both the dirtiness and the age of the machine. We derive three propositions: one signing the partial-equilibrium comparative statics of stranding losses, one characterizing general-equilibrium pass-through, and one connecting policy exposure to risk premia and the cross-section of expected returns.

2.1 Technology

A continuum of intermediate-good firms is indexed by *dirtiness* $\phi \in [0, 1]$. Each firm operates a fleet of machines (*vintages*). A machine installed at date s by firm ϕ is designed with fixed capital ratio $k_s(\phi)$ and fuel ratio $f_s(\phi)$; if operated, it uses one unit of labor, k_s units of capital services, and f_s units of fossil fuel. These ratios are chosen at installation and cannot subsequently be adjusted (*putty-clay*).

Machine-level output is Cobb–Douglas:

$$y_{s,t}(\phi) = \theta_s \, Z_t \, k_s(\phi)^{\alpha_k(\phi)} \, f_s(\phi)^{\alpha_f(\phi)}, \tag{1}$$

with $\alpha_k(\phi) + \alpha_f(\phi) + \alpha_\ell = 1$. We assume $\alpha_f(0) = 0$ and $\alpha_f'(\phi) > 0$, so the clean boundary has zero fossil-fuel use and dirtiness monotonically increases fuel intensity. The capital share $\alpha_k(\phi) = 1 - \alpha_\ell - \alpha_f(\phi)$ is decreasing, and the labor share α_ℓ is common across firms *within a sector*. A common within-sector labor share ensures that all firms in a given sector share a single balanced-growth rate. When we compare capital-intensive and capital-light industries below, we interpret the model as a family of sectoral economies indexed by sector-specific labor shares: lower α_ℓ sectors correspond to more capital-intensive industries in the data.

Machine-level TFP has two components. *Embodied* productivity θ_s grows at rate $g_\theta \geq 0$ and is fixed at the date of installation: newer vintages are more productive. *Disembodied* aggregate TFP Z_t grows at rate $g_z \geq 0$ and benefits all machines regardless of vintage. Machines die deterministically at age $A > 0$.

A competitive final-good producer aggregates intermediates with CES technology. Each intermediate firm takes its output price $P(\phi)$ as given; the specific elasticity of substitution does not enter the partial-equilibrium comparative statics below, though it matters for the magnitude of general-equilibrium pass-through.

2.2 Balanced growth path

On the balanced growth path (BGP), embodied productivity grows at g_θ, aggregate TFP at g_z, and installed factor ratios grow at the common rate

$$g_y = \frac{g_\theta + g_z}{\alpha_\ell}. \tag{2}$$

This follows from balanced growth: output per frontier machine must grow at g_y to match the wage, giving $g_y \alpha_\ell = g_\theta + g_z$. The level interest rate is r. Relative output prices $P(\phi)$ and the fuel price p^f are constant on the BGP; the wage $w_t = w_0 e^{g_y t}$ grows at the output rate.

An age-a machine has installed ratios $k^* e^{-g_y a}$ and $f^* e^{-g_y a}$, where k^* and f^* are the current frontier values. Its output decays at the *net vintage decay rate*

$$\tilde{g} \equiv g_y - g_z = \frac{g_\theta + g_z(1-\alpha_\ell)}{\alpha_\ell}, \tag{3}$$

which is slower than g_y because disembodied TFP Z_t boosts old machines equally. Define frontier output per machine as $x^* \equiv \theta_0 Z_0 (k^*)^{\alpha_k} (f^*)^{\alpha_f}$. Output at age a is $x^* e^{-\tilde{g}a}$, decaying at the net vintage rate $\tilde{g}$. Fuel use at age a is $f^* e^{-g_y a}$, decaying at the faster rate g_y. The gap between revenue decay ($\tilde{g}$) and fuel-cost decay (g_y) when $g_z > 0$ is the *growth wedge* that drives many of the results below.

We define the annuity factor $\mathcal{A}_r(a) \equiv (1 - e^{-r(A-a)})/r$, measuring the present value of a unit flow over remaining life $A - a$ discounted at rate r. On the BGP, each firm installs machines at a constant rate, and with deterministic death at A the within-firm age distribution is uniform on $[0, A]$.

2.3 Installation and free entry

At installation, the firm chooses k_s and f_s to maximize the machine's net present value. Defining frontier output at installation as $x_s \equiv \theta_s Z_s k_s^{\alpha_k} f_s^{\alpha_f}$, output at age a is $x_s e^{g_z a}$ (only disembodied TFP $Z_{s+a} = Z_s e^{g_z a}$ varies with age), so

$$\max_{k_s, f_s} \left\{ \int_0^A e^{-ra} \Big[P(\phi)\, x_s\, e^{g_z a} - w_s\, e^{g_y a} - p^f f_s \Big]\, da \;-\; q_k\, k_s \right\}. \tag{4}$$

The Cobb–Douglas first-order conditions yield:

$$p^f f^*(\phi) = \alpha_f(\phi)\, P x^*\, \Lambda, \qquad \Lambda \equiv \frac{\mathcal{A}_{r-g_z}(0)}{\mathcal{A}_r(0)} \geq 1, \tag{5}$$

$$q_k\, k^*(\phi) = \alpha_k(\phi)\, P x^*\, \mathcal{A}_{r-g_z}(0), \tag{6}$$

where Λ is a forward-looking correction: disembodied growth makes the machine's output grow over its life, raising fuel's marginal revenue product above its flow cost, so $\Lambda \geq 1$ with equality when $g_z = 0$.

Competitive machine investment drives the NPV of a new machine to zero. Substituting the first-order conditions and the free-entry condition $v^0(0, \phi) = q_k k^*(\phi)$, the α_f and α_k terms cancel, yielding

$$P(\phi)\, x^*(\phi) = \frac{w\, \mathcal{A}_{r-g_y}(0)}{\alpha_\ell\, \mathcal{A}_{r-g_z}(0)}, \tag{7}$$

where $w \equiv w_0$. Because α_ℓ, w, r, g_y, g_z, and A are all common within a sector, frontier gross revenue per machine Px^* is independent of ϕ: $(Px^*)'(\phi) = 0$ (Lemma 1, Appendix A). All heterogeneity in α_f and α_k is absorbed by the cost structure: dirtier firms spend more on fuel and less on capital, but frontier revenue per machine is the same. Free entry forces the labor claim $\alpha_\ell P x^*$ to equal the same present-value wage bill for every variety. The objects derived above characterize the pre-tax BGP. We now introduce a carbon tax as an unanticipated shock to this steady state and derive its incidence across the existing capital stock.

2.4 Operating profits and machine value

A carbon tax $\tau \geq 0$ is levied per unit of fuel. The detrended operating profit of an age-a machine at firm ϕ is

$$\pi(a, \phi; \tau) = Px^* e^{-\tilde{g}a} - (p^f + \tau) f^*(\phi) e^{-g_y a} - w. \tag{8}$$

Revenue decays at $\tilde{g}$; fuel cost decays at $g_y > \tilde{g}$ (when $g_z > 0$); the wage w is constant in detrended terms.

The detrended present value of an age-a machine is

$$v^0(a, \phi) = Px^* e^{-\tilde{g}a}\, \mathcal{A}_{r-g_z}(a) - p^f f^*(\phi) e^{-g_y a}\, \mathcal{A}_r(a) - w\, \mathcal{A}_{r-g_y}(a). \tag{9}$$

The three terms use different annuity factors reflecting their different time profiles: revenue is discounted at $r - g_z$ (it grows at g_z in levels), fuel at r (constant in levels), and wages at $r - g_y$ (growing at g_y in levels).

2.5 Asset Stranding

The *partial-equilibrium asset stranding* from a marginal carbon tax increase is

$$\mathcal{S}^{PE}(a, \phi) \equiv -\frac{1}{v^0} \frac{\partial v^0}{\partial \tau}\bigg|_{\tau=0} = \frac{f^*(\phi)\, e^{-g_y a}\, \mathcal{A}_r(a)}{v^0(a, \phi)}, \tag{10}$$

the present value of fuel exposure as a fraction of machine value.[1]

The central theoretical result signs the derivatives of stranding with respect to dirtiness, age, and their interaction.

[1] All subsequent results are stated on the positive-value domain $\mathcal{D} \equiv \{(a, \phi) \in [0, A) \times [0, 1] : v^0(a, \phi) > 0\}$. For log-derivative statements we further restrict to the interior set $\mathcal{D}_+ \equiv \{(a, \phi) \in \mathcal{D} : f^*(\phi) > 0\}$, which under $\alpha_f(0) = 0$ and $\alpha_f'(\phi) > 0$ corresponds to the dirty interior $\phi > 0$.

Proposition 1 (Stranding losses)**.** *Assume $\alpha_f(0) = 0$ and $\alpha_f'(\phi) > 0$ on $[0, 1]$. Under free entry, on the interior domain $\mathcal{D}_+$:*

(i) ***Dirtiness.*** *$\partial_\phi \log \mathcal{S}^{PE}(a, \phi) > 0$ for all $(a, \phi) \in \mathcal{D}_+$. Dirtier firms face strictly higher PE stranding at every age.*

(ii) ***Sign equivalence.*** *At each age $a \in [0, A)$:*

$$\operatorname{sgn}\Big(\partial_a \log \mathcal{S}^{PE}\Big) = \operatorname{sgn}\Big(\partial^2_{a\phi} \log \mathcal{S}^{PE}\Big) = -\operatorname{sgn}\ B(a),$$

where

$$B(a) \equiv \underbrace{Px^* e^{-\tilde{g}a}\Big[\mathcal{A}_{r-g_z}(a) - \mathcal{A}_r(a)\Big]}_{\text{growth-wedge component} \geq 0 \text{ when } g_z > 0} + \underbrace{w\Big[\mathcal{A}_r(a) - \mathcal{A}_{r-g_y}(a)\Big]}_{\text{operating-leverage component} \leq 0}. \tag{11}$$

Moreover, $\partial^2_{a\phi} \log \mathcal{S}^{PE} = \mathcal{M}(a, \phi) \cdot \partial_a \log \mathcal{S}^{PE}$ with amplification factor

$$\mathcal{M}(a, \phi) = \frac{p^f (f^*)'(\phi)\, e^{-g_y a}\, \mathcal{A}_r(a)}{v^0(a, \phi)} > 0, \tag{12}$$

which depends on fuel-intensity dispersion $(f^*)'$ but not on B.

(iii) ***Single crossing.*** *When $g_z > 0$, $r > g_y$, and*

$$\frac{1}{\alpha_\ell}\left(1 - \frac{\mathcal{A}_r(0)}{\mathcal{A}_{r-g_z}(0)}\right) > 1 - \frac{\mathcal{A}_r(0)}{\mathcal{A}_{r-g_y}(0)}, \tag{13}$$

then $B(0) > 0$ and there exists a unique $a^* \in (0, A)$ such that

$$B(a) > 0 \text{ for } a \in [0, a^*), \qquad B(a) < 0 \text{ for } a \in (a^*, A).$$

Young dirty machines are the most exposed for $a < a^*$ (growth wedge dominates); old dirty machines are the most exposed for $a > a^*$ (operating leverage dominates near end of life).

(iv) ***Clean-side placebo.*** *At the clean boundary $\phi = 0$, $f^*(0) = 0$ and $\mathcal{S}^{PE}(a, 0) = 0$ at every age. This exact zero is a partial-equilibrium statement about the direct fuel-cost channel.*

All proofs are in Appendix A.

The two forces. The sign of the age gradient and the cross-partial both reduce to the sign of a single scalar function $B(a)$, which decomposes into two competing forces. The *growth wedge* ($g_z > 0$): revenue decays at $\tilde{g} < g_y$, so old machines' revenue holds up better than their fuel cost. This makes old machines relatively valuable compared to their fuel exposure, so *young* machines are more exposed as a fraction of value. *Operating leverage*: the wage w is constant in detrended terms and does not respond to the carbon tax. It compresses old machines' profit margins more heavily (their revenue has decayed but the wage has not), making *old* machines disproportionately squeezed. When the growth wedge dominates ($B > 0$), young dirty machines bear the highest stranding losses; when leverage dominates ($B < 0$), old dirty machines do.

Intuition for the growth wedge. The growth wedge has a transparent economic logic. Fuel costs decay at rate g_y as machines age; older machines were installed further back on the frontier and therefore locked in smaller fuel ratios in current detrended units. Revenue decays at the slower rate $\tilde{g} = g_y - g_z$, because disembodied TFP Z_t boosts old machines' output even as their embodied productivity falls behind. The stranding ratio $\mathcal{S}^{PE}$ is the present value of fuel exposure divided by machine value. For young machines, fuel exposure is large (they locked in high fuel ratios with long remaining lives) while their value includes a revenue stream that holds up better than fuel costs. For old machines, fuel exposure has largely run off, and what value remains is increasingly dominated by the wage claim. When the gap $g_y - \tilde{g} = g_z$ is large enough, the numerator of $\mathcal{S}^{PE}$ decays faster with age than the denominator, making young machines more exposed. Operating leverage works against this: the constant wage erodes old machines' margins, raising $\mathcal{S}^{PE}$ for old machines. The sign of $B(a)$ determines which force wins.

Embodied versus disembodied growth. The distinction between embodied and disembodied growth matters. With pure embodied growth ($g_z = 0$), the growth wedge is absent: revenue and fuel cost decay at the same rate g_y, so the first term of B vanishes and only leverage remains. In that case $B(a) < 0$ for all a, so old dirty machines are the most exposed. Adding disembodied growth ($g_z > 0$) creates a wedge between the way revenue and fuel commitments age. This wedge favors young machines because their fuel commitments have the longest remaining lives. Higher g_z strengthens that wedge. Higher g_θ, holding g_z fixed, raises the common growth rate g_y and strengthens leverage without creating a new wedge. Put differently, disembodied growth is what makes young dirty capital relatively exposed in the baseline model; setting $g_\theta = 0$ alone does not eliminate the age interaction if $g_z > 0$ persists.

The role of capital intensity. Condition (13) is easier to satisfy in sectoral calibrations with low α_ℓ (less operating leverage), high g_z/g_θ (more disembodied growth), and moderate machine life A. The labor share plays a double role. First, it determines the strength of operating leverage: the wage term in $B(a)$ is proportional to w, which scales with α_ℓ through free entry. Low α_ℓ shrinks this term directly. Second, in more capital-intensive sectors, physical capital is a large share of firm value, so the putty-clay irreversibility binds most tightly. We therefore interpret the capital-intensity prediction as a *cross-sector* comparative static: within any one sector α_ℓ is common across firms, but sectors with lower common labor shares correspond to more capital-intensive industries in the data. Even when the PE growth wedge does not fully dominate (that is, even when B crosses zero at some $a^* < A$), the crossing age a^* is largest where α_ℓ is smallest, making the young-dirty region widest in low-labor-share sectors.

The clean-side placebo. The exact zero on the clean side in Part (iv) is a PE statement about the direct fuel-cost channel: any confound that generates fuel-cost stranding for all capital (e.g., a generic "young capital is riskier" story) would contradict $\mathcal{S}^{PE} \equiv 0$ at $\phi = 0$. Once GE pass-through is admitted, clean machines can still move through common price and wage channels, so the empirical benchmark is weaker age dependence on the clean side

rather than a literal structural zero. The event-study design in Section 3 exploits this by comparing the age gradient across fossil and renewable generating assets within the same firms and events.

Illustrative calibration. Figure 1 illustrates the shape of $B(a)$ and the implied stranding exposure for one calibration. Panel (a) plots the normalized age gradient $B(a)/B(0)$. With pure disembodied growth ($\eta = 1$) and substantial fuel intensity, the growth wedge dominates for young machines ($a < a^*$) and operating leverage takes over beyond $a^* \approx 7$ years. Panel (b) translates this into the fractional stranding exposure $\mathcal{S}^{PE}(a)/\mathcal{S}^{PE}(0)$, expressed as a percentage deviation from the newborn-machine benchmark, and overlays the empirical distribution of recursive capital age among capital-intensive Compustat firms in 2022–24. The key feature is that $\mathcal{S}^{PE}$ peaks within the empirically relevant age range: at the median firm age of roughly 4.7 years, exposure is near its maximum and declining, so that older capital bears higher fractional stranding losses. This is the mechanism behind the negative emissions × age interaction estimated in Section 2.6. Eighty percent of firms sit to the left of a^*, squarely in the region where the operating-leverage channel strengthens with age.

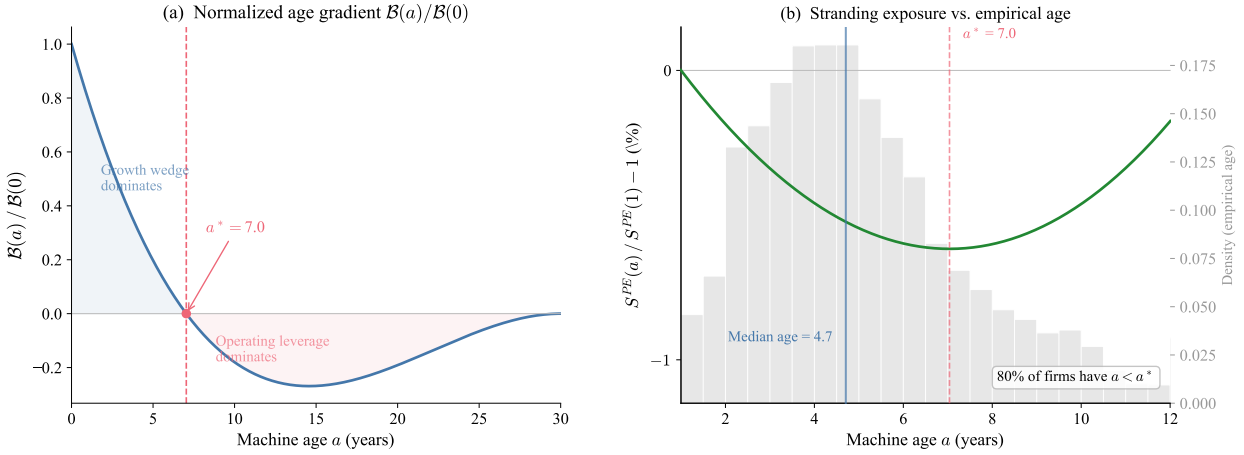


Figure 1: Calibrated age gradient and stranding exposure.

Notes: Panel (a) plots $B(a)/B(0)$, the normalized age gradient of machine-level stranding exposure defined in Proposition 1. The blue-shaded region ($B > 0$) corresponds to ages where the growth-wedge effect dominates; the red-shaded region ($B < 0$) to ages where operating leverage dominates. Panel (b) plots the fractional stranding exposure $\mathcal{S}^{PE}(a)/\mathcal{S}^{PE}(0) - 1$ in percent, overlaid on the empirical distribution of recursive capital age (Lin, Palazzo, and Yang 2020) for capital-intensive firms in Compustat, 2022–2024. The vertical blue line marks the empirical median age; the dashed red line marks a^*. Parameters: $\eta = 1$ (pure disembodied growth), $A = 30$ years, $\alpha_\ell = 0.30$, $\alpha_f = 0.50$, $r = 0.02$, $g_y = 0.03$. Capital-intensive firms are defined by the balanced SIC-2 median PPEGT/AT classification described in Section 2.6.

General equilibrium. In partial equilibrium, operating leverage works only through the fuel-cost channel, compressing old machines' margins. In general equilibrium, the output price and wage also respond to the tax, introducing two additional channels that both reinforce the young-dirty prediction. Old machines, with thin margins, have revenue and

wages that are large multiples of value; when the output price rises ($\hat{P} > 0$) or the wage falls ($\hat{w} < 0$), old machines benefit disproportionately. Both GE channels therefore partially offset the direct fuel-cost hit for old machines, making young-dirty dominance more likely than the PE analysis alone suggests. The exact clean-side zero of Part (iv) does not survive in GE: clean machines can move through common price and wage channels, so the empirical benchmark is weaker age dependence on the clean side rather than a literal zero. Appendix C states the formal GE decomposition (Proposition 4) and derives the corrections to each derivative.

2.6 Risk premia and the cross-section of expected returns

The policy-exposure comparative statics carry over directly to expected returns under a simple pricing structure.

Proposition 2 (Risk premia)**.** *Suppose a permanent carbon tax of size τ arrives at Poisson rate $\lambda > 0$, and the stochastic discount factor jumps by a factor $1 + \xi$ ($\xi > 0$) upon arrival. Let*

$$S(a, \phi) \equiv -\frac{1}{V(a, \phi; 0)} \frac{\partial V(a, \phi; \tau)}{\partial \tau}\bigg|_{\tau=0}$$

denote the relevant semi-elasticity of value to the tax shock. In PE, $S(a, \phi) = \mathcal{S}^{PE}(a, \phi)$; with price and wage pass-through, $S(a, \phi) = S^{GE}(a, \phi)$. To first order in τ:

$$E[r^e(a, \phi)] \approx \lambda \xi \tau S(a, \phi). \tag{14}$$

Consequently,

$$\partial_\phi \log E[r^e] \approx \partial_\phi \log S, \qquad \partial_a \log E[r^e] \approx \partial_a \log S, \qquad \partial^2_{a\phi} \log E[r^e] \approx \partial^2_{a\phi} \log S. \tag{15}$$

Under PE, all sign results of Proposition 1 transfer directly. Under GE, the same pricing mapping applies to S^{GE}, so the age and interaction signs combine the PE forces from Proposition 1 with the pass-through corrections in Proposition 4.

Proof. See Appendix A.

Firm-level aggregation. In PE, a firm of type ϕ with age density $n(a)$ has direct fuel-cost exposure

$$\mathcal{S}^{firm}(\phi) = \frac{\text{FE}(\phi)}{V(\phi)} = \frac{\int_0^A f^*(\phi)\, e^{-g_y a}\, \mathcal{A}_r(a)\, n(a)\, da}{\int_0^A v^0(a, \phi)\, n(a)\, da}, \tag{16}$$

the ratio of the firm's total dollar fuel exposure to its value. The numerator is multiplicatively separable: $f^*(\phi)$ times a ϕ-independent integral. Hence $\partial_\phi \mathcal{S}^{firm} > 0$ unconditionally: dirtier firms have higher PE fuel exposure. For the age gradient, the dollar fuel exposure $D(a) \equiv f^* e^{-g_y a} \mathcal{A}_r(a)$ is *unconditionally* decreasing in a: younger fleets have higher total fuel exposure regardless of the sign of B. Whether this translates to higher *percentage* stranding depends on whether fuel exposure declines faster than firm value when the fleet ages, a condition reinforced by GE price and wage adjustments (Proposition 4). Under GE, the pricing-relevant firm-level object adds the corresponding aggregated revenue and wage pass-through terms in the obvious way. Cross-firm variation in $\bar{a}_i$ arises from heterogeneous investment histories that the model treats as exogenous initial conditions.

Cross-sectional regression. Linearizing the relevant exposure object ($\log \mathcal{S}^{PE}$ in PE or $\log S^{GE}$ once pass-through is included) around a reference point and aggregating across the firm's vintage portfolio, the model suggests a firm-level regression of the form

$$E[r_i^e] \approx c + \beta_1 \, \phi_i + \beta_2 \, \bar{a}_i + \beta_3 \, (\phi_i \times \bar{a}_i), \tag{17}$$

where $\bar{a}_i$ is the value-weighted average age of firm i's capital stock. The model signs the coefficients as follows:

1. $\beta_1 > 0$: dirtier firms earn higher expected returns (the carbon premium). This is unconditional.
2. $\text{sgn}(\beta_3)$ depends on the balance of forces. In PE, β_3 is governed by the sign of B, which reflects the growth-wedge versus leverage tradeoff (Proposition 1). In GE, price pass-through and wage decline both push β_3 toward negative values (Proposition 4). The interaction should be strongest in low-labor-share sectoral calibrations, which we map to capital-intensive industries in the data.
3. For clean firms ($\phi \approx 0$), the direct fuel-cost component is zero at all ages. With GE spillovers, residual clean-side age dependence can arise through common price and wage channels, but it should be much weaker than on the dirty side.
4. Equation (12) is exact at the machine level. After aggregation and linear projection it becomes a qualitative amplification result rather than an exact coefficient restriction: large fuel-intensity dispersion can generate a sizable interaction even when the reduced-form age main effect is small.

Reinforcing channels. Two additional forces reinforce the GE channels. First, variable utilization (Appendix B): old machines near the shutdown boundary reduce operations, limiting their fuel exposure. Young machines, far from shutdown, bear the full tax burden. Second, the combination of weak leverage (low α_ℓ in the sectoral calibrations we map to capital-intensive industries) and GE price and wage adjustments (Proposition 4) tips the balance toward young-dirty dominance precisely where the model predicts the interaction should appear. The PE growth wedge brings the age gradient close to zero in capital-intensive industries; the GE corrections and utilization close the remaining gap.

Alternative policies. The age $\times$ dirty interaction is specific to cost-side instruments (carbon taxes, cap-and-trade, emissions regulations) that raise operating costs proportionally to fuel intensity (Appendix D). Green technology improvements and subsidies strand dirty capital indirectly through the revenue and wage channels, which have long-duration annuity structures closely matched to the machine's own value duration; the resulting age dependence is weak. Dirty investment taxes and ESG-driven cost-of-capital wedges can actually *benefit* existing dirty incumbents by reducing competitive entry, reversing the sign of the value change. The cross-sectional return pattern around a policy event is therefore diagnostic of the type of stranding mechanism the market is pricing: a strong age $\times$ dirty interaction signals cost-side risk; its absence points to technology, subsidy, or investment-channel mechanisms. This taxonomy motivates our focus on cost-side regulatory events in Sections 3 and 4. While

the Propositions hold for any firm with putty-clay technology, the empirical bite of the mechanism depends on physical capital being a substantial share of firm value. In capital-light industries, where intangible assets dominate, stranding losses on installed equipment are negligible relative to total firm value. The capital-intensity split therefore serves as a built-in scope test in Section 4.

From model to data. In summary, the model identifies two competing forces, the growth wedge and operating leverage, whose balance determines the sign of the age $\times$ dirty interaction. The testable predictions are:

1. $\beta_1 > 0$: dirtier firms earn higher expected returns (unconditional).
2. The emissions $\times$ age interaction concentrates in *capital-intensive* industries, consistent with sector-level low-α_ℓ comparative statics and because physical capital is quantitatively important there. In capital-light industries, the putty-clay channel should be much weaker, providing a cross-sector scope test.
3. The age gradient should be much weaker for clean capital ($\phi = 0$) than for dirty capital because the direct fuel-cost channel is absent on the clean side.
4. A note on signs: in event time, older dirty machines earn *higher* returns when regulation tightens (they lose less). In the unconditional cross-section, younger dirty machines earn *higher* expected returns (compensation for crash-state exposure). These are consistent when cost-side stranding shocks are contractionary ($\Delta C < 0$): the risk premium decays with capital age as fuel exposure diminishes.

We now take these predictions to the data: first to the utility sector, where plant-level vintage data provide a direct test of age patterns across fossil and renewable generators (Section 3), and then to the broad cross-section of returns (Section 4).

3 Utility-Sector Event Study

We begin by testing the model's predictions in the sector where they are sharpest. Plant-level data from the U.S. Energy Information Administration (EIA) allow us to observe the vintage composition of each utility's generating fleet directly, providing a close empirical counterpart to the model's capital age. We exploit 19 cost-side regulatory events (EPA rules and Supreme Court decisions that directly govern power plant emissions) and test whether the model's vintage predictions hold in narrow event windows where confounding variation is limited.

3.1 Setting and data

Our sample consists of 48 publicly traded electric utilities matched to plant-level data from EIA Form 860 over the period 2005–2024. For each firm, we compute capacity-weighted average fleet age for fossil, renewable (solar and wind only, excluding hydro and nuclear), and total generating assets using the most recent annual EIA vintage observation available

at each event date. Daily stock returns come from CRSP. The EIA–CRSP crosswalk links utility holding companies to CRSP permanent identifiers (PERMNOs) via Compustat, with explicit handling of four major utility mergers during the sample period (Progress Energy–Duke Energy, Allegheny Energy–FirstEnergy, Constellation Energy–Exelon, and Dynegy–Vistra). Full details on sample construction, vintage measures, and orthogonalization are in Appendix H.

Orthogonalized vintage. A firm with an overall old fleet tends to have both old fossil and old renewable assets. To isolate the *compositional* variation (a firm that has selectively modernized its renewable fleet while retaining old fossil plants), we orthogonalize each technology-specific age on total fleet age in the event-date cross-section. The orthogonalized fossil vintage has a cross-sectional correlation of 0.77 with raw fossil age, confirming that the rank ordering is preserved.

Event selection. We assemble a calendar of candidate events spanning EPA regulations and Supreme Court decisions over 2006–2024. From this universe, we retain the 19 events that satisfy two criteria: (i) the event has a direct, plausible channel to the cost of operating fossil-fuel power plants, and (ii) the event is not rendered uninterpretable by nearby confounder that might specifically affect stock prices in the utility sector (for example, one event is roughly contemporaneous with the Fukushima nuclear disaster). Each event is signed $d_e \in +1, -1$ for whether it tightens or loosens regulation. Appendix F lists all events with dates, directions, and quality assessments.

3.2 Portfolio event study

For each event, we sort firms into terciles by orthogonalized fossil (or renewable) vintage and form value-weighted "Old" and "Young" portfolios. The signed Old$-$Young return spread, $d_e \times (R^O_{e,t} - R^Y_{e,t})$, measures how much more old-vintage firms gain (or lose less) than young-vintage firms when cost-side regulation tightens. The model predicts this spread is positive for fossil vintage and near zero for renewable vintage, since the direct fuel-cost channel is absent on the clean side (Section 2.5). We cumulate the cross-event average and report t-statistics based on the full cross-event covariance matrix, which accounts for both serial correlation in average abnormal returns and cross-event dependence (Appendix H).

Figure 2 presents the results. Panel (a) plots the signed cumulative average abnormal return (CAAR) spread for fossil vintage across the 19 cost-side events. The spread is economically small and statistically insignificant in the pre-event window, then rises sharply after day 0, reaching +0.73% by day +10 ($t = 3.28$). Firms with older fossil fleets significantly outperform those with younger fossil fleets when cost-side regulation tightens, precisely the pattern predicted by the model under growth-wedge dominance, where young machines' large remaining fuel commitments make them more exposed than old machines whose fuel exposure has largely run off.

Panel (b) shows the renewable placebo. The Old$-$Young spread for renewable vintage is +0.16% by day +10 ($t = 0.40$), economically small and statistically indistinguishable from

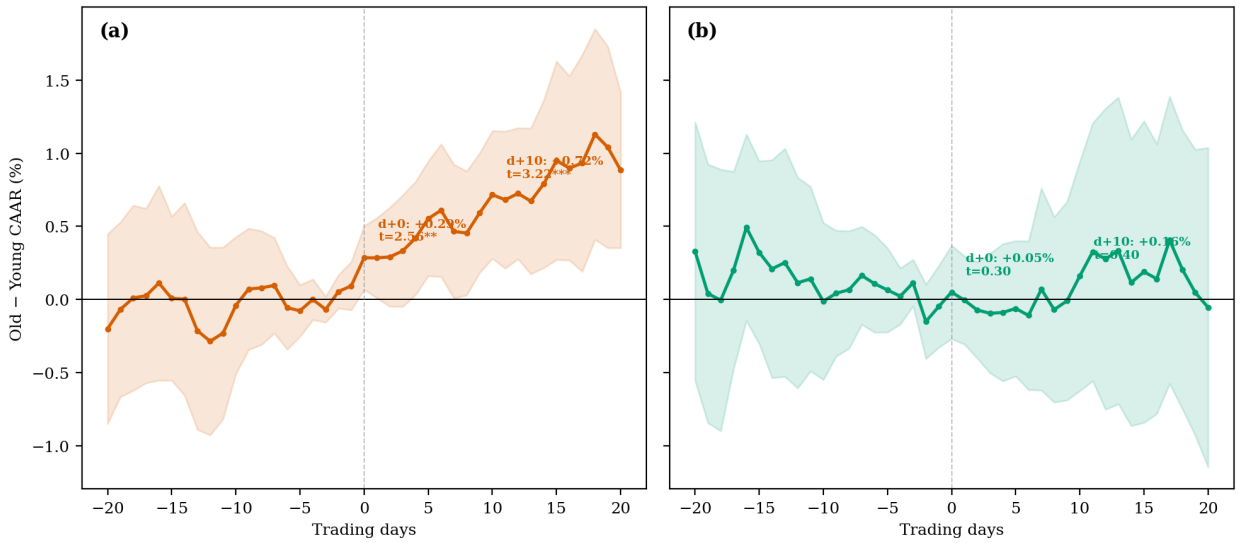


Figure 2: Old−Young CAAR spread around cost-side regulatory events.

Notes: Panel (a): fossil vintage (orthogonalized on total fleet age); Panel (b): renewable vintage (orthogonalized on total fleet age). Value-weighted tercile portfolios, 19 events, signed by regulatory direction. Shaded bands show 95% confidence intervals based on cross-event covariance matrix standard errors. The fossil spread is significantly positive (old fossil outperforms young fossil when regulation tightens); the renewable spread is indistinguishable from zero. The near-zero renewable spread is the model's clean-side benchmark: clean generators have no direct fuel-cost exposure to strand, so any age gradient should be much weaker.

zero. This is consistent with the model's clean-side benchmark: solar and wind generators have no direct fossil-fuel operating costs, so the carbon tax does not create differential fuel-cost exposure across vintages. The contrast between panels is a useful identification result: any confound that generates an age gradient for all capital would typically produce nonzero spreads in both panels. The near-zero renewable spread therefore supports the view that the fossil-vintage result operates through the fuel-cost channel, not through a generic age or investment-timing effect.

Robustness. The results are robust to (i) equal-weighted portfolios, (ii) median rather than tercile splits, (iii) BMP-standardized test statistics that down-weight high-volatility events, and (iv) restricting to $Q \leq 2$ (low-confounder) events. Leave-one-out analysis confirms that no single event drives the result: the fossil CAAR at day +10 remains significant at the 5% level when any individual event is dropped.

3.3 Cross-sectional regressions

The portfolio sorts document the direction and significance of the Old−Young spread but do not control for firm characteristics. We now estimate a regression in the spirit of the model's linearized return equation (17), with the dependent variable being the firm-level signed cumulative abnormal return (industry-adjusted, using the value-weighted utility sec-

Table 1: Cross-sectional CAR regressions: generator age

	Panel A: Dirty generator age			Panel B: Solar+Wind generator age		
	$(0, 5)$	$(0, 10)$	$(0, 15)$	$(0, 5)$	$(0, 10)$	$(0, 15)$
Generator age	$+0.0320^{**}$	$+0.0306^{*}$	$+0.0284$	$+0.0060$	-0.1389	-0.1577
	(2.23)	(1.71)	(0.82)	(0.05)	(-0.76)	(-0.74)
Generator age^2	$+0.0008$	$+0.0011$	$+0.0018$	-0.0216	$+0.0082$	$+0.0080$
	(0.74)	(0.78)	(0.80)	(-0.68)	(0.17)	(0.15)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Event FEs	Yes	Yes	Yes	Yes	Yes	Yes
N	744	744	744	101	101	101
R^2	0.113	0.110	0.111	0.378	0.467	0.441

Cross-sectional regressions of signed industry-adjusted CAR (in percentage points) on orthogonalized generator age. 19 cost-side events. Panel A: dirty generator age = capacity-weighted fossil fleet age, orthogonalized on total fleet age (per-event cross-section). Panel B: solar+wind generator age = capacity-weighted solar and wind fleet age, orthogonalized on total fleet age; restricted to firms with solar+wind capacity share $> 10\%$. Controls: fossil generation share, log generation capacity, log market capitalization, book-to-market ratio, and leverage. All specifications estimated by OLS with event fixed effects. t-statistics in parentheses are event-clustered. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

tor portfolio as benchmark) and the key regressor being orthogonalized fossil vintage:

$$\mathrm{CAR}_{i,e}(0, h) = \alpha_e + \beta_1 \, \mathrm{age}_{i,e}^{\mathrm{fossil},\perp} + \beta_2 \left(\mathrm{age}_{i,e}^{\mathrm{fossil},\perp}\right)^2 + \boldsymbol{\gamma}' \mathbf{X}_{i,e} + \varepsilon_{i,e}, \tag{18}$$

where α_e are event fixed effects and $\mathbf{X}_{i,e}$ includes fossil generation share, log generation capacity, log market capitalization, book-to-market, and leverage. The quadratic term allows for the hump-shaped stranding profile discussed in Section 2.5: the model's stranding ratio $\mathcal{S}^{PE}(a, \phi)$ is nonlinear in age due to the interaction of exponential decay in fuel exposure with the annuity structure of machine value. Standard errors are clustered by event.

Table 1 reports industry-adjusted estimates across three CAR windows $\{(0, 5), (0, 10), (0, 15)\}$. The coefficient on orthogonalized fossil vintage is positive and significant: older fossil fleets earn higher signed CARs when cost-side regulation tightens. The quadratic term allows for nonlinearity in the age–CAR relationship; it is individually insignificant, suggesting a predominantly linear age effect over the range of the data. The results strengthen when abnormal returns are computed using Fama-French three- and five-factor benchmarks (Appendix Tables OA.1 and OA.2), and are robust to firm-clustered and heteroskedasticity-robust standard errors, to dropping the quadratic term (Appendix Table OA.3), and to restricting the sample to the $Q \leq 2$ (low-confounder) events described in Appendix F.[2]

[2]With event fixed effects the industry- and market-adjusted benchmarks yield algebraically identical cross-sectional coefficients, since the benchmark return on each event-day is constant across firms and absorbed by the fixed effect.

4 Broad Market Evidence

We now extend the analysis from the utility sector to the broad stock market. In the utility event study, plant-level EIA data allowed direct measurement of capital vintage. In the broad cross-section, we proxy capital age with an accounting-based recursive measure and use Trucost Scope 1 emissions as the empirical counterpart of the model's dirtiness parameter ϕ.

4.1 Data and measurement

Our sample merges three datasets. Corporate carbon emissions come from S&P Trucost, which reports Scope 1 greenhouse gas emissions (in tons CO_2e) for approximately 4,200 publicly listed U.S. firms. We match these to monthly stock returns from CRSP and annual accounting data from Compustat. The merged panel covers January 2005 through December 2024 and contains approximately 366,000 firm-month observations after imposing the standard filters described below.

Capital age. In the broad panel we do not observe plant-level vintage directly. Following Lin et al. (2020) and Geelen et al. (2024), we proxy the model's investment-weighted average vintage $\bar{a}_i$ (equation (17)) with a recursive capital age measure:

$$\text{Age}_{i,t} = \frac{(1-\delta_j)\, K_{i,t-1}\,(\text{Age}_{i,t-1}+1) + I^{\text{gross}}_{i,t-1}}{K_{i,t}}, \tag{19}$$

where $K_{i,t} = (1-\delta_j)K_{i,t-1} + I^{\text{gross}}_{i,t-1}$. This is the surviving-capital-weighted average of the aged existing stock and new investment (with age 1). High values indicate old capital stocks; low values indicate recent investment. As a validation check, Appendix Figure 6 compares this accounting-based measure to the directly observed EIA generator age for electric utilities, the one sector where both are available. The two measures have a near-unit slope ($\hat{\beta} = 0.99$, $p < 0.001$), suggesting that the recursive measure tracks physical capital vintage reliably, at least in the subsample where ground truth is observable.

Emissions measures. To bring the model's dirtiness parameter ϕ_i to the data, we need an empirical counterpart. A natural measure is emissions per unit of productive capital. Consider the log decomposition:

$$\ln E_i = \ln\left(\frac{E_i}{\text{PPE}_i}\right) + \ln \text{PPE}_i, \tag{20}$$

where E_i is Scope 1 GHG emissions and PPE_i is gross PP&E. The first term is the intensive-margin counterpart of ϕ (emissions per unit of capital); the second scales with firm size. The *primary* specification follows Bolton and Kacperczyk (2021) and enters ln(Scope 1 GHG) with ln(market cap) as a separate control, which avoids ratio bias and lets the data determine the scaling exponent. As an alternative, we use Scope 1 GHG / gross PP&E (Compustat `ppegt`), which places a fully predetermined quantity in the denominator at the cost of requiring book-to-market as an additional control. We scale by gross PP&E rather than

revenue because the model's dirtiness parameter is fuel per unit of capital, not fuel per unit of output, and because revenue responds endogenously to a carbon tax through price pass-through, which would attenuate the measured interaction. We use gross rather than net PP&E because gross PP&E isolates tangible capital, the structural object in the model, without the intangibles, goodwill, and cash that inflate total assets, and because net PP&E would create a mechanical correlation with capital age (which is itself a function of accumulated depreciation on `ppent`). Both emissions and age variables are z-scored within each cross-section (calendar month) to ensure comparability of coefficients. Following Bolton and Kacperczyk (2021), all firm characteristics are winsorized at the 2nd and 98th percentiles; capital age is winsorized more conservatively at the 1st and 99th percentiles to preserve cross-sectional variation in the key regressor.

Validation. For the 45 utilities with matched EIA generator data, recursive capital age correlates positively with capacity-weighted physical generator age ($\beta = 0.99$, $r = 0.21$, $p < 0.001$; Appendix Figure 6), confirming that the accounting proxy captures genuine variation in the vintage composition of physical assets.

Capital-intensity classification. Industries are classified as "capital-intensive" or "capital-light" based on the median ratio of gross PP&E to total assets at the SIC-2 level, measured from Compustat over 1985–1995 (before the rise of intangible-intensive firms). The cutoff is chosen to balance the number of firm-month observations in each group (Appendix I). The model predicts that putty-clay irreversibility, and hence the bite of stranding risk, matters only where capital is physically substantial, making this split a built-in placebo.

4.2 Cross-sectional regressions

Our baseline regression is

$$r_{i,t} = \alpha + \beta_1 \,\text{Emissions}_{i,t} + \beta_2 \,\text{Age}_{i,t} + \beta_3 \,(\text{Emissions} \times \text{Age})_{i,t} + \Gamma' X_{i,t} + \text{FE} + \varepsilon_{i,t}, \qquad (21)$$

where $r_{i,t}$ is the monthly excess stock return (CRSP return minus the risk-free rate), and $X_{i,t}$ is a vector of controls. The coefficient β_1 captures the carbon premium of Bolton and Kacperczyk (2021); β_3 is the new object, the age $\times$ dirty interaction, that the model predicts should be negative when cost-side stranding risk is important and the growth-wedge term dominates. A negative β_3 means that, among dirty firms, those with younger capital earn higher returns, consistent with compensation for bearing greater stranding risk.

We control for the full set of firm characteristics used by Bolton and Kacperczyk (2021): ln(market cap), book-to-market, momentum, ROE, leverage, investment-to-assets, ln(PPE), sales growth, EPS growth, the Herfindahl–Hirschman index, market beta, and return volatility. Annual accounting variables are mapped to monthly returns using the Fama and MacBeth (1973) June $t+1$ convention, which ensures a minimum six-month gap between the fiscal year-end and the first return observation. Monthly controls (size, beta, volatility) are lagged by one month. All regressions include year-month fixed effects; the baseline also includes SIC-2 industry fixed effects.

We estimate pooled OLS with fixed effects absorbed via iterated Frisch–Waugh demeaning. Standard errors are two-way clustered by firm and calendar year, following Petersen (2009). This clustering accounts for both persistent firm-level characteristics and common time-series shocks to returns.

Specification progression. Table 3 presents a seven-column progression that builds from the Bolton and Kacperczyk (2021) replication (column 1) to our full specification (column 4) and then tests robustness. Column (1) enters emissions alone with all controls and year-month + SIC-2 fixed effects, replicating the baseline B&K specification. Column (2) enters capital age alone. Column (3) runs the horse race: emissions and age jointly, without the interaction. Column (4) adds the interaction term, our main result. Column (5) drops industry fixed effects to assess whether the interaction operates across or within industries. Column (6) drops all firm-level controls, retaining only the three treatment variables and fixed effects, to test whether the interaction is driven by omitted controls. Column (7) replaces separate year-month and SIC-2 fixed effects with fully interacted industry × year-month fixed effects, which absorb all time-varying industry-level shocks and identifies the coefficients purely from within-industry-month variation.

Where the putty-clay channel operates. Table 2 presents the main specification (emissions, capital age, and their interaction with the full set of B&K controls and year-month plus SIC-2 fixed effects) across two emissions measures, two time periods, and three industry samples.

Three patterns emerge. First, the emissions × age interaction concentrates in *capital-intensive* industries. In the all-industries sample, the interaction is negative but insignificant in the post-Paris period (Panel A: -0.0008, $t = -1.36$). Decomposing this into the two halves reveals the source: capital-intensive industries drive the entire signal (-0.0020, $t = -3.34$), while the capital-light subsample is an economically and statistically precise zero ($t < 1$ throughout). This is the cleanest falsification of an omitted-variable story: any confound that operates across all industries would also appear in the capital-light half. Instead, the pattern is exactly what the putty-clay model predicts: the irreversibility of installed capital, and hence the bite of stranding risk, matters only where capital is physically substantial and difficult to retrofit.

Second, the interaction strengthens in the *post-Paris* period. Moving from the full sample to post-Paris roughly doubles the interaction coefficient in capital-intensive industries (Panel A: from -0.0010 to -0.0020, $t = -3.34$). This is consistent with markets learning to price stranding risk after the Paris Agreement of December 2015, when cost-side climate policy became a more credible threat.

Third, the two emissions specifications tell the same qualitative story but differ in power. The log-level specification (Panel A) produces larger and more significant interaction coefficients than the gross-PP&E-ratio specification (Panel B), consistent with the discussion above: the log-level approach avoids ratio bias and gives the data maximum flexibility in determining the scaling of emissions with firm size.

Table 2: The Stranding Premium Concentrates in Capital-Intensive Industries

	All industries		Capital-intensive		Capital-light	
	Full sample	Post-Paris	Full sample	Post-Paris	Full sample	Post-Paris
Emissions	0.0022**	0.0031**	0.0026*	0.0040**	0.0012	0.0016
	(0.0010)	(0.0012)	(0.0013)	(0.0016)	(0.0010)	(0.0013)
Capital age	0.0009	0.0007	0.0017*	0.0017	0.0003	0.0005
	(0.0008)	(0.0009)	(0.0009)	(0.0011)	(0.0011)	(0.0013)
Emissions × Capital age	−0.0001	−0.0005	−0.0010*	−0.0019***	0.0003	0.0005
	(0.0004)	(0.0004)	(0.0005)	(0.0006)	(0.0006)	(0.0008)
Obs	321,678	231,553	166,654	114,900	155,024	116,653
Controls	Yes					
Year-month FE	Yes					
SIC-2 FE	Yes					

Notes: Each cell reports the coefficient from a pooled OLS regression of monthly excess returns on ln(Scope 1 GHG emissions), capital age, their interaction, and controls, with year-month and SIC-2 fixed effects. Controls: ln(market cap), book-to-market, momentum, ROE, leverage, investment/assets, sales growth, EPS growth, HHI, market beta, and return volatility. All emissions and age variables are within-month z-scored. *Capital age* is recursive capital age (winsorized at the 1st and 99th percentiles) (Lin, Palazzo, and Yang 2020; Geelen, Hajda, and Morellec 2024). Industries are split into "capital-intensive" and "capital-light" at the median of SIC-2 gross PP&E / total assets (Compustat, 1985–1995), chosen to balance the number of firm-month observations in each group. "Post-Paris" covers January 2016 to December 2024. Two-way clustered standard errors (firm × calendar year) in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Specification robustness. Table 3 unpacks the capital-intensive post-Paris cell, the specification that identifies the putty-clay channel most clearly, across six regression designs, for both the log-level and the gross-PP&E-ratio emissions measures (Panels A and B). The emissions coefficient is positive and highly significant across all columns: in Panel A, a one-standard-deviation increase in ln(Scope 1) is associated with 55–81 basis points higher monthly returns ($t > 3.4$), confirming the carbon premium documented by Bolton and Kacperczyk (2021).

Column (1) shows the emissions premium alone; column (2) adds capital age, which is insignificant. When the interaction enters in column (3), the picture changes. The interaction coefficient is -0.0020 ($t = -3.34$): among high-emission firms, those with older capital earn significantly lower returns. This is the paper's clearest cross-sectional finding. It lines up with the model's cost-side stranding prediction in the case where the growth-wedge term dominates: young dirty capital is more exposed, so investors require greater compensation for holding it. By contrast, the age main effect on its own is small and imprecisely estimated.

The interaction is robust. It survives dropping industry fixed effects (column 4, $t = -4.37$), dropping all firm-level controls (column 5, $t = -2.69$), and replacing separate fixed effects with SIC-2 × year-month fixed effects (column 6, $t = -2.87$). Panel B tells the same story

with the gross-PP&E-ratio emissions measure: the interaction is -0.0015 ($t = -2.65$) in the baseline and -0.0018 ($t = -3.50$) without industry fixed effects. Appendix Table 12 shows the interaction also survives the more demanding SIC-3 × year-month fixed effects.

Table 3: Older Dirty Capital Earns Lower Returns

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Emissions = ln(Scope 1 GHG)						
Emissions	0.0034**	0.0034**	0.0042**	0.0048**	0.0085***	0.0036**
	(0.0014)	(0.0014)	(0.0014)	(0.0016)	(0.0018)	(0.0014)
Capital age		0.0008	0.0019	0.0020	0.0020	0.0018
		(0.0010)	(0.0012)	(0.0011)	(0.0013)	(0.0011)
Emissions × Capital age			−0.0020***	−0.0020***	−0.0018***	−0.0018**
			(0.0006)	(0.0005)	(0.0006)	(0.0006)
Obs	114,010	114,010	114,010	114,010	141,453	114,010
Panel B: Emissions = ln(Scope 1 GHG/Gross PP&E)						
Emissions	0.0021**	0.0021**	0.0023**	0.0025**	0.0039***	0.0019**
	(0.0008)	(0.0008)	(0.0009)	(0.0009)	(0.0011)	(0.0009)
Capital age		0.0008	0.0011	0.0010	0.0036**	0.0011
		(0.0010)	(0.0011)	(0.0010)	(0.0016)	(0.0009)
Emissions × Capital age			−0.0014**	−0.0013**	−0.0009	−0.0012*
			(0.0006)	(0.0005)	(0.0006)	(0.0006)
Obs	114,010	114,010	114,010	114,010	140,112	114,010
B&K controls	Yes	Yes	Yes	Yes	No	Yes
SIC-2 FE	Yes	Yes	Yes	No	Yes	—
Year-month FE	Yes	Yes	Yes	Yes	Yes	—
SIC-2 × year-month FE	No	No	No	No	No	Yes

Notes: The dependent variable is the monthly excess stock return. Panel A measures emissions as ln(Scope 1 GHG); Panel B as ln(Scope 1 GHG/gross PP&E) (Compustat `ppegt`). All emissions and age variables are within-month z-scored. *Capital age* is recursive capital age (winsorized at the 1st and 99th percentiles) (Lin, Palazzo, and Yang 2020; Geelen, Hajda, and Morellec 2024). Columns (1)–(3) include year-month and SIC-2 fixed effects; column (4) drops industry fixed effects; column (5) drops all firm-level controls; column (6) replaces separate year-month and industry fixed effects with SIC-2 × year-month fixed effects. B&K controls: ln(market cap), book-to-market, momentum, ROE, leverage, investment/assets, ln(PPE), sales growth, EPS growth, HHI, market beta, and return volatility (Bolton and Kacperczyk, 2021). Annual accounting variables are mapped to monthly returns using the Fama–French June t+1 convention; monthly controls are lagged one month. Sample: capital-intensive industries (SIC-2 industries above the balanced cut-off of gross PP&E / total assets, measured 1985–1995), January 2016 to December 2024. Two-way clustered standard errors (firm × calendar year) in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Economic magnitude. Figure 3 makes the age gradient visible without imposing a functional form. Among capital-intensive firms in the post-Paris period, the carbon premium

declines monotonically from roughly 75 bps/month for the youngest capital-age quintile to near zero for the oldest. Capital-light firms cluster around zero at every age quintile, confirming that the premium operates through the putty-clay channel. The gradient is economically large: the youngest-to-oldest quintile spread for dirty capital-intensive firms is several times the average carbon premium reported by Bolton and Kacperczyk (2021), whose specification (which omits the interaction) captures the premium evaluated at the sample mean age.

We sort firms into within-month capital-age quintiles and estimate the emissions coefficient separately for each quintile via a fully saturated pooled regression with quintile × emissions × capital-intensity × post-Paris interactions. The right panel shows the pre-Paris period: capital-intensive coefficients are uniformly small and insignificant, consistent with the market not pricing stranding risk before the Paris Agreement. The aggregate four-way interaction $\phi \times A \times \text{CapInt} \times \text{Post}$ is −39 bps ($t = -2.88$; Appendix Table 16), and the result is not sensitive to the exact choice of cutoff date (Figure 4).

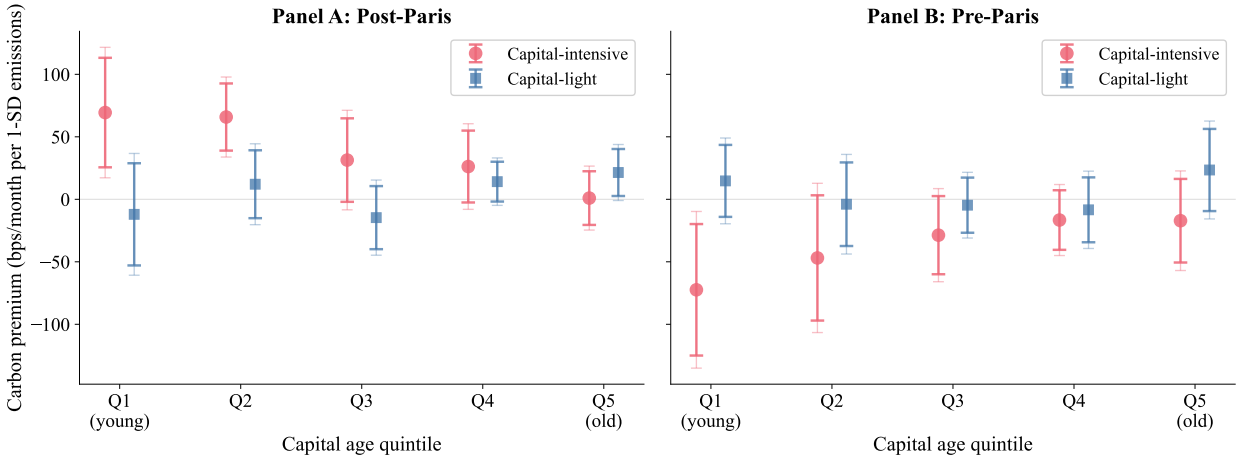


Figure 3: Carbon premium by capital-age quintile.

Notes: Estimated from a pooled regression with quintile × ln(Scope 1) × capital-intensity × post-Paris interactions. Each marker shows the emissions coefficient for one within-month age quintile (Q1 = youngest, Q5 = oldest). Thick error bars: 90% confidence intervals; thin error bars: 95% confidence intervals; both based on two-way (firm, calendar year) clustered standard errors. Red circles: capital-intensive industries; blue squares: capital-light industries (placebo). Left panel: post-Paris (2016–2024); right panel: pre-Paris (2005–2015). Controls and fixed effects as in Table 3. The sample includes all firm-months with non-missing Scope 1 emissions and recursive capital age. Age quintiles are formed within each calendar month. The regression includes the full Bolton and Kacperczyk (2021) control set, year-month fixed effects, and SIC-2 × year-month fixed effects, with all lower-order interaction terms fully saturated. $N = 277{,}684$. Appendix Table 17 reports the numerical coefficients; Table 16 reports the aggregate four-way interaction; and Figure 4 examines sensitivity to the post-Paris cutoff date (with COVID robustness in Appendix Figure OA.1).

The ratio $-\hat{\beta}_3/\hat{\beta}_1$ provides a useful summary of how fast the carbon premium decays with capital age, free of scaling assumptions. In the baseline specification, $\hat{\beta}_3/\hat{\beta}_1 \approx -0.30$: a one-standard-deviation increase in capital age erodes the emissions premium by roughly 30%.

Converting to calendar time using the cross-sectional standard deviation of recursive capital age ($\sigma_A \approx 2.3$ years), the premium decays at approximately 13% per year of capital age, implying a half-life of $\ln 2/0.13 \approx 5$ years. A capital-intensive firm with eight-year-old capital (90th percentile) bears roughly 40% of the stranding premium of an otherwise-identical firm with new capital; at fifteen years the share falls to about 12%.

This half-life has a natural structural interpretation. In the model, the numerator of the stranding ratio $\mathcal{S}^{PE}$ (the present value of fuel exposure) decays at rate g_y as fuel commitments run off with age. A 5-year half-life is consistent with effective depreciation rates in the range of 10–15% per year, broadly in line with BEA estimates for industrial equipment and power-generation assets. The interaction term thus reveals that the carbon premium is not a uniform tax on dirty firms but a vintage-dependent price of stranding risk, concentrated in firms whose capital commitments are recent.

Interpreting the sign of the interaction. The negative interaction coefficient $\beta_3 < 0$ has a natural model-based interpretation. The utility event study (Section 3) established that older fossil fleets lose less around cost-side events, a positive event-time dirty-age interaction. If the same class of cost-side stranding shocks drives both the event-study results and the unconditional cross-section, and if aggregation across vintages does not overturn the sign, then the positive event-time interaction together with $\beta_3 < 0$ implies that the stranding shocks priced in unconditional returns are on balance contractionary ($\Delta C < 0$): dirty assets lose value in high-marginal-utility states, generating a positive risk premium that decays with capital age as fuel exposure diminishes. We therefore read the unconditional evidence as consistent with cost-side policy risk (carbon taxes, cap-and-trade, emissions regulations) dominating the pricing of stranding risk, as opposed to technology-driven stranding which would imply expansionary shocks ($\Delta C > 0$) and $\beta_3 > 0$. This is consistent with the findings of

Machine-level amplification and firm-level regressions. Equation (12) delivers an exact proportionality only at the machine level: the cross-partial is proportional to the age gradient with factor $\mathcal{M}$. After aggregation across vintages and projection onto a firm-level linear regression, this need not survive as an exact restriction such as $\beta_3 = \bar{\mathcal{M}}\beta_2$. The useful implication is qualitative. Capital-intensive industries exhibit substantial cross-sectional variation in emissions per unit of capital (for example, fossil generators, gas turbines, and renewable operators within broad industry groupings), so even a modest reduced-form age effect can coexist with a sizable negative interaction. That is the pattern in the data: $\hat{\beta}_2$ is small and imprecisely estimated, while $\hat{\beta}_3$ is significantly negative.

When did the market start pricing stranding risk? In Figure 4 we run a single fully saturated pooled OLS on the entire panel (all firms, all months, all industries) with four interacted variables: emissions (ϕ), capital age ($\bar{a}$), capital intensity (CapInt), and a post-period dummy (Post). The regression includes all 15 lower-order interactions, the full Bolton and Kacperczyk (2021) controls, and year-month plus SIC-2 fixed effects. We then vary the cutoff date that defines the Post dummy from 2010-Q1 to 2023-Q4, producing a coefficient path for each cutoff. Panel (a) plots the three-way interaction $\phi \times \text{CapInt} \times \text{Post}$:

the incremental carbon premium in capital-intensive industries in the post period. Panel (b) plots the four-way interaction $\phi \times \bar{a} \times \text{CapInt} \times \text{Post}$: the incremental age gradient of stranding risk.

Each coefficient measures the average premium in all years *after* the cutoff minus the average in all years *before* it, so the coefficient path identifies the cutoff that best separates the low-premium and high-premium eras. Both panels peak when the cutoff aligns with the Paris Agreement (2015–2016), identifying this as the break point at which cost-side policy became a credible threat. A second local peak around 2020–2021, separated from the first by a trough, suggests a further repricing coinciding with the Biden administration's renewed regulatory agenda (EPA Power Plant Rule, IRA). The pattern is unchanged when we add COVID-period controls, a dummy for March 2020 through December 2021 fully interacted with emissions, age, and capital intensity, confirming that the 2020–2021 strengthening reflects policy repricing rather than pandemic-related return mechanics (Appendix Figure OA.1).

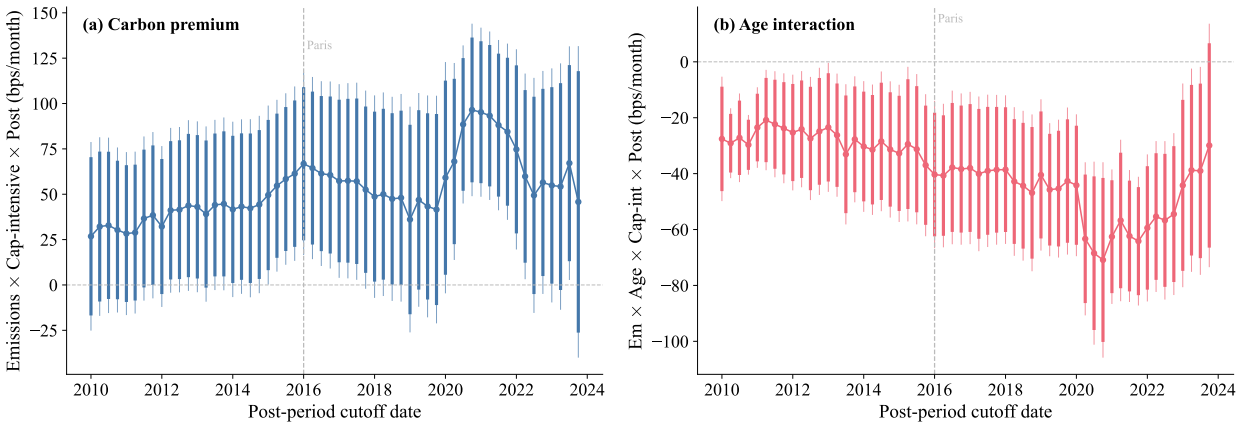


Figure 4: Timing of stranding-risk pricing.

Notes: Each point is a coefficient from a single fully saturated pooled OLS on the full panel (all industries, all months), varying only the cutoff date that defines the Post dummy. Panel (a): the three-way interaction $\phi \times \text{CapInt} \times \text{Post}$ (incremental carbon premium in capital-intensive industries). Panel (b): the four-way interaction $\phi \times \bar{a} \times \text{CapInt} \times \text{Post}$ (incremental age gradient). Thick error bars: 90% CI; thin error bars: 95% CI; both based on two-way (firm, calendar year) clustered standard errors. Dashed vertical line marks the Paris Agreement (December 2015). Controls and fixed effects as in Table 3.

4.3 Event study

The cross-sectional regressions establish that the emissions–age interaction is priced in unconditional monthly returns. A natural concern is that this correlation reflects omitted firm characteristics rather than stranding risk per se. To provide closer-to-causal evidence that discrete increases in stranding risk drive the premium, we test whether the same interaction appears in narrow windows around cost-side climate-policy events, where confounding variation is limited.

We expand the candidate set beyond the utility-specific EPA events of Section 3 to all

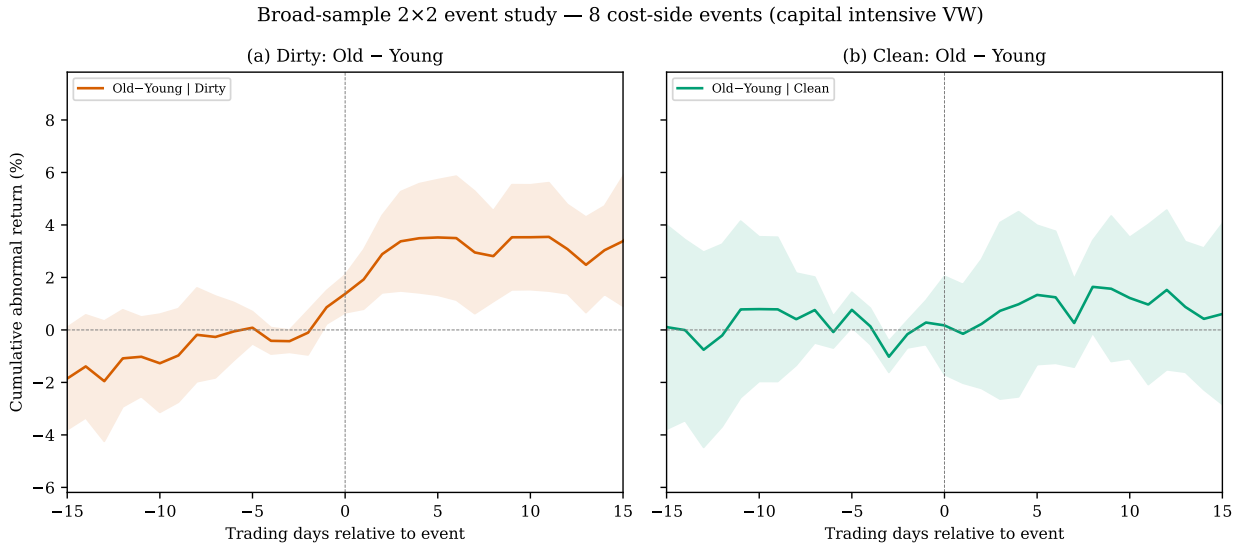


Figure 5: Old−Young CAAR spread in capital-intensive industries around eight cost-side climate-policy events.

Notes: Panel (a): dirty firms (top emissions quartile); Panel (b): clean firms (bottom emissions quartile). Value-weighted portfolios formed on quartiles of Scope 1 emissions and recursive capital age. Shaded bands show 95% confidence intervals based on cross-event covariance matrix standard errors.

climate-policy shocks that materially shift expectations about carbon pricing across sectors, including EPA regulations, Supreme Court decisions, elections, and international agreements. From this universe, we apply the same contamination-exclusion logic as in the utility sample: events are excluded when contaminated by contemporaneous shocks to energy markets, financial conditions, or overlapping climate summits. We additionally require $F_i < 0$ (not anticipated; Section F.3). Eight events survive these filters (Appendix Table 7).

Portfolio event study. For each event, we sort capital-intensive firms into quartiles on Scope 1 emissions and (conditionally) on recursive capital age, forming four corner portfolios: dirty-old, dirty-young, clean-old, clean-young. Value-weighted portfolio returns are computed daily using prior-June market capitalization. The signed Old−Young spread, $d_e \times (R_{e,t}^{\text{Old}} - R_{e,t}^{\text{Young}})$, is cumulated across events with cross-event covariance matrix standard errors.

Figure 5 presents the results for capital-intensive industries. Panel (a) shows the dirty Old−Young CAAR spread: old dirty firms significantly outperform young dirty firms when cost-side regulation tightens, consistent with the utility-sector finding and the model's prediction that young dirty capital bears the greatest stranding risk. Panel (b) shows the clean spread, which is insignificant; the age gradient operates only through the dirty-capital channel, as the model predicts.

Pooled daily OLS. Table 4 confirms the portfolio result parametrically using the **?** pooled daily OLS specification:

$$r_{i,t} = \alpha_{ym} + \mathbf{b}'\mathbf{Z}_i + \mathbf{d}'(\mathbf{Z}_i \times \text{EventDay}_t) + \varepsilon_{i,t},$$

where $\mathbf{Z}_i = (z(\text{Em}_i),\ z(\text{RecAge}_i),\ z(\text{Em}_i \times \text{RecAge}_i))'$ are z-scored cross-sectional exposure variables and EventDay_t equals $+1$ on tightening days, -1 on loosening days, and zero otherwise. Standard errors are two-way clustered by firm and calendar date. The triple interaction Em × RecAge × EventDay is positive and significant: +23.9 bp ($t = 2.70$) in the raw specification and +11.5 bp ($t = 2.11$) after Carhart adjustment. The capital-light subsample is insignificant across all risk adjustments ($t = 1.19$–1.63; Appendix Table 19), confirming the placebo.

Table 4: Pooled Daily OLS Event Study — Capital-Intensive Industries

	(1) Raw	(2) Mkt-adj	(3) FF3	(4) FF4	(5) FF5
EventDay	−32.5069	21.5004	53.6034	47.6429	–
	(51.0132)	(32.3244)	(35.4099)	(31.0479)	
Emissions (z)	0.1319	0.1335	−0.6755	−0.6767	−0.6898
	(0.7903)	(0.7937)	(0.5374)	(0.5357)	(0.4678)
Rec. capital age (z)	0.3588	0.3613	0.6514**	0.6618**	0.6081**
	(0.4538)	(0.4590)	(0.2800)	(0.2763)	(0.2553)
Em × RecAge	−0.0214	−0.0224	−0.0186	−0.0197	0.0138
	(0.3517)	(0.3552)	(0.2494)	(0.2485)	(0.2177)
Em × EventDay	−35.1008***	−35.7878***	−16.2419***	−16.1268***	−18.4343***
	(9.0268)	(9.3946)	(5.1726)	(5.2186)	(5.7384)
RecAge × EventDay	−12.6625	−14.8569*	−7.3020	−6.4491	−7.4831
	(7.9120)	(8.5214)	(10.7267)	(10.5356)	(9.5393)
Em × RecAge × EventDay	21.1137***	21.2849***	11.4996***	11.9722***	10.8819***
	(6.9842)	(7.2162)	(4.2299)	(4.6321)	(4.0153)
Fixed effects	YM	YM	YM	YM	Date
Clustering	Firm+Date	Firm+Date	Firm+Date	Firm+Date	Firm+Date
Observations	4,428,907	4,428,907	4,428,459	4,428,459	4,428,459
Firms	2,236	2,236	2,223	2,223	2,223
Event firm-days	7,521	7,521	7,521	7,521	7,521

5 Conclusion

This paper provides a structural foundation for the carbon risk premium. The premium compensates investors for stranding risk: the exposure of locked-in fuel commitments to cost-side climate policy. Our putty-clay model predicts that this exposure is concentrated

in young, dirty, capital-intensive firms, and we find strong empirical support in both the utility sector (where generator-level data allow direct measurement) and the broad stock market. The carbon premium declines monotonically with capital age, with an implied half-life of approximately five years, and is indistinguishable from zero for old capital and for capital-light industries.

A distinctive contribution is the taxonomy of stranding mechanisms. Cost-side instruments (carbon taxes, cap-and-trade, emissions regulations) are unique in generating a strong emissions $\times$ age interaction, because they directly raise operating costs of installed capital in proportion to fuel intensity. Green technology improvements and investment subsidies strand dirty capital indirectly through output prices, producing weaker age dependence. Dirty investment taxes and ESG-driven cost-of-capital wedges can benefit existing dirty incumbents by reducing competitive entry. The cross-sectional pattern of returns around a policy event is therefore diagnostic: a strong age $\times$ dirty interaction signals cost-side risk; its absence points to technology or investment-channel mechanisms.

These findings have important implications. For investors, the results suggest that carbon risk is not a uniform characteristic of dirty firms but varies predictably with the vintage of the capital stock; portfolio construction that conditions on capital age can sharpen the measurement of climate risk exposure. For policymakers, the concentration of stranding losses in young dirty capital means that financial markets penalize the flow of new dirty investment, complementing explicit carbon pricing in steering capital allocation. However, by the same logic, the premium does not create strong financial incentives to retire existing old dirty capital, since old capital carries negligible stranding risk.

Several limitations point to directions for future work. Our model takes the vintage distribution as given; endogenizing investment and entry decisions would allow the stranding premium itself to feed back into the composition of the capital stock. The model abstracts from firm-level dynamics such as debt overhang, refinancing risk, and managerial responses to stranding exposure, all of which could amplify or dampen the premium in practice. Finally, our empirical evidence is drawn entirely from US equity markets; extending the analysis to international settings where climate policy regimes differ would test whether the mechanism operates through the same channels globally.

References

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

Acharya, V. V., S. Giglio, E. Pastore, J. Stroebel, Z. Tan, and Z. Yong (2025): "Climate Transition Risks and the Energy Sector," Working Paper 33413, National Bureau of Economic Research.

Aghion, P., A. Dechezleprêtre, D. Hémous, R. Martin, and J. Van Reenen (2016): "Carbon taxes, path dependency, and directed technical change: evidence from the auto industry," *Journal of Political Economy*, 124, 1–51.

Aswani, J., A. Raghunandan, and S. Rajgopal (2024): "Are Carbon Emissions Associated with Stock Returns?" *Review of Finance*, 28, 75–106.

Bahk, B.-H. and M. Gort (1993): "Decomposing learning by doing in new plants," *Journal of Political Economy*, 101, 561–583.

Baker, S. R., N. Bloom, and S. J. Davis (2016): "Measuring Economic Policy Uncertainty," *The Quarterly Journal of Economics*, 131, 1593–1636.

Barclay, M. J. and C. W. Smith (1995): "The maturity structure of corporate debt," *The Journal of Finance*, 50, 609–631.

Barnett, M., W. Brock, and L. P. Hansen (2020): "Pricing uncertainty induced by climate change," *The Review of Financial Studies*, 33, 1024–1066.

Bertolotti, F., A. Lanteri, and H. Yoon (2025): "Optimal Subsidies for Capital Replacement and the Green Transition," Working paper.

Billett, M. T., T.-H. D. King, and D. C. Mauer (2007): "Growth opportunities and the choice of leverage, debt maturity, and covenants," *The Journal of Finance*, 62, 697–730.

Bolton, P. and M. Kacperczyk (2021): "Do investors care about carbon risk?" *Journal of Financial Economics*, 142, 517–549.

——— (2023): "Global pricing of carbon-transition risk," *The Journal of Finance*, 78, 3677–3754.

Campello, M., G. Kankanhalli, and H. Kim (2024): "Delayed creative destruction: How uncertainty shapes corporate assets," *Journal of Financial Economics*, 153, 103786.

Chittaro, L., M. Piazzesi, M. Sena, and M. Schneider (2025): "Asset Returns as Carbon Taxes," Working paper, Stanford University.

Crosignani, M., E. Osambela, and M. Pritsker (2025): "Understanding the Pricing of Carbon Emissions: New Evidence from the Stock Market," Working paper, Federal Reserve.

Curtis, A., M. F. Lewis-Western, and S. Toynbee (2015): "Historical cost measurement and the use of DuPont analysis by market participants," *Review of Accounting Studies*, 20, 1210–1245.

Custódio, C., M. A. Ferreira, and L. Laureano (2013): "Why are US firms using more short-term debt?" *Journal of Financial Economics*, 108, 182–212.

Darmouni, O. and A. Sutherland (2024): "Investment when new capital is hard to find," *Journal of Financial Economics*, 154, 103806.

Eisfeldt, A. L. and A. A. Rampini (2006): "Capital reallocation and liquidity," *Journal of Monetary Economics*, 53, 369–399.

——— (2007): "New or used? Investment with credit constraints," *Journal of Monetary Economics*, 54, 2656–2681.

Engle, R. F., S. Giglio, B. Kelly, H. Lee, and J. Stroebel (2020): "Hedging climate change news," *The Review of Financial Studies*, 33, 1184–1216.

Fama, E. F. and J. D. MacBeth (1973): "Risk, return, and equilibrium: empirical tests," *Journal of Political Economy*, 81, 607–636.

Gârleanu, N. and L. H. Pedersen (2025): "Climate Risk Pricing," Working paper.

Geelen, T., J. Hajda, E. Morellec, and A. Winegar (2024): "Asset life, leverage, and debt maturity matching," *Journal of Financial Economics*, 154, 103796.

Gilchrist, S., J. Martinez, and N. Rickard (2026): "The Green Energy Transition in a Putty-Clay Model of Capital," Working paper.

Gilchrist, S. and J. C. Williams (2000): "Putty-clay and investment: a business cycle analysis," *Journal of Political Economy*, 108, 928–960.

Gómez-Cram, R. and M. Olbert (2023): "Measuring the Expected Effects of the Global Tax Reform," Working paper.

Gourio, F. (2011): "Putty-clay technology and stock market volatility," *Journal of Monetary Economics*, 58, 117–131.

Greenwood, J., Z. Hercowitz, and P. Krusell (1997): "Long-run implications of investment-specific technological change," *American Economic Review*, 87, 342–362.

Hsu, P.-H., K. Li, and C.-Y. Tsou (2023): "The pollution premium," *The Journal of Finance*, 78, 1343–1392.

Hulten, C. R. and F. C. Wykoff (1981): "The measurement of economic depreciation," in *Depreciation, Inflation, and the Taxation of Income from Capital*, ed. by C. R. Hulten, Washington, D.C.: The Urban Institute Press, 81–125.

İmrohoroğlu, A. and c. Tüzel (2014): "Firm-level productivity, risk, and return," *Management Science*, 60, 2073–2090.

Johansen, L. (1959): "Substitution versus Fixed Production Coefficients in the Theory of Economic Growth: A Synthesis," *Econometrica: Journal of the Econometric Society*, 27, 157–176.

Lanteri, A. and A. A. Rampini (2022): "Financing the adoption of clean technology," *Working paper*.

Li, K., C.-Y. Tsou, and C. Xu (2023): "Learning and the capital age premium," *Journal of Monetary Economics*, 136, 76–90.

Lin, X., B. Palazzo, and F. Yang (2020): "The risks of old capital age: Asset pricing implications of technology adoption," *Journal of Monetary Economics*, 115, 145–161.

Ma, S., J. Murfin, and R. Pratt (2022): "Young firms, old capital," *Journal of Financial Economics*, 146, 331–356.

Pástor, v., R. F. Stambaugh, and L. A. Taylor (2021): "Sustainable investing in equilibrium," *Journal of Financial Economics*, 142, 550–571.

Pedersen, L. H., S. Fitzgibbons, and L. Pomorski (2021): "Responsible investing: the ESG-efficient frontier," *Journal of Financial Economics*, 142, 572–597.

Petersen, M. A. (2009): "Estimating standard errors in finance panel data sets: comparing approaches," *The Review of Financial Studies*, 22, 435–480.

Stohs, M. H. and D. C. Mauer (1996): "The determinants of corporate debt maturity structure," *Journal of Business*, 69, 279–312.

Zhang, L. (2005): "The value premium," *The Journal of Finance*, 60, 67–103.

Zhang, S. (2025): "Carbon Returns across the Globe," *The Journal of Finance*, 80, 615–645.

A Proofs

Lemma 1 (Revenue independence under free entry)**.** *On the balanced growth path with competitive machine investment, frontier gross revenue per machine $P(\phi)x^*(\phi)$ is independent of ϕ: $(Px^*)'(\phi) = 0$.*

Proof. Write the value of a frontier machine ($a = 0$) using equation (9):

$$v^0(0,\phi) = Px^*\mathcal{A}_{r-g_z}(0) - p^f f^*\mathcal{A}_r(0) - w\mathcal{A}_{r-g_y}(0).$$

Substitute the fuel FOC (5) into the fuel-cost term: $p^f f^*\mathcal{A}_r(0) = \alpha_f Px^*\mathcal{A}_{r-g_z}(0)$. So $v^0(0,\phi) = (1-\alpha_f)Px^*\mathcal{A}_{r-g_z}(0) - w\mathcal{A}_{r-g_y}(0)$. Set this equal to the capital cost via free entry ($v^0(0,\phi) = q_k k^*$) and substitute the capital FOC (6):

$$(1-\alpha_f)Px^*\mathcal{A}_{r-g_z}(0) - w\mathcal{A}_{r-g_y}(0) = \alpha_k Px^*\mathcal{A}_{r-g_z}(0).$$

Since $1-\alpha_f-\alpha_k = \alpha_\ell$: $\alpha_\ell Px^*\mathcal{A}_{r-g_z}(0) = w\mathcal{A}_{r-g_y}(0)$. Solving gives $Px^* = w\mathcal{A}_{r-g_y}(0)/[\alpha_\ell\mathcal{A}_{r-g_z}(0)]$, which depends only on common parameters. □

A.1 Proof of Proposition 1

Part (i): Dirtiness. For $(a,\phi) \in \mathcal{D}_+$,

$$\log \mathcal{S}^{PE} = \log f^* - g_y a + \log \mathcal{A}_r - \log v^0,$$

so

$$\partial_\phi \log \mathcal{S}^{PE} = \frac{(f^*)'}{f^*} - \frac{\partial_\phi v^0}{v^0}.$$

From the fuel FOC, $f^*(\phi) = \alpha_f(\phi)Px^*\Lambda/p^f$, so $(f^*)'(\phi) = \alpha_f'(\phi)Px^*\Lambda/p^f > 0$ by assumption. For the second term, since $(Px^*)' = 0$ (Lemma 1):

$$\partial_\phi v^0 = -p^f(f^*)'e^{-g_y a}\mathcal{A}_r < 0.$$

Hence $-\partial_\phi v^0/v^0 > 0$. Both terms are positive, so $\partial_\phi \log \mathcal{S}^{PE} > 0$ on $\mathcal{D}_+$.

Part (ii): Sign equivalence. Using the tail-derivative identity $\partial_a F(a) = \delta F(a) - h(a)$ for $F(a) = \int_0^{A-a} e^{-\delta t}h(a+t)\,dt$, applied to $\mathcal{A}_r(a)$ and $v^0(a,\phi)$:

$$\partial_a \log \mathcal{A}_r = r - \frac{1}{\mathcal{A}_r}, \qquad \partial_a \log v^0 = (r-g_y) - \frac{\pi}{v^0}.$$

From $\log \mathcal{S}^{PE} = \log f^* - g_y a + \log \mathcal{A}_r(a) - \log v^0$:

$$\partial_a \log \mathcal{S}^{PE} = -g_y + r - \frac{1}{\mathcal{A}_r} - (r-g_y) + \frac{\pi}{v^0} = \frac{\pi}{v^0} - \frac{1}{\mathcal{A}_r} = \frac{\pi\mathcal{A}_r - v^0}{v^0\mathcal{A}_r} = \frac{-B(a)}{v^0\mathcal{A}_r}.$$

For the cross-partial, differentiating π/v^0 with respect to ϕ using $(Px^*)' = 0$:

$$\partial_\phi\left[\frac{\pi}{v^0}\right] = \frac{(\partial_\phi\pi)v^0 - \pi(\partial_\phi v^0)}{(v^0)^2} = \frac{-p^f(f^*)'e^{-g_y a}[v^0 - \pi\mathcal{A}_r]}{(v^0)^2} = \frac{-p^f(f^*)'e^{-g_y a}\cdot B(a)}{(v^0)^2}.$$

Since the prefactor $-p^f(f^*)'e^{-g_y a} < 0$, $\operatorname{sgn}(\partial^2_{a\phi}\log\mathcal{S}^{PE}) = -\operatorname{sgn}(B)$. The proportionality $\partial^2_{a\phi} = \mathcal{M}\cdot\partial_a\log\mathcal{S}^{PE}$ follows by direct comparison.

Part (iii): Single crossing. Define $B(a) = v^0(a) - \pi(a)\mathcal{A}_r(a)$. The ODE for B is:

$$B'(a) = (r - g_y)B(a) + \mathcal{A}_r(a)\,\Psi_B(a), \qquad \Psi_B(a) \equiv g_y w - g_z P x^* e^{-\tilde{g}a}. \tag{22}$$

The forcing term $\Psi_B(a)$ is strictly increasing in a (since $e^{-\tilde{g}a}$ decreases). Near end of life, $B(a) \to \frac{(A-a)^2}{2}[g_z P x^* e^{-\tilde{g}A} - g_y w] < 0$ (the bracket is negative because $g_z/g_y < \alpha_\ell \leq w/(Px^*e^{-\tilde{g}A})$ under the model's parameter restrictions). So $B(a) < 0$ near $a = A$.

Existence of a^*: $B(0) > 0$ (by hypothesis (13)) and $B < 0$ near A, so by continuity at least one zero exists.

The first zero a_1 satisfies $a_1 \leq \hat{a}$ (where $\Psi_B(\hat{a}) = 0$): since $B(0) > 0$ and $B(a_1) = 0$, we need $B'(a_1) \leq 0$, hence $\Psi_B(a_1) \leq 0$, hence $a_1 \leq \hat{a}$.

Uniqueness: Suppose a second zero $a_2 > a_1$ exists. Then $B(a_2) = 0$ with B arriving from below, so $B'(a_2) \geq 0$, hence $\Psi_B(a_2) \geq 0$, hence $a_2 \geq \hat{a}$. After a_2, B becomes positive. For all $a > \hat{a}$ where $B(a) > 0$: both terms in $B'(a) = (r - g_y)B(a) + \mathcal{A}_r(a)\Psi_B(a)$ are positive (using the maintained assumption $r > g_y$), so B is strictly increasing. But $B < 0$ near A: a positive, increasing function cannot reach the negative region. Contradiction. □

Part (iv): Clean-side placebo. For $\phi = 0$: $f^*(0) = 0$, so $\mathcal{S}^{PE}(a, 0) = 0$ at every age. This is a PE statement; GE pass-through is handled separately in Proposition 4. □

A.2 Proof of Proposition 2

Let $V(a, \phi; \tau)$ denote the value after a permanent tax of size τ, and define

$$S(a, \phi) \equiv -\frac{1}{V(a, \phi; 0)}\frac{\partial V(a, \phi; \tau)}{\partial \tau}\bigg|_{\tau=0}.$$

Then

$$\log V(a, \phi; \tau) = \log V(a, \phi; 0) - \tau\, S(a, \phi) + o(\tau).$$

In PE, $S = \mathcal{S}^{PE}$ by definition. With price and wage responses, differentiating the GE value formula yields $S = S^{GE}$ as in Proposition 4. Under Poisson arrival at rate λ with SDF jump $1 + \xi$, standard jump-risk pricing gives the expected excess return as arrival rate times jump price times exposure:

$$E[r^e(a, \phi)] = \lambda\,\xi\,\tau\, S(a, \phi) + o(\tau).$$

The log-derivative transfer follows because $\log E[r^e] = \log(\lambda\xi\tau) + \log S + o(\tau)$, and $\log(\lambda\xi\tau)$ is independent of (a, ϕ). □

B Variable Utilization

With idiosyncratic operate/shut-down shocks drawn iid Gumbel$(0, \kappa)$, utilization is $u(a, \phi) = \sigma(\pi(a, \phi)/\kappa) \in (0, 1)$, and the expected operating payoff is $\kappa \log(1 + e^{\pi/\kappa})$. The key new force is *utilization self-insurance*: old machines near the shutdown boundary reduce operations, limiting fuel exposure. When $\partial_a \pi < 0$, this acts in the same direction as the growth wedge: young machines are far from shutdown ($u \approx 1$) and bear the full tax burden.

Proposition 3 (Uniform small-κ approximation)**.** *Fix a compact set $\mathcal{K} \subset [0, A) \times [0, 1]$ on which $\underline{\pi} \equiv \inf_{(\tilde{a},\phi)\in\mathcal{K}^+} \pi(\tilde{a}, \phi) > 0$ (where $\mathcal{K}^+$ is the continuation envelope). Then $\mathcal{S}^{PE}_\kappa(a, \phi) = \mathcal{S}^{PE}(a, \phi) + O(e^{-\underline{\pi}/\kappa})$ uniformly on $\mathcal{K}$, and the same exponential closeness holds for first derivatives.*

Corollary 1. *If a deterministic derivative of $\mathcal{S}^{PE}$ is bounded away from zero on $\mathcal{K}$, then the corresponding derivative of $\mathcal{S}^{PE}_\kappa$ has the same sign for all sufficiently small κ.*

The proofs follow from the exponential convergence $|1-u| \le e^{-\underline{\pi}/\kappa}$ and $|\kappa \log(1+e^{\pi/\kappa}) - \pi| \le \kappa e^{-\underline{\pi}/\kappa}$ on the continuation-positive set; the integrals defining J_κ, v^0_κ, and their derivatives inherit the same bounds via dominated convergence.

Effect on the three derivatives. The dirtiness derivative remains positive on continuation-positive compact sets for small κ. The age derivative is pushed toward negative (young more exposed) when $\partial_a \pi < 0$: young machines have $u \approx 1$ and bear the full fuel cost, while old machines self-insure. The cross-partial is pushed toward negative under the same monotonicity conditions, reinforcing the growth wedge.

C General Equilibrium

Let $\hat{P} \equiv (dP/d\tau)/P$ and $\hat{w} \equiv (dw/d\tau)/w$ denote the semi-elasticities of output price and wage to a marginal tax increase, and define the revenue-to-value and wage-to-value ratios $\mathcal{R}(a, \phi) \equiv Px^* e^{-\tilde{g}a} \mathcal{A}_{r-g_z}(a)/v^0(a, \phi)$ and $\mathcal{W}(a, \phi) \equiv w\, \mathcal{A}_{r-g_y}(a)/v^0(a, \phi)$.

Proposition 4 (GE adjustment to stranding)**.** *Under free entry:* (i) $\mathcal{R}(0, \phi) = 1/\alpha_k(\phi)$ *and* $\mathcal{W}(0, \phi) = \alpha_\ell/\alpha_k(\phi)$*, independent of r, g_z, g_θ, and A.* (ii) *The GE stranding ratio is* $S^{GE}(a, \phi) = \mathcal{S}^{PE}(a, \phi) - \hat{P}\,\mathcal{R}(a, \phi) + \hat{w}\,\mathcal{W}(a, \phi)$. (iii) *If profit margins thin with age ($\mathcal{R}$ and $\mathcal{W}$ increasing), and the tax raises prices ($\hat{P} > 0$) and lowers wages ($\hat{w} < 0$), then both GE channels reduce the age gradient of stranding relative to PE, making young-dirty dominance more likely.*

Proof. Part (i): from the capital FOC and free entry, $v^0(0, \phi) = \alpha_k(\phi) Px^* \mathcal{A}_{r-g_z}(0)$; dividing frontier revenue and the present-value wage bill by this expression gives the result. Parts (ii)–(iii): differentiate v^0 with respect to τ including the P and w responses; the age gradient of each GE term inherits the sign of $\partial_a \mathcal{R}$ or $\partial_a \mathcal{W}$.

Detailed channel decomposition. In general equilibrium, three channels supplement the PE fuel-cost channel: (i) wages respond ($dw = w_\tau\, d\tau$), (ii) revenue responds ($d\log(PY) = \Psi(\phi)\, d\tau$, age-independent), and (iii) discount rates may change (ignored throughout). The GE stranding decomposes as

$$S^{GE} = \mathcal{S}^{PE} + \operatorname{sgn}(w_\tau)\, \mathcal{W} - \Psi\, \mathcal{R}, \tag{23}$$

where $\mathcal{W}(a,\phi) = |w_\tau| \mathcal{A}_{r-g_y}(a)/v^0$ (wage channel) and $\mathcal{R}(a,\phi) = Px^* e^{-\tilde{g}a} \mathcal{A}_{r-g_z}(a)/v^0$ (revenue channel).

Define three "payout rates":

$$G_r \equiv \frac{\pi}{v^0} - \frac{1}{\mathcal{A}_r}, \qquad G_y \equiv \frac{\pi}{v^0} - \frac{1}{\mathcal{A}_{r-g_y}}, \qquad G_z \equiv \frac{\pi}{v^0} - \frac{1}{\mathcal{A}_{r-g_z}}.$$

The PE age gradient is $G_r = \partial_a \log \mathcal{S}^{PE}$; the wage and revenue channels have age gradients $G_y \mathcal{W}$ and $G_z \mathcal{R}$ respectively. For clean firms, $\mathcal{S}^{PE}(a,0) = 0$, so any GE response comes entirely through these pass-through terms.

Proposition 5 (GE channel formulas)**.** *The exact GE corrections to each derivative are:*

$$\partial_\phi S^{GE} = \partial_\phi \mathcal{S}^{PE} - \mathcal{M}\mathcal{W} - \mathcal{R}(\Psi' + \Psi\mathcal{M}), \tag{24}$$

$$\partial_a S^{GE} = \partial_a \mathcal{S}^{PE} - G_y \mathcal{W} - \Psi G_z \mathcal{R}, \tag{25}$$

$$\partial^2_{a\phi} S^{GE} = \partial^2_{a\phi} \mathcal{S}^{PE} - \mathcal{M}\mathcal{W}(G_r + G_y) - \mathcal{R}[\Psi' G_z + \Psi\mathcal{M}(G_r + G_z)]. \tag{26}$$

When wages fall after the tax ($w_\tau < 0$): the wage channel attenuates the PE dirtiness gradient ($-\mathcal{M}\mathcal{W} < 0$); pushes the age gradient toward negative when $G_y > 0$; and pushes the cross-partial toward negative when $G_r + G_y > 0$. The wage channel thus reinforces the growth wedge against leverage on the age and cross-partial margins. The revenue channel can either reinforce or offset these forces depending on Ψ' and Ψ.

D Alternative Policy Instruments

Different policy instruments operate through different cash-flow channels, each with a different time profile and annuity rate:

Channel	Per-period change	Annuity rate	Duration
Fuel cost (τ)	$\tau f^* e^{-g_y(a+t)}$	r	shortest
Revenue (ΔP)	$\Delta P \cdot x^* e^{-\tilde{g}(a+t)}$	$r - g_z$	medium
Wage (Δw)	Δw (constant detrended)	$r - g_y$	longest

The ordering $\mathcal{A}_r(a) < \mathcal{A}_{r-g_z}(a) < \mathcal{A}_{r-g_y}(a)$ means the fuel channel has the shortest-duration annuity and the wage channel the longest. A policy that loads on the fuel channel creates the largest *duration mismatch* with the machine's value, generating the strongest age dependence.

Green technology and subsidies. Clean technology improvements raise clean-firm productivity, shifting CES demand away from dirty goods ($\Delta P(\phi) < 0$ for dirty firms) and raising wages ($\Delta w > 0$). Crucially, no fuel-cost channel is present. The profit change $\Delta\pi^{GT} = \Delta P \cdot x^* e^{-\tilde{g}(a+t)} - \Delta w$ loads entirely on the revenue ($\mathcal{A}_{r-g_z}$) and wage ($\mathcal{A}_{r-g_y}$) annuities, both of which are long-duration and closely matched to v^0's own duration structure. The resulting duration mismatch is small, so stranding is nearly age-independent: the age gradient and cross-partial are weak. Green investment subsidies (IRA-type) produce qualitatively identical effects through quantity expansion rather than quality improvement.

Dirty investment taxes. A tax on new dirty investment reduces entry, raising the dirty output price $\Delta P^{IT}(\phi) > 0$. Existing dirty machines *gain* value. The age gradient is reversed: young machines gain more (they have higher output and more remaining life). Incumbent protection: existing dirty capital benefits from reduced competition.

ESG cost-of-capital wedge. A dirtiness-dependent required return $\xi(\phi)$ has two effects: (i) all future cash flows are discounted more heavily (value falls), and (ii) reduced dirty entry raises dirty output prices (value rises). The net effect is ambiguous. For small ξ, the profit effect dominates and incumbents gain; for large ξ, the discount effect dominates. In neither case does the fuel channel play a role, so the cross-partial is weak.

Taxonomy. Table 5 summarizes these results.

Table 5: Stranding mechanisms across policy instruments

	Carbon tax	Green tech	Green subsidy	Dirty inv. tax	ESG wedge
Active channels					
Fuel cost ($\mathcal{A}_r$)	✓	–	–	–	–
Revenue ($\mathcal{A}_{r-g_z}$)	–	✓	✓	✓	✓
Wage ($\mathcal{A}_{r-g_y}$)	–	✓	✓	–	✓
Discount rate	–	–	–	–	✓
Signs for dirty firms					
Profit change	$-$	$-$	$-$	$+$	?
Dirtiness gradient	$+$ (**strong**)	$+$ (weak)	$+$ (weak)	$-$	?
Age gradient	$-\text{sgn}(B)$	≈ 0	≈ 0	$-$	?
Cross-partial	$-\text{sgn}(B)$	≈ 0	≈ 0	$-$	?
Duration mismatch with v^0	large	small	small	small	moderate
Risk premium sign (ΔC direction)	$+$ ($\Delta C < 0$)	$-$ ($\Delta C > 0$)	$-$ ($\Delta C > 0$)	$-$ ($\Delta C < 0$)	?

The presence of a strong age × dirtiness interaction in event-study returns is diagnostic of cost-side stranding risk. Its absence suggests technology-driven or subsidy-driven mecha-

nisms; a reversed interaction points to investment-side instruments or incumbent-protection effects.

E Age-Invariant-Profit Benchmark

Suppose per-period operating profit is age-invariant: $\pi(a, \phi) = \pi(\phi)$ for all a. Then machine value at age a is $V_a(\phi) = D(a) \cdot \pi(\phi)$, where $D(a)$ is a pure duration factor. The change upon a carbon tax is $\Delta V_a(\phi) = D(a) \times \Delta\pi(\phi)$, multiplicatively separable in age and dirtiness. In percentage terms:

$$\frac{\Delta V_a(\phi)}{V_a(\phi)} = \frac{\Delta\pi(\phi)}{\pi(\phi)},$$

which is independent of a. The age coefficient β_2 and the interaction β_3 are both zero in the cross-sectional regression. There is no stranding in the economically relevant percentage dimension.

This benchmark is more restrictive than simply setting $g_\theta = 0$. In the paper's baseline equations, $g_\theta = 0$ with $g_z > 0$ still leaves a nonzero wedge between the way revenue and fuel commitments age, so age dependence can remain. What kills the interaction is not "no embodied growth" by itself, but the stronger requirement that per-period profits be age-invariant.

F Event Selection

This appendix describes the construction of the core event sample and the full calendar of climate-policy events.

F.1 Selection logic

The model's taxonomy of stranding mechanisms (Appendix D, Table 5) shows that only *cost-side* policies, those that directly raise the operating cost of existing dirty capital, generate a strong age × dirty interaction. Green subsidies and investment-side instruments operate through output prices or discount rates, producing weaker or reversed vintage effects. This distinction motivates a baseline sample restricted to events that directly govern or materially constrain EPA authority over power plant emissions.

The universe of candidate events is defined *ex ante*: we include every regulation or court decision that directly bears on power plant emissions or EPA's authority to regulate them. Specifically, our baseline sample of 19 events comprises:

1. **EPA regulatory actions** (10 events): proposed and final rules, and executive orders directing review, that impose, tighten, or roll back emissions limits on existing power plants (Endangerment Finding proposed, CSAPR finalized, Clean Power Plan announced, CPP proposed rule, CPP finalized, Trump EO to review CPP, ACE rule proposed, ACE rule finalized, EPA Power Plant Rule proposed, EPA Power Plant Rule finalized).

2. **Supreme Court proceedings** (9 events): certiorari grants, oral arguments, and decisions that expand or constrain EPA's authority over power plant emissions, including the Clean Air Act (*Massachusetts v. EPA*: cert granted, oral argument), EPA's cost-consideration obligation (*Michigan v. EPA*: cert granted, oral argument, decision), the major questions doctrine (*West Virginia v. EPA*: cert granted, oral argument, decision), and the Good Neighbor Rule (*Ohio v. EPA*: stay granted).

There is no cherry-picking on confounders: we retain all events that satisfy the selection criterion, including those with severe confounders ($Q = 3$). Our identification relies on the Old−Young *spread*, which differences out common macro shocks that hit all utilities equally. Leave-one-out analysis confirms that no single event drives the results.

We exclude events that operate through non-cost-side channels:

- **Subsidy events** (IRA, Infrastructure Act, PATH Act): the model predicts a weak or absent vintage interaction for policies that improve clean alternatives without directly penalizing dirty operating costs. These events serve as a *falsification sample*.
- **General political events** (elections, Paris Agreement, Biden climate EO): these bundle expectations across multiple policy channels and are not cleanly attributable to cost-side mechanisms.
- **Other non-baseline regulations** (for example MATS mercury standards): these may matter for utilities, but they do not fit the paper's baseline event definition centered on rules and court decisions that directly govern the current power-plant emissions channel under study.
- **Non-utility-sector events** (LNG Export Pause, Methane Rule): these target oil and gas operations rather than power plant operating costs.
- **Legislative events** (Waxman–Markey): while this bill would have established a cost-side cap-and-trade system, it passed only the House and never became law; the cost-side mechanism remained hypothetical.

Two events warrant special comment. First, *Michigan v. EPA*: while the underlying case concerns mercury regulation (MATS), the Supreme Court's holding (that EPA must consider costs before regulating) constrains EPA's authority broadly, including its ability to impose CO_2 regulations on power plants. We include all three *Michigan v. EPA* proceedings (cert granted, oral argument, decision) on these grounds. Second, *Loper Bright Enterprises v. Raimondo*: while the overturning of Chevron deference materially constrains EPA's authority, the event is excluded because the daily Economic Policy Uncertainty index (Baker et al., 2016) shows the uncertainty spike concentrated in the days *before* the decision, suggesting the market had largely anticipated the ruling by the announcement date.

F.2 Event quality scores

Each event is scored 1–3 for confounder severity: 1 = no major or severe confounders across political, economic, and market dimensions; 2 = moderate confounders or clustering with other events; 3 = at least one severe confounder (e.g., simultaneous financial crisis, FOMC

announcement, or energy-price shock). Quality scores are used to construct robustness subsamples (e.g., restricting to $Q \leq 2$ or $Q = 1$ events) but do not determine inclusion in the core sample.

F.3 Policy uncertainty validation

We validate the information content of the baseline events using the daily Economic Policy Uncertainty (EPU) index of Baker et al. (2016). For each event we compute a *front-run score*

$$F_i = \max_{t \in [-5,-1]} |z_t| - \max_{t \in [0,+4]} |z_t|,$$

where z_t is the standardized daily EPU value relative to its full-sample mean and standard deviation, and event day 0 is the announcement date. A positive F_i indicates that the largest EPU spike occurred *before* the event, consistent with market anticipation or information leakage, while a negative F_i indicates a genuine post-announcement surprise. Among the 19 baseline events, the median F_i is negative, confirming that these events are predominantly unanticipated policy shocks. The five excluded EPA and Supreme Court events (Table 6) are contaminated by contemporaneous shocks to energy markets, climate policy news, or broad financial conditions; the 19 baseline events are the uncontaminated remainder.

F.4 Event calendar

Table 6 lists all EPA regulatory actions and Supreme Court proceedings bearing on power plant emissions or EPA authority over them. Events included in the baseline sample are marked with •; excluded events show the reason for exclusion. Events are excluded when contaminated by contemporaneous shocks to energy markets, climate policy news, or broad financial conditions that would materially affect utility returns independently of the policy event. The 19 baseline events are exactly the set that remains after these exclusions.

F.5 Broad market event calendar

Table 7 lists the candidate events for the broad-market event study of Section 4.3. The candidate set comprises all climate-policy events with broad market relevance, a wider universe than the utility-specific EPA and Supreme Court proceedings in Table 6. Events are included if $F_i < 0$ (not anticipated), quality ≤ 2, and not operating through the subsidy channel.

F.6 Robustness to event-set composition

Table 8 examines how the Generator age coefficient varies as events are added one by one in ascending F_i order (least front-run first). Starting from the six highest-surprise events and building up to the full baseline of 19, the fossil vintage coefficient in the $(0, 5)$ window is positive and significant at every step. The $(0, 10)$ window is significant for the highest-surprise subsets, weakens as noisier events are added, and recovers at the full baseline. Solar+wind vintage is insignificant throughout, consistent with the clean-side benchmark.

G Capital Age Measures from Compustat

This appendix discusses the construction and interpretation of accounting-based capital age measures, and surveys their use in the finance and economics literature.

G.1 Definitions

We construct four measures of capital age from Compustat annual data. Let PPEGT denote gross property, plant, and equipment (historical cost), PPENT denote net PP&E (gross minus accumulated depreciation), DPACT denote accumulated depreciation, DP denote annual depreciation expense, and AM denote amortization of intangibles. The approximate balance-sheet identity is DPACT $\approx$ PPEGT $-$ PPENT.

1. **Depreciation ratio** (primary measure):

$$\text{DepRatio}_{it} = \frac{\text{PPEGT}_{it} - \text{PPENT}_{it}}{\text{PPEGT}_{it}} = \frac{\text{DPACT}_{it}}{\text{PPEGT}_{it}} \in [0, 1].$$

   This captures the fraction of the capital stock's useful life that has been consumed. It is a standard financial analysis ratio, widely known in the CFA curriculum as the "average age ratio" or "accumulated depreciation to fixed assets ratio."

2. **Capital age** (years):

$$\text{CapAge}_{it} = \frac{\text{DPACT}_{it}}{\text{DP}_{it} - \text{AM}_{it}}.$$

   Under straight-line depreciation, this estimates the average number of years the capital stock has been in service.

3. **Useful life** (years):

$$\text{UsefulLife}_{it} = \frac{\text{PPEGT}_{it}}{\text{DP}_{it} - \text{AM}_{it}}.$$

   This estimates the total implied useful life of the capital stock.

4. **Depreciation rate**:

$$\text{DepRate}_{it} = \frac{\text{DP}_{it}}{\text{PPENT}_{it}}.$$

   The annual rate of capital consumption. Higher values indicate shorter-lived or faster-depreciating capital.

These measures are linked by the identity:

$$\text{DepRatio} = \frac{\text{CapAge}}{\text{UsefulLife}}.$$

The depreciation ratio tells us where a firm sits on the lifecycle of its capital: zero for brand-new assets, one for fully depreciated assets.

G.2 Related Literature

The use of PP&E-based ratios to characterize the age and maturity of physical capital has a long history in corporate finance.

Debt maturity matching. Stohs and Mauer (1996) introduce PP&E-based asset maturity measures to study how firms match debt maturity to asset life. They define asset maturity as PPEGT/DP (our useful life measure) and document a strong positive relation between asset maturity and debt maturity. This approach has been widely adopted in the capital structure literature (Barclay and Smith, 1995; Billett et al., 2007; Custódio et al., 2013). Most recently, Geelen et al. (2024) build a theory of financing in which leverage and debt maturity are negatively related to capital age, with strong Compustat-based empirical support.

Capital age and asset pricing. A growing literature documents a "capital age premium": firms with older capital earn higher expected stock returns. Lin et al. (2020) find that a long–short portfolio sorted on firm-level capital age earns 9% annually, which they explain with a model in which costly technology adoption makes old-capital firms riskier. Li et al. (2023) document a similar premium (Sharpe ratio of 0.44) and attribute it to investor learning about productivity exposure. Both papers measure capital age in years using a recursive weighted-average approach initialized from Compustat's accumulated depreciation.

Capital reallocation and vintage effects. Eisfeldt and Rampini (2006, 2007) study how capital is reallocated across firms and show that financially constrained firms disproportionately purchase used (older) capital. Ma et al. (2022) document this pattern at scale using 1.56 million equipment transactions: young firms systematically acquire old capital from more mature firms. Campello et al. (2024) show that uncertainty delays creative destruction by slowing both the acquisition of new capital and the disposal of old capital. Darmouni and Sutherland (2024) study how supply disruptions for new capital shift investment toward secondary markets.

Embodied technical change. The macro literature on vintage capital emphasizes that newer capital embodies superior technology (Greenwood et al., 1997; Bahk and Gort, 1993). This provides the theoretical foundation for why capital age matters for firm value: older capital is less productive and more exposed to obsolescence risk. Hulten and Wykoff (1981) provide the foundational empirical work on the divergence between accounting (book) depreciation and economic depreciation.

Depreciation rates in production-based models. The investment-based asset pricing literature uses Compustat depreciation rates as key model inputs. Zhang (2005) uses an exogenously calibrated depreciation rate in a neoclassical model that explains the value premium. İmrohoroğlu and Tüzel (2014) estimate firm-level total factor productivity from Compustat using depreciation rates within the perpetual inventory method and show that low-productivity firms earn higher returns.

Accounting measures of asset age. Curtis et al. (2015) study how the historical cost basis of PP&E interacts with asset age to affect the informativeness of financial ratios. They measure asset age in years as accumulated depreciation divided by depreciation expense and find that asset turnover ratios are more persistent for firms with older assets.

G.3 Measurement Considerations

Several caveats apply to all accounting-based capital age measures:

- *Accounting vs. economic depreciation.* Compustat reports book depreciation (typically straight-line for financial reporting), while economic depreciation generally follows a declining-balance pattern (Hulten and Wykoff, 1981). This biases the depreciation ratio toward the midpoint of useful life. For cross-sectional analyses the ranking of firms is likely preserved, but levels should be interpreted with caution.
- *Acquisitions.* Purchase accounting revalues acquired PP&E to fair value, resetting accumulated depreciation. This makes acquired capital appear younger than it is.
- *Impairments and write-downs.* Asset impairments create discontinuous jumps in the depreciation ratio unrelated to physical aging.
- *Fully depreciated assets.* The depreciation ratio is bounded at one even if fully depreciated assets remain in productive use, creating a ceiling effect.
- *Cross-industry variation.* Capital intensity and useful lives vary across industries, requiring industry fixed effects or industry-adjusted measures in cross-sectional tests.

Our primary specification uses the recursive capital age measure of Lin et al. (2020) and Geelen et al. (2024), described in Section 4. The recursive measure responds immediately to investment, directly maps to the model's average age $\bar{j}$, and has been shown to predict the cross-section of stock returns (9% annual long-short spread) and leverage dynamics. We use the static capital age in years as a secondary measure and report results for the depreciation ratio and useful life in robustness checks.

G.4 Validation: Recursive Capital Age vs. Generator Age

Figure 6 validates the recursive capital age measure against physical asset age data from the EIA. For 45 publicly traded electric utilities with matched EIA generator data, we plot year-month-demeaned recursive capital age (computed from Compustat PP&E using equation 19) against year-month-demeaned capacity-weighted generator age (computed from EIA Form 860 operating dates). We restrict to firm-months with recursive capital age above five years, which retains 98% of observations and all 45 firms; below this threshold the accounting proxy is dominated by non-generation PP&E (transmission and distribution infrastructure) and is uninformative about generator vintage.

The binscatter confirms a positive cross-sectional relationship ($\beta = 0.99$, $r = 0.21$, $p < 0.001$): utilities with older capital stocks by the accounting measure also operate physically

older generating fleets. The slope near unity indicates that a one-year increase in accounting capital age corresponds to roughly a one-year increase in physical generator age. This validation supports interpreting cross-sectional variation in recursive capital age as reflecting genuine differences in the vintage composition of physical assets, the key object in the putty-clay model.

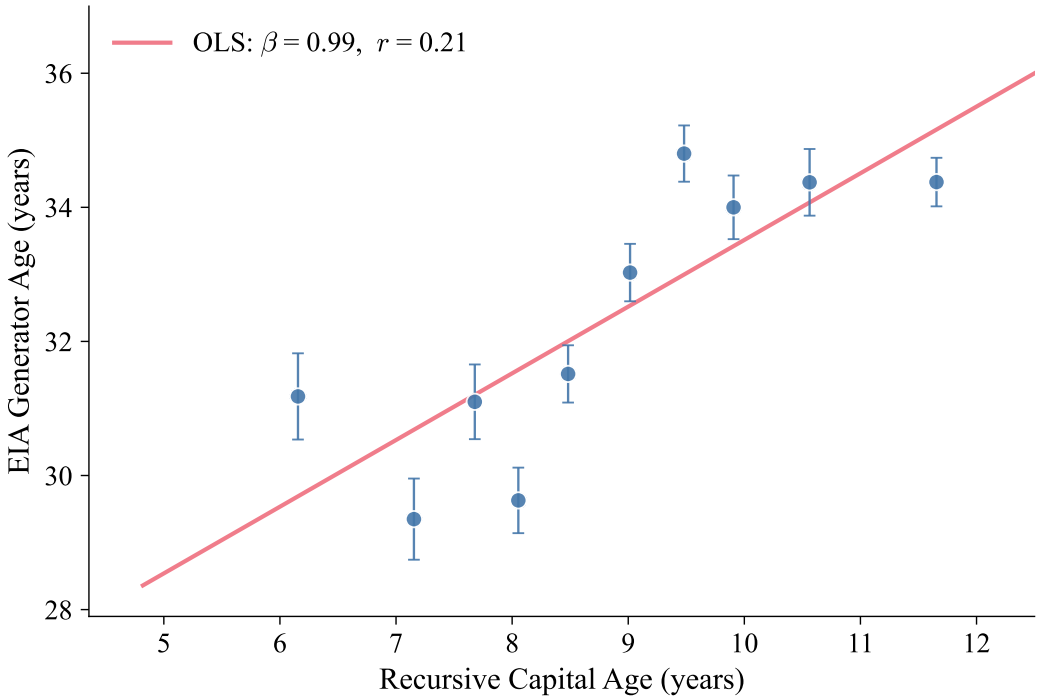


Figure 6: Recursive capital age vs. EIA generator age: electricity sector binscatter. *Notes:* Each dot represents the mean of a decile bin; error bars show 95% confidence intervals. The sample covers 8,567 firm-month observations for 45 utilities with recursive capital age $\geq$ 5 years (98% of matched observations). Both axes are demeaned by year-month to absorb common time trends; axis labels show levels (demeaned values plus sample means). Recursive capital age is computed from Compustat PP&E using the Lin et al. (2020) recursion (equation 19); generator age is the capacity-weighted mean operating age of all generators reported in EIA Form 860. The OLS slope is $\hat{\beta} = 0.99$ with Pearson $r = 0.21$ ($p < 0.001$).

H Utility Event Study: Data and Methodology

This appendix describes the construction of the utility-sector event study in Section 3.

H.1 Sample

The sample consists of 46 publicly traded electric utilities with matched plant-level data from the U.S. Energy Information Administration (EIA). We construct the EIA–CRSP crosswalk by matching utility holding companies in EIA Form 860 (generator-level data) to CRSP

permanent security numbers (PERMNOs) via Compustat identifiers. The crosswalk links each CRSP PERMNO to the set of EIA plant codes operated by the corresponding utility holding company. Sample firms span investor-owned utilities with significant generating capacity, excluding pure transmission-and-distribution companies and independent power producers without public equity.

Daily stock returns and market capitalizations are obtained from CRSP for the period 2005–2024. Fama–French three-factor returns are from Ken French's data library via WRDS.

H.2 Capital vintage measures

For each utility, we compute technology-specific capital age measures from EIA Form 860 generator-level data. Each generator reports its operating date, nameplate capacity (MW), and primary fuel type. We classify generators into three technology groups:

- **Fossil:** coal, natural gas, petroleum, and other fossil-fuel generators.
- **Green (all renewables):** wind, solar (photovoltaic and thermal), hydro, geothermal, biomass, and other renewable sources.
- **Renewable (solar + wind):** the subset of green generators using solar or wind technology only. This is the definition of "renewable vintage" used in the main text and figures; it excludes hydro, nuclear, and other legacy technologies whose vintage profiles reflect different investment dynamics.

For each firm i and event date e, the technology-specific capital age is the capacity-weighted average age of active generators in each group:

$$\text{age}_{i,e}^{k} = \frac{\sum_{g \in \mathcal{G}_{i,e}^{k}} \text{capacity}_g \times (\text{event year} - \text{operating year}_g)}{\sum_{g \in \mathcal{G}_{i,e}^{k}} \text{capacity}_g},$$

where $k \in \{\text{fossil}, \text{green}, \text{renewable}\}$ and $\mathcal{G}_{i,e}^{k}$ is the set of active generators of type k for firm i at event e. Because the EIA data are updated annually, we use the most recent vintage panel observation on or before the event date.

H.3 Orthogonalization

A firm with an old overall fleet tends to have both older fossil and older renewable generators, creating a mechanical correlation between technology-specific ages. To isolate variation in the *composition* of the vintage (for example, a firm that has selectively modernized its renewable fleet while retaining old fossil plants), we orthogonalize each technology-specific age on total fleet age.

For each event e, we regress $\text{age}_{i,e}^{k}$ on total fleet age $\text{age}_{i,e}^{\text{total}}$ across firms and take the residual:

$$\text{age}_{i,e}^{k,\perp} = \text{age}_{i,e}^{k} - \hat{\alpha}_e - \hat{\beta}_e \cdot \text{age}_{i,e}^{\text{total}}.$$

The orthogonalized measure $\text{age}^{k,\perp}$ captures whether a firm's fossil (or green) fleet is older or

younger *than expected given its overall capital maturity.* We orthogonalize within each event to allow the firm composition and the correlation of technology-specific ages with total age to vary across events. In the sample, orthogonalized fossil age has a cross-sectional correlation of 0.78 with raw fossil age, confirming that the orthogonalization preserves the rank ordering while removing the total-age component.

H.4 Portfolio construction

For each event, we sort firms into terciles by $\text{age}_{i,e}^{k,\perp}$ and form "Old" (top tercile) and "Young" (bottom tercile) portfolios. Value-weighted portfolios use day -1 market capitalization as weights. With 46 firms and tercile sorts, the Old and Young portfolios each contain approximately 15 firms per event.

H.5 CAAR estimation and inference

Let $R_{e,t}^{O}$ and $R_{e,t}^{Y}$ denote the portfolio returns on day t relative to event e for the Old and Young portfolios, and let $d_e \in \{+1, -1\}$ denote the event direction ($+1$ for tightenings, -1 for loosenings). The signed abnormal return difference for event e on day t is:

$$AR_{e,t} = d_e \cdot \left(R_{e,t}^{O} - R_{e,t}^{Y} \right).$$

The cross-event average abnormal return and CAAR are:

$$\overline{AR}_t = \frac{1}{E} \sum_{e=1}^{E} AR_{e,t}, \qquad \text{CAAR}_t = \sum_{s=-20}^{t} \overline{AR}_s.$$

To remove any pre-event level, we demean by the average CAAR over days $[-5, -1]$:

$$\widetilde{\text{CAAR}}_t = \text{CAAR}_t - \frac{1}{5} \sum_{s=-5}^{-1} \text{CAAR}_s.$$

Cross-event covariance matrix SEs. Because events may share common shocks (macroeconomic news, sector trends) that induce cross-event correlation in abnormal returns, we estimate the full $T \times T$ cross-event covariance matrix of the AAR vector:

$$\hat{\Omega}_{s,u} = \frac{1}{E(E-1)} \sum_{e=1}^{E} \left(AR_{e,s} - \overline{AR}_s \right) \left(AR_{e,u} - \overline{AR}_u \right).$$

The variance of $\widetilde{\text{CAAR}}_t$ is $\mathbf{w}_t' \hat{\Omega}\, \mathbf{w}_t$, where $\mathbf{w}_t$ is the weight vector encoding both cumulation (from day -20 to t) and the demeaning adjustment. This approach, which follows **?**, accounts for both serial correlation in AAR and cross-event dependence.

BMP standardization. As a robustness check, we also report BMP-standardized test statistics (**?**). For each event e, we estimate the pre-event volatility $\hat{\sigma}_e$ from the daily signed

return difference $AR_{e,t}$ over the window $t \in [-120, -21]$ (at least 20 observations required). The standardized abnormal return is $SAR_{e,t} = AR_{e,t}/\hat{\sigma}_e$. We then form the average SAR and its CAAR, with inference based on the cross-event covariance matrix of the standardized series. This accounts for the fact that some events occur in high-volatility periods and would otherwise receive excessive weight in the cross-event average.

I Capital-Intensity Classification

Figure 7 displays the year-by-year share of market capitalization and firm count held by capital-intensive industries under the balanced SIC-2 classification. The partition is approximately even throughout the sample: capital-intensive firms account for 49–57% of firms and 50–60% of total market capitalization, with no systematic size imbalance between the two groups.

Figure 8 shows the distribution of recursive capital age (non-winsorized) separately for capital-intensive and capital-light firms, at the start (2005–2007) and end (2022–2024) of the sample. The distributions are right-skewed in both groups, with most mass between 2 and 8 years. Capital-intensive firms exhibit a heavier right tail, consistent with longer-lived assets and lumpier investment. Importantly, the age distributions are broadly stable across the two periods, indicating that the cross-sectional variation in capital age reflects persistent industry-level depreciation patterns rather than a transient investment cycle.

J Pooled OLS Robustness

All-industries specification. Table 9 replicates the specification progression of Table 3 using all industries rather than restricting to capital-intensive sectors. The interaction is negative but insignificant ($t = -1.36$), consistent with the putty-clay channel operating primarily through physical capital.

Alternative capital-intensity cutoff. Table 10 uses the 25th-percentile cutoff of SIC-2 gross PP&E/total assets (rather than the median), which retains more industries as "capital-intensive." The interaction remains significant ($t = -2.68$).

Emissions scope robustness. Table 11 shows that the emissions × capital age interaction is significant across all four Trucost emissions scopes: Scope 1 (direct), Scope 2 (purchased electricity), Scope 3 (upstream supply chain), and Scope 1+2 (combined).

Finer industry fixed effects. Table 12 replaces SIC-2 × year-month fixed effects with SIC-3 × year-month fixed effects. The interaction coefficients are virtually unchanged.

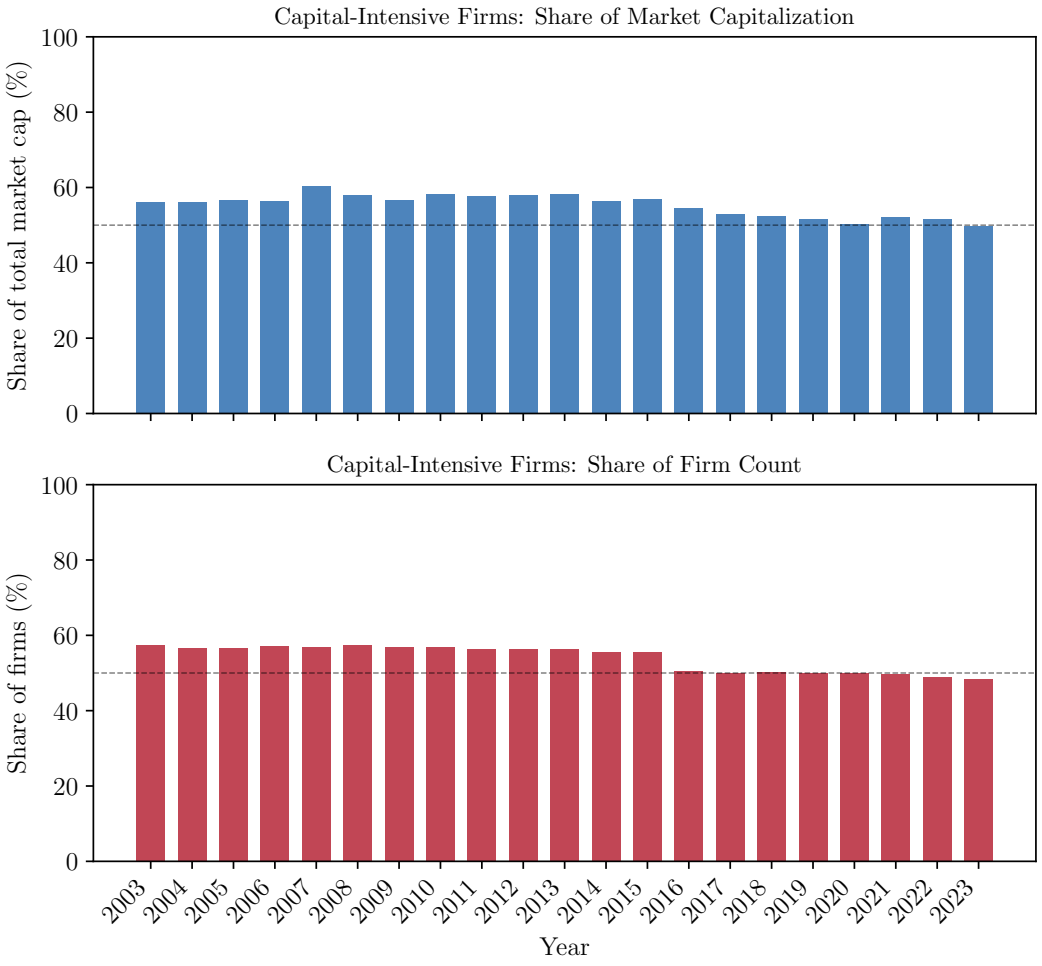


Figure 7: Share of market capitalization and firm count held by capital-intensive industries. *Notes:* Top panel: share of total market capitalization; bottom panel: share of firm count. The classification assigns entire SIC-2 industries to the capital-intensive group if their 1985–1995 median gross PP&E/total assets exceeds the sample-wide median (0.41). The dashed line marks a 50% share.

K Fama–MacBeth Robustness

The cross-sectional regressions in Section 4.2 use pooled OLS with two-way clustered standard errors, following Bolton and Kacperczyk (2021). Pooled OLS stacks all firm-months and runs a single regression, so months with more firms (later in the sample, as Trucost coverage expands) receive proportionally greater weight. As a robustness check, we re-estimate all specifications using Fama and MacBeth (1973) cross-sectional regressions, which give each month equal weight regardless of the number of firms. This addresses the concern that the Trucost coverage expansion (from roughly 400 firms in 2005 to over 1,400 by 2024) could mechanically influence the pooled estimates.

The Fama–MacBeth procedure runs a separate cross-sectional OLS regression each month, with SIC-2 industry dummies included as regressors, and collects the time series of slope

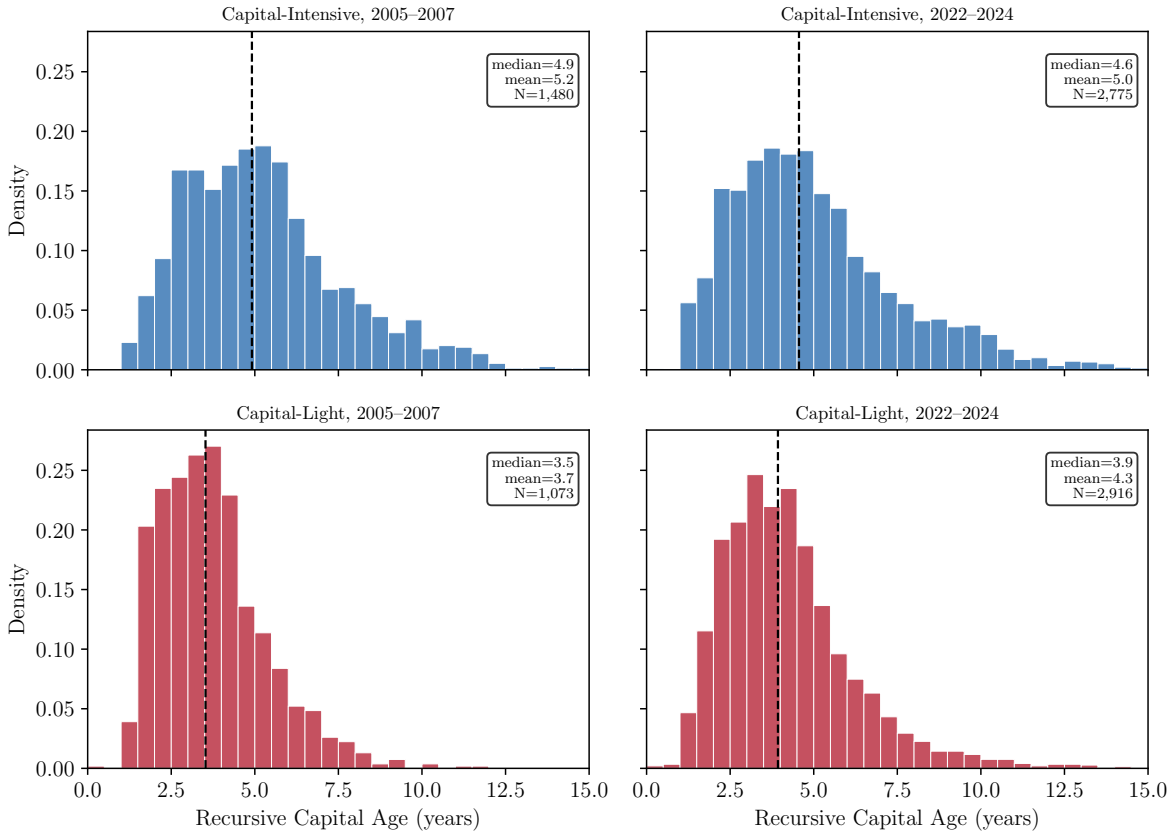


Figure 8: Distribution of recursive capital age by capital intensity.

Notes: Capital-intensive firms (top) and capital-light firms (bottom), at the start (left, 2005–2007) and end (right, 2022–2024) of the sample. Dashed vertical lines mark the within-group median. Age is computed using the **?** recursion on the full Compustat history and is not winsorized.

coefficients. Reported coefficients are time-series averages; t-statistics use Newey–West standard errors with 12 lags to account for serial correlation in the monthly slopes. Including SIC-2 dummies in each monthly cross-section is equivalent to allowing SIC-2 intercepts to vary freely by month; that is, it corresponds to SIC-2 × year-month fixed effects in a pooled representation, the same stringency as the baseline pooled OLS tables. Column (6) replaces SIC-2 with SIC-3 industry dummies, providing a genuinely more demanding specification that absorbs within-SIC-2 industry heterogeneity.

Main specification. Table 13 replicates the six-column specification progression of Table 3 using Fama–MacBeth, with both the log-level (Panel A) and log-ratio (Panel B) emissions measures. The interaction coefficients are negative throughout and of comparable magnitude to the pooled estimates. With the scaled emissions measure (Panel B), the interaction remains significant at the 5% level in the baseline specification (column 3, -0.0011, $t = -2.10$).

Table 14 repeats the Fama–MacBeth progression for all industries. The interaction is weaker

than in capital-intensive industries but directionally consistent.

Robustness across samples. Table 15 presents the Fama–MacBeth version of the capital-intensity split table. The key patterns survive: the interaction concentrates in capital-intensive industries and in the post-Paris period.

L Quadruple Interaction and Paris-Timing Sensitivity

The quadruple interaction regression decomposes the carbon premium into four groups: capital-intensive × post-Paris, capital-intensive × pre-Paris, capital-light × post-Paris, and capital-light × pre-Paris. This allows the emissions × age gradient to differ across both industry type and time period in a single regression.

Aggregate quadruple interaction. Table 16 reports the aggregate quadruple interaction, with the four-way term $\phi \times A \times \text{CapInt} \times \text{Post}$ as the key coefficient. The four-way interaction is -0.0039 ($t = -2.88$): the emissions × age gradient is significantly more negative for capital-intensive firms in the post-Paris period.

Quintile decomposition. Table 17 reports the emissions coefficient for each capital-age quintile within each of the four groups. These are the numerical values underlying Figure 3. The monotonic decline from Q1 to Q5 in the post-Paris capital-intensive column ($+100.2 \rightarrow +10.0$ bps) confirms that the age gradient is not driven by outliers or a single quintile.

Paris-timing sensitivity. Figure 4 in the main text examines whether the result depends on the specific choice of January 2016 as the post-Paris cutoff. The four-way coefficient is negative at every cutoff and statistically significant from 2010 onward, strengthening as the cutoff moves later. The stranding premium is not a knife-edge artifact of the Paris Agreement; rather, it reflects a trend in the pricing of capital-age-dependent stranding risk that intensifies over the sample period, with a notable further strengthening around 2021 coinciding with the Biden administration's regulatory agenda.

Figure OA.1 (Online Appendix) repeats the exercise with COVID-period controls, confirming that the 2021–2022 strengthening reflects policy repricing rather than pandemic-related return mechanics.

M Pooled Daily OLS Event Study: Additional Specifications

This appendix reports additional specifications for the pooled daily OLS event study of Section 4.3. See Section 4.3 for the specification, event selection, and capital-intensive results.

All industries. Table 18 reports the pooled daily OLS for the full sample (all industries). The triple interaction Em × RecAge × EventDay is positive and significant, confirming that the event-study pattern extends beyond capital-intensive industries, though with less statistical power.

Capital-light industries (placebo). Table 19 restricts to capital-light industries. The triple interaction is insignificant across all risk adjustments ($t = 1.19$–1.63), confirming that the stranding channel operates through physical capital.

Subsidy falsification. Table 20 uses five subsidy events (quality ≤ 3, $F_i \leq -0.04$) on the capital-intensive sample. The triple interaction is insignificant in the raw and CAPM specifications ($t < 1$) but marginally significant after FF3/FF4 adjustment ($t \approx 1.85$–1.88).

N Fossil Fuel Intensity per Worker

Figure 9 contrasts two measures of industrial fossil fuel intensity over the last four decades. The left panel plots fossil fuel consumption per goods-producing worker: the ratio rises from roughly 0.73 trillion BTU per thousand workers in the mid-1980s to approximately 1.0 by the late 2000s, a 37% increase (0.63% per year), and has remained near that level since. The upward trend reflects the fact that goods-producing employment has declined (driven by automation and offshoring) faster than industrial fossil fuel consumption, so each remaining worker is associated with an increasing quantity of fossil fuel use. The right panel plots fossil fuel consumption per unit of real industrial output (the Federal Reserve's Industrial Production Index): this ratio falls by roughly 45% over the same period, from 335 to 185 trillion BTU per index point. Industry is becoming more energy-efficient per unit of output, but the fossil fuel burden per worker is rising.

Table 21 decomposes the 1998–2018 changes in fossil fuel intensity across 21 NAICS-3 manufacturing subsectors using a symmetric shift-share decomposition. For the fuel-per-worker ratio (F/L), the aggregate increase is dominated by between-industry reallocation (74%): light-fossil sectors such as textiles and apparel shed workers far faster than heavy-fossil sectors such as petroleum refining and chemicals, raising the employment-weighted average. For the fuel-per-real-output ratio (F/Y), the aggregate decline is driven overwhelmingly by within-industry efficiency gains (114%), partially offset by modest between-industry reallocation toward fuel-intensive sectors.

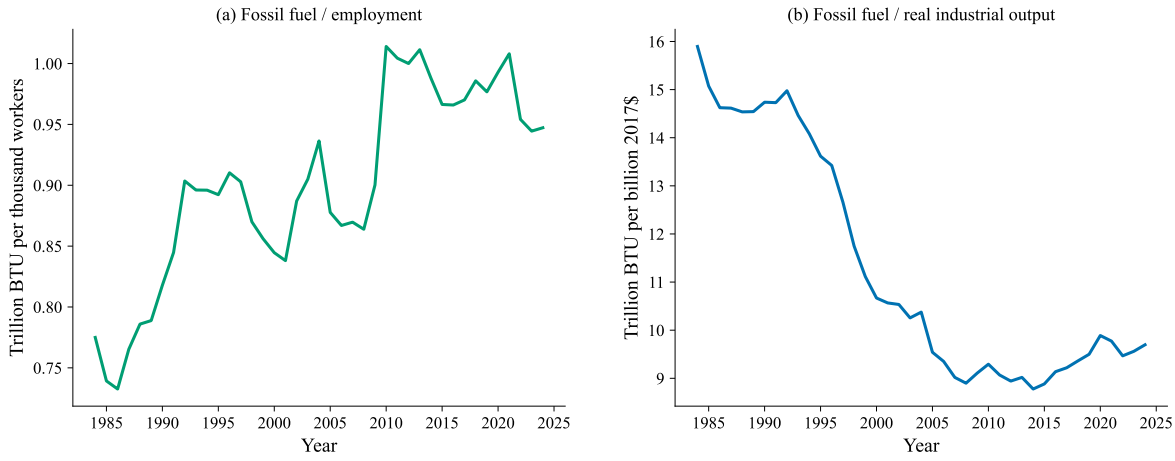


Figure 9: Industrial fossil fuel intensity: per worker (left) vs. per real dollar of output (right), 1984–2024.

Notes: The numerator in both panels is total fossil fuel consumption by the U.S. industrial sector (coal, petroleum, and natural gas) in trillion BTU from the EIA Monthly Energy Review (series FFICBUS). In the left panel, the denominator is average annual goods-producing employment in thousands from the BLS Current Employment Statistics survey (FRED series USGOOD). In the right panel, the denominator is real manufacturing output in trillions of chained 2017 dollars, constructed by rescaling the Federal Reserve's Industrial Production Index (FRED series INDPRO) using BEA real manufacturing GDP at the 2017 benchmark ($2.11 T).

Table 6: Event calendar: EPA regulatory actions and Supreme Court proceedings

Event	Date	Dir	Type	Core / Exclusion reason
Baseline cost-side sample (19 events)				
Mass v. EPA: cert granted	2006-06-26	+	Supreme Court	•
Mass v. EPA: oral argument	2006-11-29	+	Supreme Court	•
Endangerment Finding proposed	2009-04-17	+	EPA regulation	•
CSAPR finalized	2011-07-06	+	EPA regulation	•
Clean Power Plan announced	2013-06-25	+	EPA regulation	•
CPP proposed rule	2014-06-02	+	EPA regulation	•
Michigan v. EPA: cert granted	2014-11-25	−	Supreme Court	•
Michigan v. EPA: oral argument	2015-03-25	−	Supreme Court	•
Michigan v. EPA (MATS struck down)	2015-06-29	−	Supreme Court	•
CPP finalized	2015-08-03	+	EPA regulation	•
Trump EO to review CPP	2017-03-28	−	EPA regulation	•
ACE rule proposed	2018-08-21	−	EPA regulation	•
ACE rule finalized	2019-06-19	−	EPA regulation	•
WV v. EPA: cert granted	2021-10-29	−	Supreme Court	•
WV v. EPA: oral argument	2022-02-28	−	Supreme Court	•
West Virginia v. EPA	2022-06-30	−	Supreme Court	•
EPA Power Plant Rule proposed	2023-05-11	+	EPA regulation	•
EPA Power Plant Rule Finalized	2024-04-25	+	EPA regulation	•
Ohio v. EPA (Good Neighbor stayed)	2024-06-27	−	Supreme Court	•
Excluded EPA and Supreme Court events				
Massachusetts v. EPA	2007-04-02	+	Supreme Court	Iran-UK sailor crisis (oil spike)
EPA Endangerment Finding	2009-12-07	+	EPA regulation	Copenhagen COP15 opened same day
MATS proposed rule	2011-03-16	+	EPA regulation	Fukushima nuclear disaster
EPA MATS finalized	2011-12-21	+	EPA regulation	Buildup to Strait of Hormuz oil crisis
Supreme Court stays CPP	2016-02-09	−	Supreme Court	Scalia death (Feb 13); oil at $27

Notes: The table lists all EPA regulatory actions and Supreme Court proceedings that directly bear on power plant CO_2 emissions or EPA's authority to regulate them. Direction indicates whether the event tightens (+) or loosens (−) regulation. Events are excluded when contaminated by contemporaneous shocks to energy markets, climate policy news, or broad financial conditions that would materially affect utility returns independently of the policy event. MATS events regulate mercury emissions rather than CO_2 but are included in the candidate set because they raise utility operating costs through a related regulatory channel. Subsidy events (IRA, Infrastructure Act, PATH Act), elections, legislative actions (Waxman–Markey), and non-EPA regulations (FERC) are excluded by design and not shown; they serve as falsification samples.

Table 7: Broad market event calendar

Event	Date	Dir	Type	F_i	Core / Exclusion reason
Baseline broad cost-side sample (8 events)					
Trump 2016 election	2016-11-09	−	General	−1.91	•
EPA Final Methane Rule	2023-12-04	+	EPA regulation	−1.40	•
Trump 2024 election	2024-11-06	−	General	−1.07	•
Endangerment Finding proposed	2009-04-17	+	EPA regulation	−0.57	•
Loper Bright (Chevron overturned)	2024-06-28	−	Supreme Court	−0.50	•
Biden LNG Export Pause	2024-01-26	+	EPA regulation	−0.39	•
Mass v. EPA: oral argument	2006-11-29	+	Supreme Court	−0.02	•
Paris Agreement adopted	2015-12-14	+	General	−0.01	•
Excluded events					
Loper Bright: cert granted	2023-05-01	−	Supreme Court	−2.40	First Republic bankruptcy
Infrastructure: Senate passage	2021-08-10	+	Subsidy	−1.68	Subsidy channel
Manchin kills climate bill	2022-07-15	−	Subsidy	−0.05	Subsidy channel
Mass v. EPA: cert granted	2006-06-26	+	Supreme Court	+0.09	Anticipated ($F_i > 0$)
Massachusetts v. EPA (decision)	2007-04-02	+	Supreme Court	+0.02	Subprime crisis
EPA Endangerment Finding	2009-12-07	+	EPA regulation	+0.16	Copenhagen COP15
Manchin kills BBB	2021-12-20	−	Subsidy	+0.15	Subsidy; Omicron
US rejoins Paris Agreement	2021-02-19	+	General	+0.09	Texas winter storm
IRA announced	2022-07-28	+	Subsidy	+0.39	Subsidy channel
Biden climate EO	2021-01-27	+	General	+0.43	GameStop; FOMC
Infrastructure bill signed	2021-11-15	+	Subsidy	+0.65	Subsidy; largely priced in
Loper Bright: oral argument	2024-01-17	−	Supreme Court	+0.76	Anticipated ($F_i > 0$)
Waxman–Markey: committee passage	2009-05-21	+	Legislative	+0.81	Legislative; anticipated ($F_i > 0$)
Waxman–Markey passes House	2009-06-29	+	Legislative	+1.96	Legislative; anticipated ($F_i > 0$)
Biden 2020 election	2020-11-09	+	General	+2.10	Pfizer vaccine same day

Notes: Direction indicates whether the event tightens (+) or loosens (−) climate policy. F_i is the front-run score (Section F.3): negative = post-announcement surprise. Events are included if $F_i < 0$, quality ≤ 2, and not operating through the subsidy channel. Subsidy events serve as a separate falsification sample (Appendix Table 20).

Table 8: Generator age coefficients: recursive add-one-in by ascending F_i

N	$F_{i,\max}$	Fossil $(0,5)$	Fossil $(0,10)$	S+W $(0,5)$	S+W $(0,10)$
6	-0.478	+0.0516 (0.0143)***	+0.0610 (0.0219)***	-0.0710 (0.2731)	-0.0885 (0.3687)
7	-0.461	+0.0502 (0.0129)***	+0.0481 (0.0216)**	-0.0228 (0.2280)	+0.0294 (0.3267)
8	-0.378	+0.0394 (0.0158)**	+0.0385 (0.0214)*	-0.0096 (0.1600)	+0.0046 (0.2300)
9	-0.277	+0.0394 (0.0146)***	+0.0371 (0.0201)*	+0.0211 (0.1507)	+0.0601 (0.2613)
10	-0.276	+0.0434 (0.0139)***	+0.0403 (0.0191)**	+0.0414 (0.1428)	-0.0092 (0.2300)
11	-0.258	+0.0429 (0.0129)***	+0.0299 (0.0217)	+0.0799 (0.1427)	+0.0510 (0.2550)
12	-0.016	+0.0397 (0.0134)***	+0.0239 (0.0223)	+0.0799 (0.1427)	+0.0510 (0.2550)
13	0.072	+0.0383 (0.0129)***	+0.0258 (0.0217)	+0.0992 (0.1359)	+0.0351 (0.2340)
14	0.081	+0.0333 (0.0127)***	+0.0214 (0.0202)	+0.0849 (0.1369)	-0.0003 (–)
15	0.087	+0.0314 (0.0127)**	+0.0225 (0.0201)	+0.0849 (0.1392)	-0.0003 (–)
16	0.131	+0.0367 (0.0131)***	+0.0264 (0.0193)	+0.0588 (0.1307)	-0.0418 (0.2322)
17	0.135	+0.0265 (0.0160)*	+0.0203 (0.0193)	-0.0056 (0.1400)	-0.1238 (0.2172)
18	0.146	+0.0293 (0.0157)*	+0.0263 (0.0196)	-0.0151 (0.1258)	-0.1379 (0.2089)
19	0.651	+0.0320 (0.0143)**	+0.0306 (0.0179)*	+0.0060 (0.1200)	-0.1389 (0.1828)

Events added recursively in ascending F_i order (lowest front-running first). Coefficients from industry-adjusted CAR regressions (equation 18) with event fixed effects and controls. Implied clustered SE in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 9: Older Dirty Capital Earns Lower Returns (All Industries)

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Emissions = ln(Scope 1 GHG)						
Emissions	0.0020*	0.0020*	0.0022*	0.0014	0.0075***	0.0017
	(0.0010)	(0.0010)	(0.0010)	(0.0009)	(0.0015)	(0.0010)
Capital age		0.0002	0.0005	0.0002	0.0009	0.0005
		(0.0011)	(0.0011)	(0.0012)	(0.0011)	(0.0010)
Emissions × Capital age			−0.0006	−0.0005	−0.0007*	−0.0006
			(0.0005)	(0.0005)	(0.0003)	(0.0005)
Obs	204,137	204,137	204,137	204,137	282,105	204,137
Panel B: Emissions = ln(Scope 1 GHG/Gross PP&E)						
Emissions	0.0014*	0.0013*	0.0013*	0.0009	0.0016**	0.0010
	(0.0006)	(0.0006)	(0.0007)	(0.0005)	(0.0007)	(0.0006)
Capital age		0.0002	0.0003	−0.0001	0.0027	0.0003
		(0.0011)	(0.0011)	(0.0011)	(0.0015)	(0.0009)
Emissions × Capital age			−0.0005	−0.0001	0.0002	−0.0004
			(0.0004)	(0.0003)	(0.0004)	(0.0004)
Obs	204,137	204,137	204,137	204,137	249,697	204,137
B&K controls	Yes	Yes	Yes	Yes	No	Yes
SIC-2 FE	Yes	Yes	Yes	No	Yes	—
Year-month FE	Yes	Yes	Yes	Yes	Yes	—
SIC-2 × year-month FE	No	No	No	No	No	Yes

Notes: All-industries version of Table 3. The dependent variable is the monthly excess stock return. Panel A measures emissions as ln(Scope 1 GHG); Panel B as ln(Scope 1 GHG/gross PP&E) (Compustat `ppegt`). All emissions and age variables are within-month z-scored. *Capital age* is recursive capital age (winsorized at the 1st and 99th percentiles) (Lin, Palazzo, and Yang 2020; Geelen, Hajda, and Morellec 2024). Columns (1)–(3) include year-month and SIC-2 fixed effects; column (4) drops industry fixed effects; column (5) drops all firm-level controls; column (6) replaces separate year-month and industry fixed effects with SIC-2 × year-month fixed effects. B&K controls: ln(market cap), book-to-market, momentum, ROE, leverage, investment/assets, ln(PPE), sales growth, EPS growth, HHI, market beta, and return volatility (Bolton and Kacperczyk, 2021). Annual accounting variables are mapped to monthly returns using the Fama–French June t+1 convention; monthly controls are lagged one month. Sample: all industries, January 2016 to December 2024. Two-way clustered standard errors (firm × calendar year) in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 10: Older Dirty Capital Earns Lower Returns (Broader Capital-Intensity Cutoff)

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Emissions = ln(Scope 1 GHG)						
Emissions	0.0031**	0.0030**	0.0037***	0.0042**	0.0082***	0.0032**
	(0.0011)	(0.0011)	(0.0011)	(0.0014)	(0.0017)	(0.0011)
Capital age		0.0007	0.0014	0.0015	0.0018	0.0013
		(0.0011)	(0.0012)	(0.0011)	(0.0012)	(0.0010)
Emissions × Capital age			−0.0016**	−0.0017***	−0.0017***	−0.0015**
			(0.0006)	(0.0005)	(0.0005)	(0.0006)
Obs	145,641	145,641	145,641	145,641	179,053	145,641
Panel B: Emissions = ln(Scope 1 GHG/Gross PP&E)						
Emissions	0.0020**	0.0019**	0.0021**	0.0023**	0.0030***	0.0018**
	(0.0007)	(0.0007)	(0.0007)	(0.0008)	(0.0009)	(0.0007)
Capital age		0.0006	0.0009	0.0007	0.0034*	0.0008
		(0.0011)	(0.0011)	(0.0010)	(0.0016)	(0.0010)
Emissions × Capital age			−0.0012**	−0.0011**	−0.0007	−0.0012*
			(0.0005)	(0.0005)	(0.0005)	(0.0006)
Obs	145,641	145,641	145,641	145,641	177,670	145,641
B&K controls	Yes	Yes	Yes	Yes	No	Yes
SIC-2 FE	Yes	Yes	Yes	No	Yes	—
Year-month FE	Yes	Yes	Yes	Yes	Yes	—
SIC-2 × year-month FE	No	No	No	No	No	Yes

Notes: The dependent variable is the monthly excess stock return. Panel A measures emissions as ln(Scope 1 GHG); Panel B as ln(Scope 1 GHG/gross PP&E) (Compustat `ppegt`). All emissions and age variables are within-month z-scored. *Capital age* is recursive capital age (winsorized at the 1st and 99th percentiles) (Lin, Palazzo, and Yang 2020; Geelen, Hajda, and Morellec 2024). Columns (1)–(3) include year-month and SIC-2 fixed effects; column (4) drops industry fixed effects; column (5) drops all firm-level controls; column (6) replaces separate year-month and industry fixed effects with SIC-2 × year-month fixed effects. B&K controls: ln(market cap), book-to-market, momentum, ROE, leverage, investment/assets, ln(PPE), sales growth, EPS growth, HHI, market beta, and return volatility (Bolton and Kacperczyk, 2021). Annual accounting variables are mapped to monthly returns using the Fama–French June t+1 convention; monthly controls are lagged one month. Sample: capital-intensive industries (SIC-2 industries above the 25th percentile of gross PP&E / total assets, measured 1985–1995), January 2016 to December 2024. Two-way clustered standard errors (firm × calendar year) in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 11: The Stranding Premium Holds Across Emissions Scopes

	Scope 1	Scope 2	Scope 3	Scope 1+2
Emissions	0.0042**	0.0051***	0.0039	0.0056***
	(0.0014)	(0.0016)	(0.0023)	(0.0015)
Capital age	0.0019	0.0017	0.0016	0.0020
	(0.0012)	(0.0012)	(0.0011)	(0.0012)
Emissions × Capital age	−0.0020***	−0.0022***	−0.0022***	−0.0022***
	(0.0006)	(0.0005)	(0.0005)	(0.0006)
Obs	114,010	113,986	114,010	114,010
B&K controls		Yes		
Year-month FE		Yes		
SIC-2 FE		Yes		

Notes: Each column reports the coefficient from a pooled OLS regression of monthly excess returns on emissions, capital age, their interaction, and the full set of B&K controls, with year-month and SIC-2 fixed effects. Scope 1 = direct GHG emissions; Scope 2 = indirect emissions from purchased electricity (location-based); Scope 3 = upstream supply-chain emissions; Scope 1+2 = sum of Scope 1 and Scope 2. All emissions and age variables are within-month z-scored. Sample: capital-intensive industries (SIC-2 industries above the balanced cutoff of gross PP&E / total assets, measured 1985–1995), January 2016 to December 2024. Two-way clustered standard errors (firm × calendar year) in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 12: Robustness to Finer Industry Fixed Effects (SIC-3 × Year-Month)

	(1) SIC-2 × YM	(2) SIC-3 × YM
Panel A: Emissions = ln(Scope 1 GHG)		
Emissions	0.0036**	0.0036**
	(0.0014)	(0.0014)
Capital age	0.0018	0.0017
	(0.0011)	(0.0014)
Emissions × Capital age	−0.0018**	−0.0019**
	(0.0006)	(0.0006)
Obs	114,010	114,010
Panel B: Emissions = ln(Scope 1 GHG/Gross PP&E)		
Emissions	0.0019**	0.0018*
	(0.0009)	(0.0009)
Capital age	0.0011	0.0010
	(0.0009)	(0.0012)
Emissions × Capital age	−0.0012*	−0.0012*
	(0.0006)	(0.0007)
Obs	114,010	114,010
B&K controls	Yes	

Notes: Robustness of Table 3 to finer industry fixed effects. Column (1) uses SIC-2 × year-month fixed effects (equivalent to column (6) of Table 3); column (2) uses SIC-3 × year-month fixed effects. Sample: capital-intensive industries, post-Paris (January 2016 to December 2024). Two-way clustered standard errors (firm × calendar year) in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 13: Older Dirty Capital Earns Lower Returns (Fama–MacBeth)

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Emissions = ln(Scope 1 GHG)						
Emissions	0.0023*	0.0021	0.0028*	0.0036**	0.0075***	0.0026*
	(0.0014)	(0.0013)	(0.0014)	(0.0014)	(0.0022)	(0.0014)
Capital age		0.0012	0.0020**	0.0020*	0.0014	0.0019
		(0.0007)	(0.0010)	(0.0011)	(0.0011)	(0.0013)
Emissions × Capital age			−0.0016**	−0.0015**	−0.0017**	−0.0017**
			(0.0007)	(0.0006)	(0.0007)	(0.0008)
Obs	114,010	114,010	114,010	114,010	141,453	114,010
Panel B: Emissions = ln(Scope 1 GHG/Gross PP&E)						
Emissions	0.0013*	0.0012*	0.0013	0.0018**	0.0031**	0.0011
	(0.0008)	(0.0008)	(0.0008)	(0.0007)	(0.0012)	(0.0008)
Capital age		0.0012	0.0014*	0.0012	0.0030**	0.0012
		(0.0007)	(0.0008)	(0.0010)	(0.0015)	(0.0011)
Emissions × Capital age			−0.0011*	−0.0009	−0.0010	−0.0010
			(0.0006)	(0.0006)	(0.0006)	(0.0007)
Obs	114,010	114,010	114,010	114,010	140,112	114,010
B&K controls	Yes	Yes	Yes	Yes	No	Yes
SIC-2 FE	Yes	Yes	Yes	No	Yes	—
SIC-3 FE	No	No	No	No	No	Yes

Notes: Fama–MacBeth (1973) version of Table 3. Each month's cross-section is a separate OLS regression; reported coefficients are time-series averages of monthly slopes. t-statistics use Newey–West standard errors with 12 lags. Panel A measures emissions as ln(Scope 1 GHG); Panel B as ln(Scope 1 GHG/gross PP&E) (Compustat `ppegt`). All emissions and age variables are within-month z-scored. *Capital age* is recursive capital age (winsorized at the 1st and 99th percentiles) (Lin, Palazzo, and Yang 2020; Geelen, Hajda, and Morellec 2024). Columns (1)–(3) include SIC-2 industry dummies; column (4) drops industry dummies; column (5) drops all firm-level controls; column (6) replaces SIC-2 with SIC-3 industry dummies. B&K controls: ln(market cap), book-to-market, momentum, ROE, leverage, investment/assets, ln(PPE), sales growth, EPS growth, HHI, market beta, and return volatility (Bolton and Kacperczyk, 2021). Annual accounting variables are mapped to monthly returns using the Fama–French June t+1 convention; monthly controls are lagged one month. Sample: capital-intensive industries (SIC-2 industries above the balanced cutoff of gross PP&E / total assets, measured 1985–1995), January 2016 to December 2024. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 14: Older Dirty Capital Earns Lower Returns (Fama–MacBeth, All Industries)

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Emissions = ln(Scope 1 GHG)						
Emissions	0.0007	0.0006	0.0007	0.0009	0.0063***	0.0008
	(0.0011)	(0.0011)	(0.0011)	(0.0010)	(0.0019)	(0.0010)
Capital age		0.0006	0.0008	0.0006	0.0006	0.0008
		(0.0008)	(0.0009)	(0.0010)	(0.0008)	(0.0010)
Emissions × Capital age			−0.0006	−0.0004	−0.0007*	−0.0006
			(0.0005)	(0.0005)	(0.0004)	(0.0005)
Obs	204,137	204,137	204,137	204,137	282,105	204,137
Panel B: Emissions = ln(Scope 1 GHG/Gross PP&E)						
Emissions	0.0004	0.0004	0.0003	0.0005	0.0010	0.0003
	(0.0006)	(0.0006)	(0.0006)	(0.0005)	(0.0008)	(0.0006)
Capital age		0.0006	0.0006	0.0004	0.0022	0.0007
		(0.0008)	(0.0008)	(0.0010)	(0.0014)	(0.0009)
Emissions × Capital age			−0.0003	−0.0000	0.0000	−0.0003
			(0.0004)	(0.0004)	(0.0004)	(0.0005)
Obs	204,137	204,137	204,137	204,137	249,697	204,137
B&K controls	Yes	Yes	Yes	Yes	No	Yes
SIC-2 FE	Yes	Yes	Yes	No	Yes	—
SIC-3 FE	No	No	No	No	No	Yes

Notes: All-industries Fama–MacBeth version of Table 3. Each month's cross-section is a separate OLS regression; reported coefficients are time-series averages of monthly slopes. t-statistics use Newey–West standard errors with 12 lags. Panel A measures emissions as ln(Scope 1 GHG); Panel B as ln(Scope 1 GHG/gross PP&E) (Compustat `ppegt`). All emissions and age variables are within-month z-scored. *Capital age* is recursive capital age (winsorized at the 1st and 99th percentiles) (Lin, Palazzo, and Yang 2020; Geelen, Hajda, and Morellec 2024). Columns (1)–(3) include SIC-2 industry dummies; column (4) drops industry dummies; column (5) drops all firm-level controls; column (6) replaces SIC-2 with SIC-3 industry dummies. B&K controls: ln(market cap), book-to-market, momentum, ROE, leverage, investment/assets, ln(PPE), sales growth, EPS growth, HHI, market beta, and return volatility (Bolton and Kacperczyk, 2021). Sample: all industries, January 2016 to December 2024. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 15: The Stranding Premium Concentrates in Capital-Intensive Industries (Fama–MacBeth)

	All industries		Capital-intensive		Capital-light	
	Full sample	Post-Paris	Full sample	Post-Paris	Full sample	Post-Paris
Emissions	0.0002	0.0006	0.0005	0.0016	−0.0006	−0.0011
	(0.0008)	(0.0014)	(0.0010)	(0.0016)	(0.0009)	(0.0015)
Capital age	0.0002	0.0005	0.0007	0.0012	−0.0004	0.0004
	(0.0005)	(0.0007)	(0.0006)	(0.0009)	(0.0008)	(0.0010)
Emissions × Capital age	−0.0003	−0.0004	−0.0008*	−0.0014**	−0.0002	0.0003
	(0.0003)	(0.0004)	(0.0005)	(0.0007)	(0.0004)	(0.0006)
Obs	321,678	231,553	166,654	114,900	155,024	116,653
Months	234	120	234	120	234	120
Controls	Yes					
SIC-2 FE	Yes					

Notes: Fama–MacBeth (1973) version of Table 2. Each cell reports the time-series average of monthly cross-sectional regression slopes. t-statistics use Newey–West standard errors with 12 lags. Emissions are measured as ln(Scope 1 GHG emissions). All emissions and age variables are within-month z-scored. *Capital age* is recursive capital age (winsorized at the 1st and 99th percentiles) (Lin, Palazzo, and Yang 2020; Geelen, Hajda, and Morellec 2024). Each cross-sectional regression includes SIC-2 industry dummies. Controls: ln(market cap), book-to-market, momentum, ROE, leverage, investment/assets, sales growth, EPS growth, HHI, market beta, and return volatility. Industries are split into "capital-intensive" and "capital-light" at the median of SIC-2 gross PP&E / total assets (Compustat, 1985–1995), chosen to balance the number of firm-month observations in each group. "Post-Paris" covers January 2016 to December 2024. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 16: Aggregate quadruple interaction: Emissions × Age × Capital-Intensity × Post-Paris

Emissions (ϕ)	0.0002
	(0.0010)
Capital age (A)	−0.0005
	(0.0011)
$\phi \times A$	0.0004
	(0.0007)
$\phi \times$ CapInt	−0.0034*
	(0.0017)
$\phi \times$ Post	0.0004
	(0.0016)
$\phi \times A \times$ CapInt	0.0007
	(0.0009)
$\phi \times A \times$ Post	0.0009
	(0.0011)
$\phi \times$ CapInt $\times$ Post	0.0067**
	(0.0026)
$\phi \times A \times$ CapInt $\times$ Post	−0.0040***
	(0.0013)
Observations	289,310
Controls	B&K full set
Fixed effects	Year-month × SIC-2

Notes: Pooled OLS of monthly excess returns on ln(Scope 1), recursive capital age, a capital-intensity indicator, and a post-Paris indicator (January 2016 onward), with all lower-order interaction terms. The B&K control set includes ln(market cap), book-to-market, momentum, ROE, leverage, investment-to-assets, ln(PPE), sales growth, EPS growth, HHI, beta, and volatility. Standard errors (in parentheses) are two-way clustered by firm and calendar year. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 17: Carbon premium by capital-age quintile: quadruple interaction [ln(Scope 1)]

	Capital-intensive		Capital-light	
Age quintile	Post-Paris	Pre-Paris	Post-Paris	Pre-Paris
Q1 (youngest)	69.4***	−72.4**	−12.0	14.7
	(26.6)	(32.0)	(24.9)	(17.5)
Q2	65.8***	−46.9	12.1	−3.9
	(16.3)	(30.5)	(16.5)	(20.3)
Q3	31.4	−28.7	−14.7	−4.7
	(20.3)	(19.0)	(15.3)	(13.4)
Q4	26.2	−16.6	14.1	−8.4
	(17.5)	(14.5)	(9.7)	(15.8)
Q5 (oldest)	1.0	−17.1	21.4*	23.4
	(13.1)	(20.3)	(11.4)	(20.0)
Observations	289,310			
Controls	B&K full set			
Fixed effects	Year-month × SIC-2			

Notes: Each cell reports the coefficient on Emissions × Q_k (in basis points per month per 1-SD emissions) from a single pooled regression with fully saturated quintile × emissions × capital-intensity × post-Paris interactions. Capital-age quintiles are formed within each calendar month. Standard errors (in parentheses) are two-way clustered by firm and calendar year. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 18: Pooled Daily OLS Event Study — All Industries

	(1) Raw	(2) Mkt-adj	(3) FF3	(4) FF4	(5) FF5
EventDay	−79.9137	−23.1423	22.5438	16.9917	–
	(76.0859)	(46.3971)	(15.5804)	(12.3301)	
Emissions (z)	0.0111	0.0112	−0.4980	−0.4972	−0.5108
	(0.4646)	(0.4673)	(0.3934)	(0.3892)	(0.3537)
Rec. capital age (z)	0.0990	0.1010	0.3640*	0.3732*	0.3353*
	(0.4260)	(0.4286)	(0.2196)	(0.2168)	(0.2028)
Em × RecAge	0.1109	0.1101	0.1223	0.1213	0.1473
	(0.2265)	(0.2296)	(0.1713)	(0.1707)	(0.1468)
Em × EventDay	8.4749	8.5526	13.8455	12.3702	9.7022
	(17.8586)	(17.9223)	(14.6023)	(13.7286)	(14.9577)
RecAge × EventDay	−24.5667	−24.3181	−11.0930	−10.8554	−13.1424
	(14.9739)	(14.9452)	(12.0774)	(11.9650)	(11.6350)
Em × RecAge × EventDay	21.3264**	19.1752**	10.0304**	11.2554**	10.2450***
	(9.8805)	(8.9787)	(4.1342)	(4.7492)	(3.8300)
Fixed effects	YM	YM	YM	YM	Date
Clustering	Firm+Date	Firm+Date	Firm+Date	Firm+Date	Firm+Date
Observations	8,425,156	8,425,156	8,424,298	8,424,298	8,424,298
Firms	4,251	4,251	4,226	4,226	4,226
Event firm-days	14,778	14,778	14,778	14,778	14,778

Table 19: Pooled Daily OLS Event Study — Capital-Light Industries (Placebo)

	(1) Raw	(2) Mkt-adj	(3) FF3	(4) FF4	(5) FF5
EventDay	−106.3936	−45.4141	7.8234	1.2223	–
	(93.1282)	(58.8591)	(8.0938)	(8.6468)	
Emissions (z)	0.2182	0.2163	−0.3009	−0.2939	−0.3184
	(0.5315)	(0.5375)	(0.4584)	(0.4528)	(0.4301)
Rec. capital age (z)	0.0855	0.0859	0.1736	0.1861	0.1681
	(0.5607)	(0.5683)	(0.2625)	(0.2589)	(0.2401)
Em × RecAge	0.3863	0.3846	0.2501	0.2533	0.2846
	(0.2458)	(0.2615)	(0.1789)	(0.1780)	(0.1756)
Em × EventDay	34.0216	37.9047	33.8794	29.7637	29.5887
	(25.1907)	(26.9240)	(22.2296)	(20.1052)	(21.6995)
RecAge × EventDay	−37.1116*	−34.0109*	−12.1576	−13.4451	−15.0831
	(21.9469)	(20.4768)	(12.4798)	(13.1162)	(12.3761)
Em × RecAge × EventDay	12.9398	10.7760	8.8653	9.2918	6.3862
	(9.8745)	(9.0894)	(5.4476)	(5.6546)	(4.7609)
Fixed effects	YM	YM	YM	YM	Date
Clustering	Firm+Date	Firm+Date	Firm+Date	Firm+Date	Firm+Date
Observations	3,996,249	3,996,249	3,995,839	3,995,839	3,995,839
Firms	2,142	2,142	2,130	2,130	2,130
Event firm-days	7,257	7,257	7,257	7,257	7,257

Table 20: Pooled Daily OLS Event Study — Subsidy Events, Capital-Intensive (Falsification)

	(1) Raw	(2) Mkt-adj	(3) FF3	(4) FF4	(5) FF5
EventDay	−30.3078	11.4471	1.4780	2.9416	–
	(34.9423)	(25.6335)	(15.1131)	(16.5898)	
Emissions (z)	0.1242	0.1262	−0.6796	−0.6807	−0.6887
	(0.7908)	(0.7941)	(0.5375)	(0.5359)	(0.4679)
Rec. capital age (z)	0.3552	0.3582	0.6505**	0.6612**	0.6093**
	(0.4539)	(0.4591)	(0.2801)	(0.2763)	(0.2553)
Em × RecAge	−0.0194	−0.0213	−0.0196	−0.0206	0.0112
	(0.3520)	(0.3556)	(0.2495)	(0.2487)	(0.2177)
Em × EventDay	−1.1042	−3.0526	0.1861	0.2411	−18.2645*
	(22.0015)	(22.6217)	(8.6279)	(8.9609)	(9.7044)
RecAge × EventDay	6.1173	3.1134	0.6294	0.2378	−4.8860
	(7.1169)	(7.4781)	(7.5798)	(7.6804)	(6.5678)
Em × RecAge × EventDay	5.4415	8.1007	9.2419**	9.1517**	14.0007***
	(8.3664)	(8.3803)	(4.5906)	(4.5582)	(3.5092)
Fixed effects	YM	YM	YM	YM	Date
Clustering	Firm+Date	Firm+Date	Firm+Date	Firm+Date	Firm+Date
Observations	4,428,907	4,428,907	4,428,459	4,428,459	4,428,459
Firms	2,236	2,236	2,223	2,223	2,223
Event firm-days	5,128	5,128	5,127	5,127	5,127

Table 21: Shift-share decomposition of aggregate fossil fuel intensity, 1998–2018

	F/L	F/Y
Aggregate 1998	0.865	6.347
Aggregate 2018	1.030	3.865
Total change	+0.165 (+100.0%)	−2.482 (+100.0%)
Within	+0.044 (+26.4%)	−2.817 (+113.5%)
Between	+0.121 (+73.6%)	+0.335 (−13.5%)
Cross	+0.002 (+1.1%)	−1.245 (+50.1%)

Notes: Symmetric (average-weight) decomposition across 21 NAICS-3 manufacturing subsectors. F = fossil fuel consumption (EIA MECS, Trillion BTU); L = employment (Census CBP); Y = real value added (NBER-CES, double-deflated, \$B). Within = $\sum_i \bar{s}_i \, \Delta(F/X)_i$; Between = $\sum_i \Delta s_i \, \overline{(F/X)}_i$; Cross = $\sum_i \Delta s_i \, \Delta(F/X)_i$.

Online Appendix

OA.1 COVID Robustness for Paris-Timing Sensitivity

Figure OA.1 repeats the Paris-timing sensitivity analysis of Figure 4 with COVID-period controls. A dummy for March 2020 through December 2021 is fully interacted with emissions, age, and capital intensity, yielding 8 additional controls that absorb all COVID-period variation in the key variables (the COVID dummy plus all two-way and three-way interactions with ϕ, $\bar{a}$, and CapInt). The coefficient paths are virtually unchanged relative to the uncontrolled version, confirming that the 2021–2022 strengthening in both the carbon premium and its age gradient reflects policy repricing, most likely the Biden administration's renewed regulatory agenda, rather than pandemic-related crash/recovery mechanics.

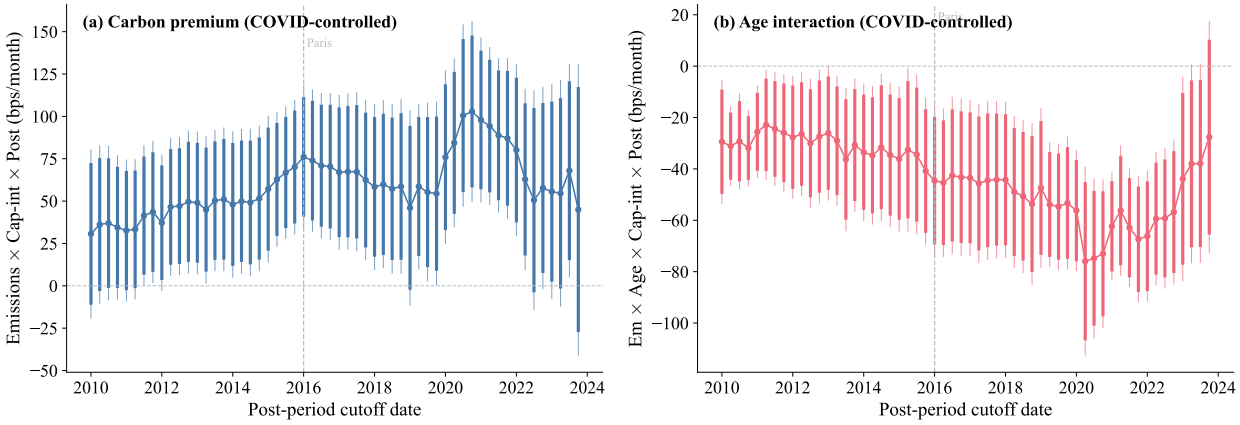


Figure OA.1: Paris-timing sensitivity with COVID controls.

Notes: Same specification as Figure 4, with the addition of a COVID-period dummy (March 2020–December 2021) fully interacted with emissions, age, and capital intensity. Panel (a): three-way interaction $\phi \times \text{CapInt} \times \text{Post}$ (incremental carbon premium). Panel (b): four-way interaction $\phi \times \bar{a} \times \text{CapInt} \times \text{Post}$ (incremental age gradient). Thick error bars: 90% CI; thin error bars: 95% CI (two-way clustered by firm and calendar year). The coefficient paths are virtually unchanged relative to the uncontrolled version in Figure 4.

OA.2 Factor-Model-Adjusted CAR Regressions

Tables OA.1 and OA.2 replicate the cross-sectional CAR regressions of Table 1 using Fama-French three- and five-factor-adjusted abnormal returns. Firm-level factor betas are estimated from the $[-250, -30]$ pre-event window. The fossil vintage coefficient strengthens under both factor adjustments, consistent with the age effect reflecting stranding-risk exposure rather than mechanical factor loading.

Table OA.3 drops the quadratic term and estimates a purely linear age effect. The fossil vintage coefficient remains positive, confirming that the result does not depend on the quadratic specification.

Table OA.1: Cross-sectional CAR regressions: FF3-adjusted

	Panel A: Dirty generator age			Panel B: Solar+Wind generator age		
	$(0, 5)$	$(0, 10)$	$(0, 15)$	$(0, 5)$	$(0, 10)$	$(0, 15)$
Generator age	$+0.0341^{***}$	$+0.0332^{**}$	$+0.0272$	$+0.0858$	$+0.0062$	$+0.0245$
	(2.62)	(2.00)	(0.87)	(1.00)	(0.05)	(0.16)
Generator age^2	$+0.0008$	$+0.0012$	$+0.0016$	-0.0350	-0.0178	-0.0332
	(0.74)	(0.80)	(0.73)	(-1.27)	(-0.46)	(-0.74)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Event FEs	Yes	Yes	Yes	Yes	Yes	Yes
N	744	744	744	101	101	101
R^2	0.546	0.471	0.381	0.779	0.783	0.761

Cross-sectional regressions of signed Fama-French three-factor-adjusted CAR (in percentage points) on orthogonalized generator age. 19 cost-side events. Firm-level factor betas estimated from $[-250, -30]$ pre-event window. Panel A: dirty generator age. Panel B: solar+wind generator age (firms with solar+wind capacity share $> 10\%$). Controls, event fixed effects, and event-clustered t-statistics as in Table 1. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table OA.2: Cross-sectional CAR regressions: FF5-adjusted

	Panel A: Dirty generator age			Panel B: Solar+Wind generator age		
	$(0, 5)$	$(0, 10)$	$(0, 15)$	$(0, 5)$	$(0, 10)$	$(0, 15)$
Generator age	$+0.0375^{**}$	$+0.0330^{*}$	$+0.0287$	$+0.0627$	-0.0042	$+0.0115$
	(2.56)	(1.86)	(0.85)	(0.69)	(-0.03)	(0.08)
Generator age^2	$+0.0011$	$+0.0011$	$+0.0017$	-0.0329	-0.0197	-0.0313
	(0.94)	(0.75)	(0.73)	(-1.12)	(-0.48)	(-0.66)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Event FEs	Yes	Yes	Yes	Yes	Yes	Yes
N	744	744	744	101	101	101
R^2	0.563	0.468	0.353	0.812	0.782	0.734

Cross-sectional regressions of signed Fama-French five-factor-adjusted CAR (in percentage points) on orthogonalized generator age. 19 cost-side events. Firm-level factor betas estimated from $[-250, -30]$ pre-event window. Panel A: dirty generator age. Panel B: solar+wind generator age (firms with solar+wind capacity share $> 10\%$). Controls, event fixed effects, and event-clustered t-statistics as in Table 1. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table OA.3: Cross-sectional CAR regressions: generator age (linear)

	Panel A: Dirty generator age			Panel B: Solar+Wind generator age		
	$(0, 5)$	$(0, 10)$	$(0, 15)$	$(0, 5)$	$(0, 10)$	$(0, 15)$
Generator age	+0.0224	+0.0174	+0.0075	−0.0548	−0.1158	−0.1352
	(1.44)	(1.02)	(0.30)	(−0.64)	(−1.06)	(−0.95)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Event FEs	Yes	Yes	Yes	Yes	Yes	Yes
N	744	744	744	101	101	101
R^2	0.113	0.110	0.110	0.376	0.467	0.441

Cross-sectional regressions of signed industry-adjusted CAR (in percentage points) on orthogonalized generator age. 19 cost-side events. Panel A: dirty generator age = capacity-weighted fossil fleet age, orthogonalized on total fleet age (per-event cross-section). Panel B: solar+wind generator age = capacity-weighted solar and wind fleet age, orthogonalized on total fleet age; restricted to firms with solar+wind capacity share > 10%. Controls: fossil generation share, log generation capacity, log market capitalization, book-to-market ratio, and leverage. All specifications estimated by OLS with event fixed effects. t-statistics in parentheses are event-clustered. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.