Capital Vintage and the Carbon Premium

Simon Gilchrist* Joseba Martinez† Natalie Rickard‡§

June 2026

Abstract

We argue theoretically and show empirically that the carbon premium prices exposure to stranded fossil-fuel capital. In a vintage-capital model in which installed machines are locked into fuel requirements, policies that raise fossil-fuel operating costs lower firm value in proportion to expected future fuel use. The model predicts that policy-risk premia should be largest for higher-emissions firms with young physical capital because their fuel commitments have the longest remaining lives; by contrast, the premium should be weaker for firms with older capital and in sectors where physical capital is a less important factor of production. We document three facts that are consistent with this prediction. First, in the U.S. electricity sector, where generator-level data provide a direct measure of capital age, firms with younger fossil fleets underperform those with older fleets in response to climate-policy events that raise expected operating costs for fossil assets. Second, in the broad cross-section of U.S. public firms, a consistently positive carbon risk premium appears after the 2016 Paris Agreement and is concentrated in capital-intensive industries. Within these industries, the emissions premium is largest for firms with the youngest capital and indistinguishable from zero in the oldest capital-age quintile. Third, rolling estimates show that the young-capital premium is negative before the Paris Agreement, turns positive afterward, and rises when climate-policy attention is high. Finally, we form a portfolio that carries stranding risk by isolating young-minus-old exposure among high-emission firms relative to clean firms within capital-intensive industries. In a historical portfolio exercise, signal-timed exposure that goes long stranding risk when climate-policy salience is high and short when it is low delivers a post-Paris Agreement annualized Sharpe ratio of 1.01 and significant factor-adjusted alpha.

*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

§We thank Patrick Bolton, Victor DeMiguel, Roberto Gómez-Cram, Marcin Kacperczyk, and Dimitris Papanikolaou for feedback and suggestions. 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

A central question in climate finance is whether investors in high-emissions firms earn a premium for bearing transition risk. Influential studies (Bolton and Kacperczyk, 2021, 2023) provide evidence that carbon emissions carry a positive average premium in the cross-section of equity returns. We show that this premium varies systematically with the age of installed physical capital and the state of climate-policy risk: it is concentrated among high-emissions firms with young physical capital, especially after the Paris Agreement and during periods of high climate-policy salience. By contrast, we find no comparable premium among high-emissions firms with old capital or among firms whose industries rely less heavily on physical capital. These patterns point to a stranding-risk channel in which emissions are priced when they are embedded in long-lived physical assets with substantial remaining policy exposure.

The paper is organized around a putty-clay model of installed capital (Johansen, 1959; Gilchrist and Williams, 2000). Three features of the model are central to the analysis: capital is designed with fixed fuel-to-labor ratios at the time of installation, newer vintages embody more productive technology, and firms have the option to idle capital when operating conditions are unfavorable. A carbon risk premium arises because policies that increase the cost of fossil-fuel use, such as carbon taxes, cap-and-trade systems, and emissions standards, generate valuation losses that scale with expected future fuel use. Whether the premium is larger for young or old capital depends on a tradeoff. Young capital has longer remaining fixed fuel commitments and therefore a larger dollar exposure to a fuel-cost shock. Older capital has lower productivity, which translates into thinner profit margins and a lower value, raising the proportional loss for any given dollar exposure. The theory shows that the proportional loss is larger for young capital when the option value of operating old machines is sufficiently high.

We examine the mechanism in U.S.-listed electric utilities. Utilities are a natural testing ground because generator-level data let us separately measure the capacity-weighted age of each public utility's fossil and renewable fleets. We calibrate the model using hourly operating data for U.S. electricity generators matched to plant-level cost data, and find that the empirically relevant parameter region is the one in which the carbon premium is largest for firms with young dirty capital. We then study portfolio returns around climate-policy events that raise expected operating costs for fossil generation. In the post-event window, firms with younger fossil fleets earn significantly lower cumulative returns than firms with older fleets. The same exercise using renewable-fleet age shows no such pattern. The contrast is informative because, in the model, capital vintage operates as a measure of remaining fuel exposure, not as a generic firm characteristic.

We then ask whether the same mechanism appears in the broad cross-section of U.S. public firms. As emphasized by Kacperczyk (2026), asset-pricing regressions should use regressors that proxy for priced cash-flow exposure to transition risk, and this exposure is typically proxied by current emissions. In our framework, current emissions alone do not capture remaining fuel commitments, but the interaction between current emissions and capital age does. This motivates a return regression in which the emissions slope varies with capital age and capital intensity.

The broad-market evidence is consistent with the model's cross-sectional predictions. Using Trucost Scope 1 emissions and a balance-sheet-based proxy for capital age, we find that the carbon premium is concentrated in capital-intensive industries and in time periods after the Paris Agreement. Within these industries, the carbon premium is largest in the youngest capital-age quintile and indistinguishable from zero in the oldest quintile. In capital-light industries, where the putty-clay mechanism should have little bite, the premium is insignificant at every age, in the full sample, and in the post-Paris Agreement subsample.

Time variation provides further evidence for interpreting the carbon premium as a policy-induced asset-stranding premium. Rolling estimates show that the carbon premium in the youngest capital-quintile among capital-intensive firms is negative on average before the Paris Agreement and positive on average after, with a statistical break dated March 2016, within months of the adoption of the Agreement in December 2015. The premium for the oldest quintile remains close to zero in both regimes.

The same structure appears in response to climate-policy salience. When text-based measures of climate-policy attention (Engle et al., 2020; Ardia et al., 2023; Arteaga-Garavito et al., 2023; DeMiguel et al., 2026) are low, the carbon premium is indistinguishable from zero in every capital-age quintile. When salience is elevated, the youngest-quintile premium is large and positive, while the oldest-quintile premium remains indistinguishable from zero. This pattern is consistent with survey evidence that market participants view regulatory risk as the leading near-term climate risk (Stroebel and Wurgler, 2021). The carbon premium appears only when investors perceive policy risk as salient, and only among firms with the longest remaining fossil exposure.

Our cash-flow channel is distinct from preference-based mechanisms (Pástor et al., 2021; Pedersen et al., 2021), which can generate a brown-minus-green spread but do not obviously predict its concentration in firms with high young-capital intensity. In our mechanism, the age-emissions interaction emerges when the policy risk is elevated, and the premium fades when the risk is low. The negative premiums that we observe in low-salience periods are consistent with mechanisms based on richer SDF structure (Gârleanu and Pedersen, 2025) or on cash-flow hedging of climate-damage states (Kacperczyk, 2026).

The final part of the paper studies portfolio implications. Recent work shows that realized brown-minus-green portfolio returns are fragile, with small unconditional Sharpe ratios (Eskildsen et al., 2026). We find the same in our sample. A standard brown-minus-green portfolio has a modest Sharpe ratio and insignificant factor-adjusted alphas. The model suggests why. If the carbon premium is concentrated in young dirty capital and appears only when policy risk is salient, then unconditional brown-minus-green spreads average over many firms and months in which the relevant risk is weakly priced. We therefore form a stranding-risk portfolio that isolates young-minus-old exposure among dirty firms relative to clean firms within capital-intensive industries, and condition exposure to this portfolio on a measure of climate-policy salience. This signal-timed strategy has a post-Paris annualized Sharpe ratio of 1.01 and a factor-adjusted alpha of 3.7% per year.

The cross-sectional structure we document has implications beyond asset pricing. That transition risk is priced into installed fossil capital, and concentrated in the youngest and most fuel-committed vintages, implies that the pricing of climate policy operates through assets whose composition cannot be quickly rebalanced. This connects naturally to macroeconomic models in which capital is costly to reallocate and the composition of investment matters for aggregate adjustment, as in our related paper Gilchrist et al. (2026), as well as Lanteri and Rampini (2022); Bertolotti et al. (2025). Unlike in frictionless economies, where return premia replicate the incentives of a carbon tax on emissions (Chittaro et al., 2025), our findings suggest that the carbon premium disproportionately taxes new capital relative to old capital, and in this sense resembles a tax on marginal dirty investment rather than a tax on fossil fuel use. This has implications for innovation. A forward-looking firm contemplating new dirty investment faces the same prospective stranding tax, raising the cost of capital and reinforcing the price signals emphasized in models of directed technical change (Acemoglu et al., 2012; Aghion et al., 2016). More broadly, Acharya et al. (2025) shows how transition risk reallocates value across firms in equilibrium, and our evidence identifies the cross-sectional shape of that reallocation.

2 Model

We build a putty-clay model with variable capital utilization. The distinguishing feature of putty-clay technology is that firms choose factor proportions when capital is installed, and those factor proportions cannot be adjusted freely afterward. Putty-clay models have been used to study how technological transitions diffuse gradually through embodied capital (Atkeson and Kehoe, 2007), how irreversible capital shapes asset-price volatility (Gourio, 2011), how embodied technology affects labor demand and factor shares (Martinez, 2026), and the welfare effects of climate policy and other forces driving the energy transition in general equilibrium (Wei, 2003; Gilchrist et al., 2026). Here, we study how capital fixity determines which firms bear priced exposure to policy-induced stranding. In our model, a carbon tax strands dirty capital because it loads on locked-in fuel commitments, but the size of that stranding loss depends on how fuel-intensive the machine is and how much future fuel utilization remains relative to future value.

2.1 Technology

A continuum of intermediate-good firms is indexed by dirtiness $\phi \in [0, 1]$. Each firm operates a fleet of machines. An age-a machine at date t in firm ϕ is designed with fixed capital-to-labor ratio $k(a, \phi)$ and fuel-to-labor ratio $f(a, \phi)$. If operated, it uses a fixed amount of labor, which we normalize to one unit for convenience. Machine-level output conditional on operation is Cobb–Douglas:

$$y_t(a, \phi) = \theta_{t-a}\, Z_t\, k(a, \phi)^{\alpha_k(\phi)}\, f(a, \phi)^{\alpha_f(\phi)}, \tag{1}$$

with $\alpha_k(\phi) + \alpha_f(\phi) + \alpha_\ell = 1$. We assume $\alpha'_f(\phi) > 0$, so greater dirtiness monotonically translates to greater fuel intensity. We further assume that the labor share α_ℓ is common to all firms, which implies that the capital share $\alpha_k(\phi) = 1 - \alpha_\ell - \alpha_f(\phi)$ is decreasing in ϕ. Assuming a common labor share also implies that firms have the same balanced-growth rate, which simplifies the analysis but is not important for our findings. Machine productivity has an embodied component θ_{t-a}, fixed when the machine is installed, and a disembodied component Z_t, common to all vintages. Machines die stochastically at Poisson rate $\delta > 0$.

Each period, the firm chooses which of its surviving machines to utilize. In levels, the deterministic payoff to firm ϕ from operating an age-a is

$$\pi_t(a, \phi; \tau_t) = P_t(\phi)\, y_t(a, \phi) - \left(p_t^f + \tau_t\right) f(a, \phi) - W_t, \tag{2}$$

where $P_t(\phi)$ is the output price of intermediate good ϕ, p_t^f is the fuel price, W_t is the wage, and τ_t is a carbon tax. Each machine draws iid Type-I extreme-value shocks $\varepsilon_{1,t}$ and $\varepsilon_{0,t}$ on the operate (1) and shut-down (0) options, each with scale parameter $\kappa_t > 0$ in every period. The machine therefore operates whenever

$$\pi_t(a, \phi; \tau_t) + \varepsilon_{1,t} \geq \varepsilon_{0,t}.$$

This formulation is a reduced-form way to capture variable utilization arising from transitory machine-level operating conditions. Let $\chi_t(a, \phi) \in \{0, 1\}$ denote the operate/shut-down decision. Under the extreme-value assumption, the difference $\varepsilon_1 - \varepsilon_0$ is logistic, which gives the conditional-logit choice probability (McFadden, 1974)

$$u_t(a, \phi; \tau_t) = \Pr\Big(\chi_t(a, \phi) = 1\Big) = \frac{1}{1 + e^{-\pi_t(a,\phi;\tau_t)/\kappa_t}}, \tag{3}$$

increasing in operating profitability. The corresponding expected per-period profit is

$$\Pi_t(a, \phi; \tau_t) = \kappa_t \log\Big(1 + e^{\pi_t(a,\phi;\tau_t)/\kappa_t}\Big). \tag{4}$$

Equation (4) is the expected per-period value of the machine when the operator chooses optimally between running and idling. It captures the option value of being able to idle in bad states. When operating conditions are more variable (a larger κ_t), the run/idle decision is less tightly tied to current operating profit. As κ_t shrinks, the operator approaches the deterministic rule of running the machine whenever its operating profit is positive.

Let $n_t(a, \phi)$ denote the mass of surviving age-a machines at firm ϕ. Firm-level output is

$$Y_t(\phi) = \int_0^\infty u_t(a, \phi)\, y_t(a, \phi)\, n_t(a, \phi)\, da. \tag{5}$$

We assume that firms produce differentiated intermediate goods and face an inverse-demand schedule $P_t(\phi)$. The exact demand system matters for pass-through and general-equilibrium quantities, but it does not enter the machine-level comparative statics below. Appendix B states the CES aggregation structure that rationalizes this price schedule and discusses the associated general-equilibrium channels.

We formulate the policy instrument as a per-unit carbon tax for tractability. The direct operating-cost channel that drives our results is common to any regulation that increases the cost of fuel use, such as emissions standards or cap-and-trade systems. Other transition policies, such as clean-energy subsidies or technology mandates, can also induce stranding, but they operate indirectly through output prices or entry conditions rather than through the operating cost of installed capital. We discuss how alternative policy instruments map into the framework in Appendix B.

2.2 Balanced growth path

On a balanced growth path, embodied and disembodied technical change jointly determine the growth rates of machines at the technological frontier (that is, machines created at any instant t that embody the latest technology θ_t) and of aggregate output. Let embodied productivity θ_t grow at rate g_θ, and let disembodied productivity Z_t grow at rate g_z. The economy and frontier machine output grow at the same rate

$$g_y = \frac{g_\theta + g_z}{\alpha_\ell}. \tag{6}$$

We use stars without time subscripts to denote stationary frontier objects, so $k_t(0,\phi) = k_t^*(\phi) = e^{g_y t} k^*(\phi)$, $f_t(0,\phi) = f_t^*(\phi) = e^{g_y t} f^*(\phi)$, and $y_t(0,\phi) = y_t^*(\phi) = e^{g_y t} y^*(\phi)$. The wage and the idiosyncratic shock scale grow with the same factor, $W_t = e^{g_y t} w$ and $\kappa_t = \kappa e^{g_y t}$, while the prices of output and fuel are stationary, $P_t(\phi) = P(\phi)$ and $p_t^f = p^f$. Balanced growth also requires that the carbon tax is constant, $\tau_t = \tau$.[1]

Detrending the operating payoff in (2) by $e^{g_y t}$ makes the utilization problem stationary. For an age-a machine, the fuel requirement was fixed at installation, so detrended fuel use decays at rate g_y. Output also benefits from the common disembodied component Z_t, so detrended revenue decays more slowly, at rate $g_y - g_z$. Define this net vintage decay rate as

$$\tilde{g} \equiv g_y - g_z. \tag{7}$$

The balanced-growth counterpart of the level payoff is therefore

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

The stationary utilization rule $u(a,\phi;\tau)$ and detrended expected profit $\Pi(a,\phi;\tau)$ are given by replacing $\pi_t = e^{g_y t}\pi$, $\kappa_t = e^{g_y t}\kappa$ into Equations (3) and (4). For convenience, in what follows we suppress the dependence of functions on τ to denote their value at $\tau = 0$.

At installation, firms choose machine design under free entry, taking future utilization into account. The resulting equilibrium pins down detrended frontier revenue $P(\phi)y^*(\phi)$, fuel commitment $f^*(\phi)$ and machine size $k^*(\phi)$. Appendix B derives these objects from the machine-design problem and shows the free-entry conditions. For the results below, the key point is that the equilibrium delivers smooth, bounded functions $P(\phi)y^*(\phi)$ and $f^*(\phi)$.

The detrended value of an age-a machine is given by the present value of the detrended cash flows the machine generates when it is utilized optimally:

$$v(a,\phi) = \int_0^\infty e^{-(r+\delta)h}\Pi(a+h,\phi)\, dh, \tag{9}$$

where r is the discount rate.

[1] We allow κ to scale with the economy to maintain balanced growth. An equivalent formulation would be to denominate the run/idle shocks in wage units.

2.3 Asset stranding

We now study how machine value responds to the introduction of an unexpected carbon tax. Starting from Equation (9), a marginal increase in the tax changes machine value through the operating component of the flow payoff. By the envelope theorem for the operating problem,

$$\left.\frac{\partial \Pi(a+h,\phi;\tau)}{\partial \tau}\right|_{\tau=0} = -u(a+h,\phi)\, f^*(\phi) e^{-g_y(a+h)},$$

because a one-unit increase in the carbon tax raises the fuel cost of an operating machine one-for-one with its fuel use, while utilization is already optimally chosen. Integrating this flow effect forward implies

$$\left.\frac{\partial v(a,\phi;\tau)}{\partial \tau}\right|_{\tau=0} = -f^*(\phi) \int_0^\infty e^{-(r+\delta)h} u(a+h,\phi) e^{-g_y(a+h)}\, dh = -f^*(\phi) I(a,\phi),$$

where

$$I(a,\phi) \equiv \int_0^\infty e^{-(r+\delta)h} u(a+h,\phi) e^{-g_y(a+h)}\, dh \tag{10}$$

is the discounted amount of future operating time over which the machine's fuel requirement is exposed to the tax.

We define the asset stranding loss from a marginal carbon tax increase as

$$\mathsf{s}(a,\phi) \equiv -\frac{1}{v(a,\phi)} \left.\frac{\partial v(a,\phi;\tau)}{\partial \tau}\right|_{\tau=0} = \frac{f^*(\phi)\, I(a,\phi)}{v(a,\phi)}, \tag{11}$$

that is, the present value of expected fuel exposure as a fraction of machine value. This is the machine-level object that we map into firm-level stranding and, in turn, into expected returns. This is a partial-equilibrium loss because we do not consider the effect of the carbon tax on output prices, wages and the interest rate. In what follows, we discuss how general equilibrium forces affect stranding losses.

Equation (11) expresses the stranding loss as the ratio of dollar exposure to machine value. The numerator $f^*(\phi)\, I(a,\phi)$ is the present value of fuel exposed to the tax. It rises with the per-period fuel commitment $f^*(\phi)$ and with the discounted future utilization, $I(a,\phi)$. The denominator $v(a,\phi)$ converts dollar exposure into a proportional loss. Both dirtiness and age move the numerator and denominator at the same time. A dirtier machine has a higher f^*, but lower flow profit pulls down utilization and value. A younger machine has more remaining utilization and so larger I, but also larger remaining cash flow and hence larger v. In each direction, whether stranding rises depends on which component of the ratio responds more.

The utilization parameter κ governs the relative responses of future fuel exposure and value. When κ is small, utilization responds sharply to per-period profit, so older and dirtier machines sit close to the shutdown margin, and the value denominator is highly sensitive to dirtiness and age. When κ is larger, the utilization decision is smoother, and machines retain greater option value as they age, so cross-sectional variation in stranding is determined to a greater extent by the numerator. Dirtier machines have larger f^*, and younger machines have more remaining fuel exposure. We derive the model's central comparative statics in this parameter region.

Proposition 1 (Stranding Comparative Statics)**.** *For any compact range of ϕ and all ages $a \geq 0$ there exists a finite threshold $\bar{\kappa}$ such that for every $\kappa \geq \bar{\kappa}$,*

$$\partial_\phi \mathsf{s}(a, \phi) > 0.$$

and

$$\partial_{a\phi} \mathsf{s}(a, \phi) < 0.$$

Proof. See Appendix A.1.

We assess whether this threshold condition is empirically relevant by calibrating the model using data from the U.S. fossil-fuel electricity generation sector, a setting in which the model's key primitives can be measured directly. Power plants are long-lived assets with fuel requirements largely fixed by technology and vintage, making them a close empirical counterpart to the putty-clay machines in the model. The sector also provides detailed data on utilization, vintage, and operating costs. In this calibration exercise, we find that the 95% confidence interval for the estimated κ sits above the model-implied $\bar{\kappa}$, placing the electricity-generation sector inside the parameter region of Proposition 1. Appendix C provides details on the data and the calibration methodology.

Figure 1 plots $\mathsf{s}(a, \phi)$ at the data-implied calibration for three fuel-intensity levels, $\phi_L < \phi_M < \phi_H$, and illustrates the partial derivatives signed by Proposition 1. At every age, stranding is highest for the dirtiest machines and lowest for the cleanest ($\partial_\phi \mathsf{s} > 0$). All three profiles decline with age as remaining fuel exposure decays, and the gap between profiles narrows with age ($\partial_{a\phi} \mathsf{s} < 0$). Stranding losses are therefore largest for young dirty capital.

2.4 Firm aggregation and expected returns

The preceding section characterized stranding at the machine level. To connect this to observable firm-level returns, we follow three steps: aggregate machine-level stranding to the firm level; show that this aggregate exposure is priced in expected returns; and map the model to the reduced-form regression that we take to the data.

Define the share of firm value attributable to vintage a as

$$\omega(a, \phi) \equiv \frac{v(a, \phi)\, n(a, \phi)}{\int_0^\infty v(a', \phi)\, n(a', \phi)\, da'},$$

where the denominator is the total value of the firm's capital, $V(\phi) \equiv \int_0^\infty v(a', \phi)\, n(a', \phi)\, da'$. The firm-level stranding ratio is then the value-weighted average of machine-level stranding:

$$\mathsf{S}(\phi) = \int_0^\infty \mathsf{s}(a, \phi)\, \omega(a, \phi)\, da = \frac{\int_0^\infty f^*(\phi) I(a, \phi) n(a, \phi)\, da}{V(\phi)}. \tag{12}$$

Equation (12) is the firm-level exposure that carries through the asset-pricing implications. In a neighborhood of the parameter region where Proposition 1 applies, machine values are

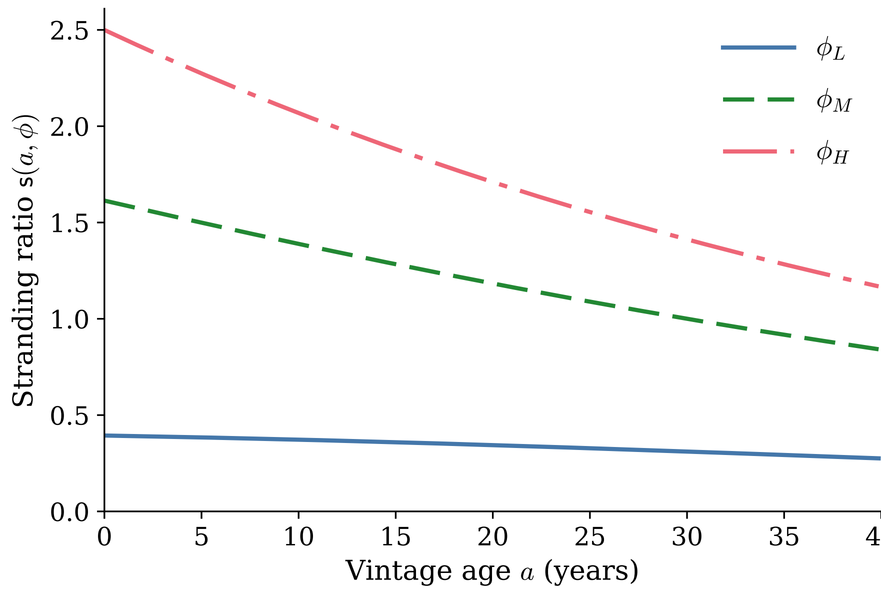

Figure 1: Stranding losses by vintage age and dirtiness

Notes: Machine-level stranding loss $\mathsf{s}(a, \phi)$ for three dirtiness levels $\phi_L < \phi_M < \phi_H$, corresponding to fuel shares $\alpha_f = 0.25$, 0.65, and 0.85 respectively. The calibration uses $\kappa = 1.13$, $Py^\star/w = 5.42$, $\alpha_\ell = 0.15$, $r = 0.05$, $\delta = 0.02$, and $g_y = 0.02$ (see Appendix C for details on the calibration).

approximately uniform across (a, ϕ) so the weights ω are determined by the distribution n, and the machine-level signs carry directly to the firm level. Therefore Proposition 1 implies that more fossil fuel-intensive firms have greater stranding exposure, and that the gradient is larger the younger the installed capital stock.

To connect stranding exposure to returns, suppose transition policy arrives as a priced jump risk. A permanent carbon tax of size τ arrives at Poisson rate $\lambda > 0$, and the stochastic discount factor jumps by a factor of $1 + \xi$ when the tax arrives.

Proposition 2 (Stranding risk premia)**.** *Under the jump-risk environment above, expected excess returns satisfy, to first order in* τ*,*

$$E[r^e(\phi)] \approx \lambda\, \xi\, \tau\, \mathsf{S}(\phi). \tag{13}$$

Proof. See Appendix A.

Equation (13) combines exposure and the price of risk. The exposure term $\mathsf{S}(\phi)$ is determined by installed capital. The sign of the expected-return premium depends on ξ, the price of the policy-arrival state. When $\xi > 0$, policy arrives in high-marginal-utility states, so high-S firms must offer higher expected returns before the shock. Combining this pricing condition with Proposition 1 gives the model's risk-premium prediction. Conditional on a positive price of transition risk, the carbon premium is largest for young dirty capital and declines with capital age.

The condition $\xi > 0$ holds when climate policy is contractionary on impact. Gilchrist et al. (2026) show this in a general-equilibrium model with putty-clay technology, where a carbon tax produces a sharp drop in aggregate consumption because the existing dirty capital stock cannot be costlessly reallocated. **?** documents the same pattern empirically, finding that carbon-policy shocks generate a persistent decline in aggregate consumption. The stochastic discount factor is therefore high precisely when stranding losses are realized.

The same exposure has the opposite sign in announcement returns. If a policy event raises the expected present value of cost-side regulation by $\Delta q > 0$, then

$$\Delta \log V(\phi) \simeq -\Delta q \, \mathsf{S}(\phi).$$

Thus high-exposure firms earn higher expected returns in normal times, and lose more when news raises expected costs. Tightening events should therefore produce more negative returns for young dirty capital than for old dirty capital. We test this prediction using the event study in the next section.

A macro asset-pricing interpretation requires a household sector that prices firm cash flows through a stochastic discount factor and supplies labor in general equilibrium. We leave these features of the economy deliberately unspecified in the main text because the machine-level cross-sectional mechanism does not depend on a particular specification of preferences, labor supply, or market clearing. Partial equilibrium isolates the three objects that determine stranding across machine age and dirtiness, remaining fuel exposure, endogenous utilization, and machine value.

General-equilibrium forces matter for pass-through, wages, discount rates, and the sign of the risk premium. Appendix B shows how these forces enter the same machine-value decomposition once additional structure is imposed. In general, these forces strengthen the partial-equilibrium age gradient. This interpretation is consistent with the related dynamic general-equilibrium model with carbon damages in Gilchrist et al. (2026), where dirtier firms are more exposed to transition risk and the emissions-age interaction remains central for the cross-section of returns.[2]

We now translate these exposure predictions into the regression specification used in the cross-sectional excess return regressions. Equation (13) shows that expected returns load on $S(\phi)$, which is the value-weighted aggregate of machine-level stranding across the firm's installed capital. The numerator of (12) aggregates the dollar cash-flow exposure of installed machines to a marginal carbon tax; the denominator converts that dollar exposure into a proportional value loss.

[2]The object $\mathsf{S}(\phi)$ measures stranding of assets in place. As emphasized by Papanikolaou (2011) and Kogan and Papanikolaou (2014), firm value also includes future investment, which in general would be affected by policy. In the balanced-growth environment considered here, such growth-opportunity exposure scales the level of stranding but does not change the signs of the cross-sectional comparative statics. We therefore focus on the assets-in-place channel and control empirically for standard growth-opportunity proxies in cross-sectional regressions. The putty-clay mechanism is distinct from termination-risk models in which fossil incumbents optimally reduce investment as the expected end of their business approaches, raising scarcity rents and equity values (Engle, 2024). Here, the priced exposure comes from the value loss on installed capital whose fuel requirements are fixed at installation.

This structure is closely related to the measurement problem emphasized by Kacperczyk (2026). In a general asset-pricing setting, carbon-related return premia load on dollar cash-flow exposure normalized by firm value. Our model gives that exposure object as $f^*(\phi)I(a,\phi)$, with dirtiness entering through the installed fuel requirement and age entering through the discounted future utilization.

To show how observed emissions map to this object, we re-index machines by their installed fuel requirement $f_i = f^*(\phi)$. Current emissions from an age-a machine are proportional to current fuel use,

$$e(a, f_i) \propto f_i u(a, f_i) e^{-g_y a},$$

where the constant of proportionality is the carbon content of fuel. The dollar exposure of the same machine can therefore be written as current emissions times a discounted remaining-emissions multiple relative to current emissions:

$$f_i I(a, f_i) \propto e(a, f_i) U(a, f_i)$$

where $U(a, f_i) = \int_0^\infty e^{-(r+\delta+g_y)h} \frac{u(a+h, f_i)}{u(a, f_i)}\, dh$ is a utilization growth factor. Integrating over the density of firm-i machines and scaling by value gives

$$\mathsf{S}_i \propto \frac{E_i}{V_i} \bar{U}_i, \tag{14}$$

where $E_i \equiv \int e(a, f_i) n_i(a)\, da$ is firm emissions and $\bar{U}_i \equiv \int U(a, f_i) \frac{e(a, f_i) n_i(a)}{E_i}\, da$. Thus, aggregate firm emissions measure the current-flow component of exposure, while $\bar{U}_i$ captures how current emissions map into expected future emissions exposed to the tax.

To take (14) to data we need an empirical proxy for $\bar{U}_i$. We summarize each firm by two observables, $\log E_i$ and an average capital age measure $\bar{a}_i$, and approximate $\log \bar{U}_i$ by a second-order Taylor expansion in $(\log E_i, \bar{a}_i)$ around their cross-sectional means. Dropping the pure quadratic terms for parsimony[3] gives

$$\log \bar{U}_i \approx c_u + \theta_E \left(\log E_i - \overline{\log E}\right) + \theta_a \left(\bar{a}_i - \bar{a}\right) + \theta_{Ea} \left(\log E_i - \overline{\log E}\right)(\bar{a}_i - \bar{a}).$$

Each parameter has a structural reading. θ_E is the elasticity of the emissions-weighted multiple with respect to current emissions, holding age fixed; it captures the persistence of current emissions into expected future emissions. θ_a is the age slope of the multiple at the cross-sectional mean. θ_{Ea} is the local age derivative of the emissions slope, which translates the model's cross-partial result $\partial_{a\phi}\mathsf{s} < 0$ from Proposition 1 into the empirical prediction $\theta_{Ea} < 0$. Current emissions map less sharply into priced exposure for firms with older capital.

Two final steps connect this specification to a return regression. First, taking logs of (14) and substituting the approximation for $\log \bar{U}_i$ expresses $\log \mathsf{S}_i$ as a linear function of $\log E_i$, $\bar{a}_i$, $\log V_i$, and the $\log E_i \times \bar{a}_i$ interaction. Second, expected excess returns from (13) are linear in S_i, not in $\log \mathsf{S}_i$, so we apply the local linearization $\mathsf{S}_i \approx \bar{\mathsf{S}}(1 + \widetilde{\log \mathsf{S}_i})$ around the

[3]The within-quintile young-dirty pattern is qualitatively unchanged when the dropped quadratic terms (and within-quintile age and age-squared controls) are added back.

cross-sectional mean, where $\widetilde{\log \mathsf{S}_i} \equiv \log \mathsf{S}_i - \log \bar{\mathsf{S}}$. Substituting into the pricing equation yields the regression

$$E[r_i^e] \approx c + \beta_1 \log E_i + \beta_2 \bar{a}_i + \beta_3 (\log E_i \times \bar{a}_i) + \Gamma' X_i, \tag{15}$$

The dependent variable is the equity return whereas the model's expected-return expression is per dollar of installed capital. Equity value also prices growth opportunities, technology-switching options, intangible rents, and other claims, so two firms with the same emissions and the same capital age can have very different equity values for reasons unrelated to stranding. Including controls in X for scale, PP&E, and valuation absorbs these non-installed-capital sources of equity value, so that the coefficient estimates β_1 and β_3 identify the stranding risk implied by the model.

The sign of the return premium depends on the aggregate price of exposure. Higher emissions proxy greater installed fuel exposure, so the emissions coefficient inherits the sign of the pricing scalar $\lambda\xi\tau$ in (13). It is positive when cost-side policy exposure is priced as a bad aggregate-state risk ($\beta_1 > 0$), zero if the shock is unpriced, and negative if the aggregate pricing sign reverses. The main age effect does not have a clean structural interpretation, since once current emissions are fixed, the impact of age depends on the remaining exposure, the value denominator, and other capital-age channels. Instead, the sharp prediction is the age gradient of dirty-capital exposure. Proposition 1 shows that the stranding-exposure gradient with respect to machine dirtiness declines with capital age, which maps to $\beta_3 < 0$ when the aggregate price of stranding risk is positive. This is a physical-capital prediction because stranding operates through installed assets whose fuel requirements are fixed, so the gradient should be strongest where tangible capital is central to production and weakest where it is not. The broad-market tests therefore split industries by capital intensity.

We test the model's predictions in two complementary settings. In the utility sector (Section 3), detailed micro-data allows us to measure capital age directly and examine the age gradient of returns in narrow event windows, where confounding variation is limited. In the broad cross-section (Section 4), we proxy capital age with an accounting-based measure and test whether the interaction between age and emissions appears in monthly returns.

3 Event-study evidence from electric utilities

We begin by testing the model's prediction for returns around policy announcements in the sector where they are most likely to be readily detectable. 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 regulatory events (EPA rules and Supreme Court decisions that directly govern power plant emissions) and test whether the model's predictions hold in narrow event windows where confounding variation is limited. This event-study design is related to evidence that climate-policy and climate-news shocks are rapidly capitalized into green-minus-brown returns, including U.S. election shocks in Ramelli et al. (2021), media climate-concern shocks in Ardia et al. (2023), and high-frequency legislator tweets in

DeMiguel et al. (2026). Our empirical strategy examines whether the same event-induced repricing is stronger within fossil-exposed firms whose installed capital is younger.

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 2005–2025. For each firm, we compute the 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. Full details on sample construction and vintage measures are in Appendix D.

We assemble a calendar of candidate events spanning EPA regulations and Supreme Court decisions over 2005–2024. From this universe, we retain the 19 events that satisfy two criteria: (i) the event has a direct effect on the expected cost of operating fossil-fuel power plants, and (ii) the timing of the event does not coincide with confounding factors that might affect stock prices in the utility sector (for example, one excluded event coincides with the Fukushima nuclear disaster). Each event is signed as $d_e \in \{+1, -1\}$ to indicate whether it tightens or loosens regulation. Table D1 lists all events, including exclusion reasons where applicable.

3.2 Portfolio event study

For each event, we sort firms into terciles by fossil (or renewable) vintage and form value-weighted top-tercile ("Old") and bottom-tercile ("Young") portfolios. Because firms with an overall old fleet may have both old fossil and old renewable assets, we isolate fuel-specific age by projecting it onto total fleet age, and use the orthogonal component to form portfolios. The signed Old−Young return spread, $d_e \times (R^O_{e,t} - R^Y_{e,t})$, measures the difference in cumulative returns when cost-side regulation tightens. This is the event-time counterpart of the exposure object in Section 2.3. High-exposure firms earn higher expected returns in normal times, but lose more when policy news raises expected regulatory costs. The model therefore predicts a positive signed spread for fossil vintage and a zero gradient for the renewable vintage through the direct fuel-cost channel. We cumulate the cross-event average and report t-statistics for the cumulative return window across events (Appendix D).

Figure 2 presents the results. Panel (a) plots the signed cumulative average abnormal return (CAAR) spread for fossil vintage across the 19 events. The spread is economically small and statistically insignificant in the pre-event window, then rises sharply beginning on the event day, reaching +0.73% by day +10 ($t = 3.28$). Firms with older fossil fleets significantly outperform those with younger fossil fleets when regulation tightens, as theory predicts.

Panel (b) repeats the sort using renewable vintage. The Old−Young spread for renewable vintage is +0.16% by day +10 ($t = 0.40$), statistically indistinguishable from zero, consistent with the model's prediction.

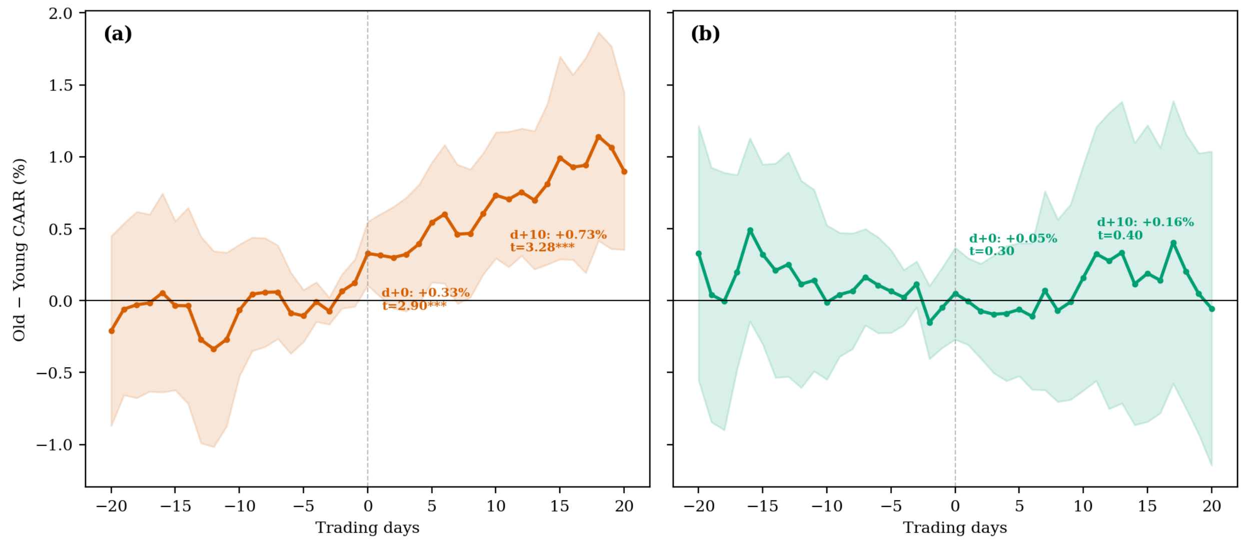

Figure 2: Old−Young CAAR spread around cost-side regulatory events.

Notes: Panels (a) and (b) plot the spread for fossil and renewable vintage, respectively. Value-weighted tercile portfolios, 19 events, signed by regulatory direction. Shaded bands show 95% confidence intervals based on the event-level covariance matrix. The fossil spread is positive and significant under equal weighting, median rather than tercile splits, BMP-standardized inference, and leave-one-event-out (Table D2).

3.3 Cross-sectional regressions

The portfolio sorts displayed in Figure 2 document the size 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 (15), with the dependent variable being the firm-level cumulative abnormal return (industry-adjusted, using the value-weighted utility sector portfolio as benchmark) and the key regressor being fossil (or renewable) fleet age:

$$\mathrm{CAR}_{i,e}(0,h) = \alpha_e + \beta_1\,\mathrm{age}^{\mathrm{fossil}}_{i,e} + \Gamma' X_{i,e} + \delta' ISO_{i,e} + \varepsilon_{i,e}, \tag{16}$$

Here $\mathrm{CAR}_{i,e}(0,h) \equiv d_e \sum_{\tau=0}^{h} \left(R_{i,e,\tau} - R^{\mathrm{util}}_{e,\tau}\right)$ is the signed cumulative abnormal return for firm i around event e, with $R^{\mathrm{util}}_{e,\tau}$ denoting the value-weighted utility-sector benchmark return. The coefficients α_e are event fixed effects, $X_{i,e}$ includes total fleet age, firm i's fossil generation share, log generation capacity, log market capitalization, log book-to-market, and log leverage; and $ISO_{i,e}$ is a vector that measures the firm's capacity-weighted share across nine ISO/RTO regions.[4]

Table 1 reports estimates for three CAR windows ($h = \{5, 10, 15\}$ days). The coefficient on fossil vintage is positive in every column, with t-statistics ranging from 2.16 to 2.20: older fossil fleets earn higher signed CARs when regulation tightens. The corresponding renewable-fleet coefficients are statistically indistinguishable from zero. When the event set is split by

[4]U.S. wholesale electricity markets are organized regionally, operated by Independent System Operators or Regional Transmission Organizations (ISOs/RTOs). These regions differ in market design, generation mix, and state-level regulatory environment, so they can respond differently to the same event. To address this, we include the firm's capacity-weighted share of generators in each ISO/RTO as controls.

Table 1: Cross-sectional regressions of post-event CAR on generator age

	Panel A: Fossil-fuel generator age			Panel B: Renewable generator age		
	$(0, 5)$	$(0, 10)$	$(0, 15)$	$(0, 5)$	$(0, 10)$	$(0, 15)$
Generator age	$+0.0326$	$+0.0442$	$+0.0522$	$+0.0284$	-0.0272	$+0.0642$
	[2.16]	[2.20]	[2.17]	[0.66]	[-0.53]	[0.68]
N	744	744	744	445	445	445
R^2	0.131	0.151	0.157	0.162	0.262	0.277

Notes: Sample: 19 regulatory events (see Appendix D). Dependent variable: industry-adjusted signed CAR, in percentage points. All regressions include event fixed effects. Controls: total fleet age, fossil generation share, log generation capacity, log market capitalization, book-to-market ratio, leverage and ISO shares. ISO shares are the firm's capacity-weighted share in each ISO/RTO bucket. t-statistics in square brackets use heteroskedasticity-consistent (HC3) standard errors.

policy direction, significance is concentrated in tightening events. Loosening-event estimates have the predicted sign but are statistically less significant, with the largest t-statistic of 1.49 in the $(0, 15)$ window (Table D4). Table D3 shows robustness to FF3- and FF5-adjusted CAR. Coefficient estimates on fossil fuel age in the factor-adjusted specifications are positive across all windows and are most precisely estimated in the shortest post-event window.

4 Broad market evidence

We now extend the analysis from the utility sector to the broad stock market. In Subsection 4.2 we consider return regressions that estimate the effect of emissions and emissions interacted with capital age, and document a negative emissions-age gradient, consistent with our model. Subsection 4.3 provides estimates obtained from rolling-window regressions to examine whether these coefficients exhibit significant time variation and document higher carbon premiums in periods of heightened policy salience. Finally, Subsection 4.4 reports the performance of a signal-timed portfolio that goes long stranding risk when climate policy-related news is salient, and short when news is not salient.

4.1 Data and variable construction

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 2025 and contains approximately 391,000 firm-month observations after imposing the standard filters described below. The broad-market analyses use returns through December 2025. Table 2 reports descriptive statistics for the main variables in the full sample and the relevant sub-samples.

Following Lin et al. (2020) and Geelen et al. (2024), we proxy the model's value-weighted average vintage $\bar{a}_i$ (equation (15)) with a recursive capital age measure constructed from Compustat balance-sheet data and defined in Appendix E.1.

As a validation check, Figure E1 plots the recursive measure against capacity-weighted generator age across utilities. The bin-scatter slope is $\hat{\beta} = 1.40$ ($p < 0.001$), meaning a one-year increase in EIA generator age corresponds to a 1.4-year increase in the recursive measure. The cross-firm Pearson r is 0.27, which we view as reasonable considering that utilities invest in many types of capital beyond generators, including transmission, distribution, etc. A further concern with capital age is that it may not reflect the emissions intensity of capital if firms green their capital stock as they invest. Table E2 shows that among capital-intensive firms, younger capital age predicts higher future-to-current Scope 1 emissions over horizons up to 5 years.[5]

The mapping in Section 2.4 shows that firm-level stranding exposure is proportional to current emissions. Scope 1 emissions therefore provide the observed flow component of exposure, while capital age proxies for how that flow maps into remaining fuel commitments.[6] PP&E, scale, and valuation controls proxy the installed-capital value denominator and its relation to equity value. As a robustness check, we also consider regressions that use the log of emissions scaled by net PPE as a measure of emissions intensity. Net PP&E is the natural scaling variable within the context of our model.

Industries are classified as "capital-intensive" or "capital-light" based on the ratio of gross PP&E to total assets at the SIC-2 level, measured from Compustat. The cutoff is chosen to balance the number of firm-month observations in each group (Appendix E.1). The model implies that capital irreversibility, and hence the bite of stranding risk, matters only where physical fossil-fuel capital is substantial. This motivates our focus on capital-intensive industries. Appendix Table E1 lists the highest- and lowest-emission SIC-2 industries within each capital-intensity group.

4.2 Cross-sectional return tests

Our baseline regression is

$$r_{i,t} = \alpha + \beta_1 \log E_{i,t} + \beta_2 \bar{a}_{i,t} + \beta_3 (\log E_{i,t} \times \bar{a}_{i,t}) + \Gamma' X_{i,t} + \text{FE} + \varepsilon_{i,t}, \tag{17}$$

where $r_{i,t}$ is the monthly excess stock return (CRSP return minus the risk-free rate), $E_{i,t}$ is Scope 1 emissions, $\bar{a}_{i,t}$ is recursive capital age, and $X_{i,t}$ is a vector of controls. Emissions and age variables are z-scored within each cross-section (calendar month). The coefficient β_1 captures the carbon premium evaluated at the sample-mean capital age; β_3 is the coefficient estimate for the age $\times$ dirty interaction that the model predicts should be negative if $\beta_1 > 0$, meaning that the emissions premium is larger for younger-capital firms and attenuates with capital age. In the cross-sectional return tables, coefficients and standard errors are reported in basis points per month.

[5] A further issue is that acquisition-driven PP&E lowers the recursive age measure, making possibly old capital look young. Table E7 shows robustness to excluding firm-years with elevated capital acquisitions.

[6] We also estimate the baseline specification using Scope 2, Scope 3, and Scope 1+2 emissions (Table E3).

Table 2: Descriptive Statistics by Capital Intensity and Capital-Age Quintile

	Full sample			Post-Paris capital-intensive by capital-age quintile					
	All	Capital-intensive	Capital-light	All	Q1	Q2	Q3	Q4	Q5
Firm-months	320,972	178,345	142,627	120,160	22,195	23,681	23,991	25,128	25,165
Excess return (% per month)									
Mean	0.58	0.45	0.74	0.31	-0.17	0.06	0.36	0.48	0.75
SD	13.74	14.24	13.09	15.74	18.62	16.74	15.78	14.29	13.09
Median	0.40	0.31	0.52	-0.03	-0.65	-0.32	0.00	0.08	0.42
P25	-6.03	-6.22	-5.78	-7.34	-9.86	-8.44	-7.32	-6.46	-5.62
P75	6.65	6.58	6.73	7.11	8.22	7.47	7.13	6.63	6.52
Capital age (years)									
Mean	4.58	5.17	3.84	5.03	2.36	3.58	4.60	5.82	8.36
SD	2.28	2.46	1.78	2.55	1.07	1.38	1.44	1.47	2.00
Median	4.16	4.80	3.52	4.60	2.06	3.28	4.30	5.47	7.97
P25	2.83	3.30	2.49	3.06	1.63	2.61	3.57	4.78	6.75
P75	5.80	6.61	4.83	6.50	2.68	3.93	5.07	6.40	9.85
ln(*Scope 1 Emissions*)									
Mean	10.20	11.07	9.12	10.34	8.92	9.84	10.25	10.87	11.63
SD	2.96	3.19	2.23	3.25	3.17	3.21	3.28	3.03	2.93
Median	10.15	11.12	9.20	10.39	8.99	10.04	10.35	10.80	11.57
P25	8.23	9.18	7.56	8.36	6.75	7.74	8.26	9.14	9.62
P75	12.07	13.21	10.63	12.58	10.83	11.97	12.40	12.92	13.65

We control for the firm characteristics used by Bolton and Kacperczyk (2021): ln(market cap), ln(PPE), book-to-market, momentum, ROE, leverage, investment-to-assets, sales growth, EPS growth, the Herfindahl–Hirschman index, market beta, and return volatility. Annual accounting variables are mapped to monthly returns using the standard Fama–French timing convention. Monthly controls are lagged by one month. The baseline specification includes year-month fixed effects and SIC-2 industry fixed effects.

Table 3 reports the results for our main specification in two time periods and three industry samples. Three patterns emerge. First, the carbon premium is concentrated in capital-intensive industries. Second, in these industries, the carbon premium decreases significantly with age. Third, these effects are most pronounced following the Paris Agreement.

At the average capital age, the within-industry carbon premium is 26 basis points per month in the full sample and 28 basis points in the capital-intensive full sample, with no significant emissions-age gradient in either. Bolton and Kacperczyk (2021) report 48 basis points per month for the same specification but a different sample. We attribute the difference to the time-varying premium documented below. In capital-intensive industries after the Paris Agreement,[7] the premium at the average age rises to 42 basis points per month and declines

[7] We use January 2016, the month after the Paris Agreement was reached, as the sub-sample start point for the post-Paris period for consistency with the existing carbon-premium literature, and return below to a formal break test and rolling-window estimates of the young-capital premium.

by 5 basis points per year of capital age. Capital-light industries do not show a carbon premium at any age.

Table 3: Full-sample carbon premium estimates

	All industries		Capital-intensive		Capital-light	
	Full sample	Post-Paris	Full sample	Post-Paris	Full sample	Post-Paris
Emissions	25.7	34.5	27.9	41.7	19.3	23.1
	[2.29]	[2.48]	[2.17]	[2.58]	[1.56]	[1.31]
Capital age	6.3	3.6	12.5	10.6	1.9	3.7
	[0.96]	[0.45]	[1.64]	[1.13]	[0.20]	[0.31]
Emissions × Capital age	0.1	−2.5	−5.9	−12.9	3.3	7.1
	[0.02]	[−0.67]	[−1.22]	[−2.11]	[0.56]	[1.01]
Obs	357,703	254,112	179,502	121,195	178,201	132,917

Notes: Each cell reports the coefficient from a regression of monthly excess returns on ln(Scope 1 Emissions), capital age, their interaction, and controls, with year-month and SIC-2 fixed effects. Coefficients are reported in basis points per month. 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; interaction terms are products of the corresponding z-scores. Industries are split into "capital-intensive" and "capital-light" using a fixed SIC-2 classification based on gross PP&E / total assets measured in Compustat over 1985–1995, adjusted to balance firm-month counts across the two groups. "Post-Paris" covers January 2016 to December 2025. t-statistics based on two-way clustered standard errors (firm × calendar year) are reported in square brackets.

Table 4 focuses on the post-Paris Agreement, capital-intensive sample in which the carbon premium is most apparent. The table provides an alternative set of estimates that examines robustness to using emissions scaled by net PP&E as the explanatory variable and includes a full set of time-by-industry fixed effects that absorb all industry-specific time variation in returns. Using scaled emissions reduces the carbon premium by one-third but leaves the emissions-by-age coefficients unchanged. As a result, the capital-age emissions channel accounts for a larger share of the overall premium. All results are robust to the full set of industry-time controls.

Figure 3 allows for a non-parametric specification of the age-emissions gradient by providing distinct estimates of the carbon premium within quintiles of firms sorted by capital age within industry in each period.[8] We estimate the emissions coefficient separately for each quintile in a pooled regression with quintile × emissions × capital-intensity × post-Paris interactions, with year-month and SIC-2 fixed effects plus the same controls as in the baseline specification.

The estimates in Figure 3 confirm our previous finding that the carbon premium is concentrated among capital-intensive industries post-Paris, and declines with age. In particular,

[8]Within-quintile emissions variation is ample. Under a year-month × SIC-2 residualization, 71% of the variation in ln(Scope 1) remains within age quintiles rather than between them. Table E9 reports the variance decomposition and cross-tabulation of firm-month observations by capital age and emissions.

Table 4: The carbon premium in the post-Paris Agreement, capital-intensive subsample

	Panel A: ln(Scope 1 Emissions)			*Panel B:* $\ln\left(\frac{\text{Scope 1 Emissions}}{\text{Net PP\&E}}\right)$		
	(1)	(2)	(3)	(1)	(2)	(3)
Emissions	36.5	28.2	35.8	25.9	18.8	21.9
	[2.47]	[1.96]	[2.42]	[2.55]	[2.02]	[2.46]
Capital age	9.2	7.5	6.8	5.7	4.7	2.7
	[0.79]	[0.74]	[0.56]	[0.52]	[0.50]	[0.24]
Emissions × Capital age	−12.7	−11.3	−13.7	−12.8	−11.1	−11.2
	[−2.00]	[−1.94]	[−2.05]	[−2.33]	[−2.03]	[−1.74]
Obs	120,160	120,160	120,160	120,160	120,160	120,160
B&K controls	Yes	Yes	Yes	Yes	Yes	Yes
SIC-2 FE	Yes	—	—	Yes	—	—
Year-month FE	Yes	—	—	Yes	—	—
SIC-2 × year-month FE	No	Yes	No	No	Yes	No
SIC-3 × year-month FE	No	No	Yes	No	No	Yes

Notes: The dependent variable is the monthly excess stock return. Coefficients are reported in basis points per month. Panel A measures emissions as ln(Scope 1 Emissions); Panel B as ln(Scope 1 Emissions/net PP&E) (Compustat `ppent`). All emissions and age variables are within-month z-scored; interaction terms are products of the corresponding z-scores. Within each panel, column (1) includes year-month and SIC-2 fixed effects; column (2) replaces separate year-month and industry fixed effects with SIC-2 × year-month fixed effects; column (3) replaces them with SIC-3 × year-month fixed effects. All regressions include ln(market cap), book-to-market, momentum, ROE, leverage, investment/assets, ln(PPE), sales growth, EPS growth, HHI, market beta, and return volatility as controls. Sample: capital-intensive industries, defined by a fixed SIC-2 classification based on gross PP&E / total assets measured in Compustat over 1985–1995 and adjusted to balance firm-month counts across capital-intensive and capital-light groups, January 2016 to December 2025. t-statistics based on two-way clustered standard errors (firm × calendar year) are reported in square brackets.

in the post-Paris period, the youngest-quintile emissions premium for capital-intensive firms is 83.4 basis points per month, while the oldest-quintile estimate is −0.1 basis points per month.[9]

Two channels discussed in the literature could potentially provide alternative explanations for the cross-sectional pattern that we uncover in our estimation of Equation (17). The first is the Papanikolaou (2011); Kogan and Papanikolaou (2013) growth opportunities channel.

[9] A consistent age gradient appears in valuations. Replacing monthly returns with annual $\log(B/M)$ as the dependent variable as in Bolton and Kacperczyk (2023), the youngest-quintile coefficient is the largest positive coefficient for capital-intensive firms (Table E8).

Our theory relates the carbon premium to an assets-in-place channel, but a possible alternative explanation for our empirical findings is that the risk of policy implementation differentially affects the value of future investment opportunities. If growth opportunities decline in response to the carbon tax and growth opportunities are larger for firms with newer capital, our estimates would also reflect this mechanism. A second concern is the financing channel studied by Geelen et al. (2024), under which debt maturity is negatively correlated with capital age. Table E5 reports the results of adding controls to capture these two channels, separately and jointly, to the pooled and per-quintile specifications. The youngest-quintile emissions coefficient attenuates mildly, and the oldest quintile remains close to zero, so the age gradient does not appear to be driven by either alternative explanation.

We also implement robustness checks that address recent measurement critiques of the carbon-premium literature. Lagging emissions by one additional fiscal year and controlling for log(sales) addresses the concern that return tests use carbon data not yet available to investors and contain forward-looking sales information (Aswani et al., 2024; Zhang, 2025). Our main finding that the carbon premium is concentrated in firms with high young-capital intensity is robust to this variation in the control set (Table E4). Our findings also help interpret the evidence in Crosignani et al. (2025), which shows that estimates of carbon pricing are highly sensitive to super-emitters, especially electric utilities. Through the lens of our model, this sensitivity is not an anomaly. Utilities are likely one of the sectors in which physical capital intensity and fleet vintage matter most, so a single emissions slope averages over firms whose stranding exposure differs sharply with the age of installed assets.

4.3 State dependence of the carbon premium

We now study the time variation in the carbon premium across the capital-age distribution for our subset of capital-intensive firms. In each month t, we form capital-age quintiles using fixed cutoffs from the full sample of such firms. We then estimate equation (17), adding quintile dummies and quintile $\times$ emissions interactions to the specification in rolling twelve-month window ($[t-11, t]$) subsamples.

Figure 4 plots the estimated rolling carbon premium for the youngest (Q1) and oldest (Q5) capital-age quintiles. The contrast between the two panels is sharp. In Panel (a), the Q1 carbon premium exhibits substantial time variation. The stranding premium increases sharply in the period between the signing of the Paris Agreement and the 2016 Presidential election, after which the carbon premium declines sharply, consistent with the election outcome generating a large downward revision in the expectation of future carbon taxes. The premium increases again around 2020–2021, coinciding with the Biden administration's regulatory push, and weakens sharply following the second Trump inauguration. In contrast, as displayed in Panel (b), the Q5 premium is consistently close to zero and shows very little variation around the Paris Agreement. These findings are consistent with our maintained hypothesis that dirty firms with long remaining fuel commitments bear the pricing of transition-risk due to regime shifts in carbon tax policy.

To formally test for regime changes in the carbon premium, we conduct Bai and Perron (1998, 2003) multiple-break tests. Applied to the rolling Q1 coefficient, the test identifies a single break with point estimate March 2016, within months of the signing of the Paris Agreement in December 2015.[10] Using this break date, the Q1 coefficient averages -26 bps/month before the break and +68 bps/month after the break. The pre-break mean is not statistically distinguishable from zero, while the post-break mean is significantly above zero.[11] The Q5 coefficient moves in the same broad direction, but its segment means are economically small and statistically indistinguishable from zero.

The finding that the young-capital carbon premium is negative prior to the Paris Agreement may be understood within the context of the frameworks proposed by Gârleanu and Pedersen (2025) and Kacperczyk (2026). In these frameworks, the observed carbon premium combines a transition risk associated with carbon taxes with a climate-hedging motive arising from investor preferences. As a result, assets whose cash flows are high in low-tax, high-damage states command lower expected returns because they hedge unspanned climate damages. Our model speaks to the exposure side of this decomposition. Young dirty capital has the largest normalized exposure to policy risk because its remaining fossil-fuel commitments are the longest lived. When policy risk is salient, this exposure carries a positive transition-risk premium; when the perceived risk of future policy is low, or policy rollback becomes more likely, this transition component weakens, and the hedging component dominates, potentially reversing the sign of the young-capital premium. We therefore view the sign changes as suggestive evidence that stranding exposure may be priced differentially across policy regimes.

Visual inspection of the rolling estimates suggests that the young-capital premium rises in periods when climate policy is more salient. Building on the literature that uses text-as-data to measure climate and political risk (Baker et al., 2016; Hassan et al., 2019; Engle et al., 2020; Sautner et al., 2023; Feng et al., 2025)To provide evidence for this conjecture, we consider how our estimates vary with an external text-based measure of climate-policy salience. We use, the Sentometrics U.S. Media Climate Change Index (MCCI), an updated monthly version of the Media Climate Change Concerns index of Ardia et al. (2023). The index is constructed from U.S. news coverage of climate change-related issues, decomposed into topic subindices. We use the Legal Actions subindex as the baseline climate-attention signal because it isolates news about legal, regulatory, and administrative actions related to climate change, which are likely to affect investors' assessments of stranded-asset risk for emissions-intensive capital. We therefore interpret it as a proxy for episodes in which climate-policy risk becomes more salient and estimate, for each capital-intensity group, a regime-dependent carbon-premium specification.[12]

Figure 5 shows the one-month-lagged MCCI Legal Actions signal after real-time detrending

[10] The LWZ criterion selects $k^* = 1$ break, and sequential sup-F tests find no evidence of additional breaks. The test uses a 15% trimming parameter with up to five breaks.

[11] Using HAC standard errors with twelve lags, the 95% confidence intervals are $[-56.3, 3.5]$ for the Q1 pre-break mean and $[33.1, 103.0]$ for the Q1 post-break mean, and the corresponding Q5 intervals are $[-28.7, 12.6]$ and $[-18.5, 29.6]$.

[12] Specifically, for $g \in \{\text{Capital-intensive}, \text{Capital-light}\}$, the estimating equation is $r_{i,t} = \sum_{q=1}^{5} \left[\beta_{gq}^{L}(1-H_t) + \beta_{gq}^{H} H_t\right] z_{i,t}^{E} \mathbf{1}\{Q_{i,t} = q\} + \sum_{q=2}^{5} \left[\delta_{gq}^{L}(1-H_t) + \delta_{gq}^{H} H_t\right] \mathbf{1}\{Q_{i,t} = q\} + \Gamma_g' X_{i,t} + \text{FE} + \varepsilon_{i,t}$. H_t equals one when the one-month-lagged MCCI Legal Actions signal is above its expanding-window median, $Q_{i,t}$ is the capital-age quintile, and $z_{i,t}^{E}$ is within-month standardized log Scope 1 emissions.

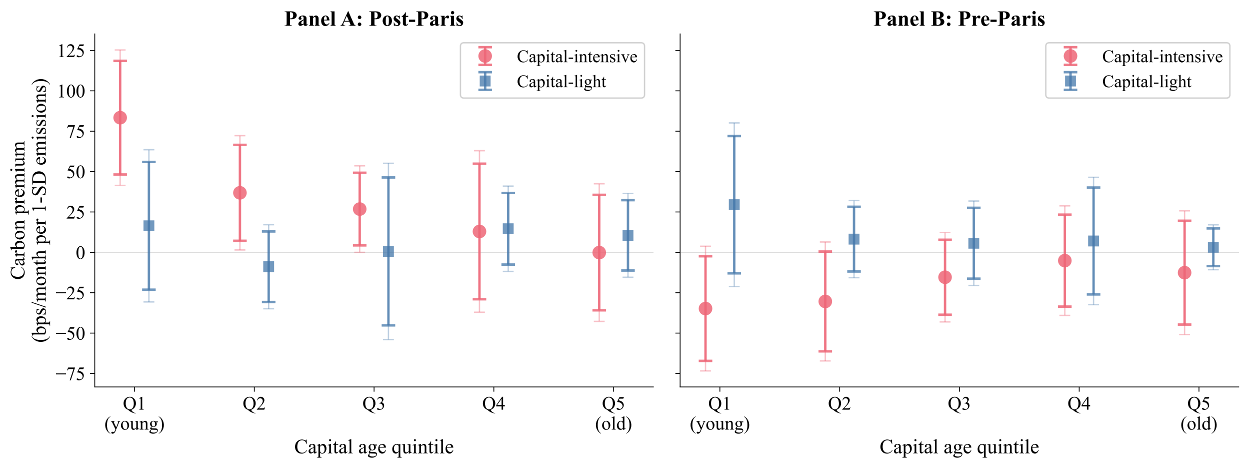

Figure 3: Carbon premium by capital-age quintile.

Notes: Estimated from a pooled regression with quintile $\times$ ln(Scope 1) $\times$ capital-intensity $\times$ post-Paris interactions. Each marker shows the emissions coefficient for one 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. Panel A: post-Paris (2016–2025); Panel B: pre-Paris (2005–2015). Controls and fixed effects as in Table 4. Age quintiles are formed using balanced cutoffs computed within each SIC-2 industry $\times$ period (pre/post-Paris). The regression includes the baseline controls (see Section 4.2), year-month and SIC-2 fixed effects. Table E6 reports the numerical coefficients and their robustness to estimating them via the Fama-MacBeth procedure.

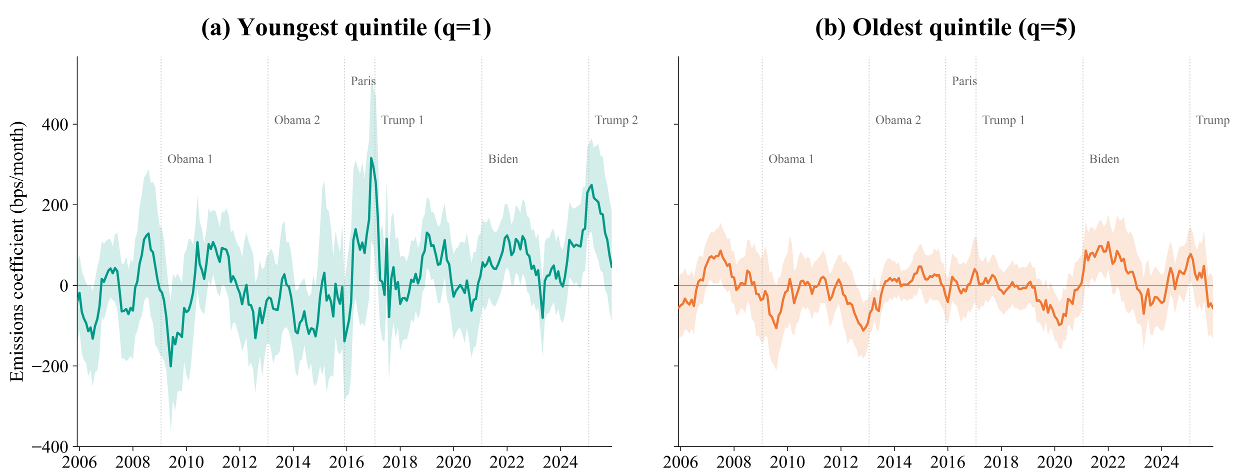

Figure 4: Rolling estimated carbon premium by capital-age quintile.

Notes: For each month t, we run a regression of excess returns on age-quintile dummies, age-quintile $\times$ $z(\log E)$ interactions, $z(\bar{a})$, baseline controls (see Section 4.2), year-month, and SIC-2 fixed effects on the twelve-month window $[t-11, t]$ within the capital-intensive subsample. Age quintiles are formed using fixed absolute cutoffs defined on the full capital-intensive sample. Standard errors are clustered by firm. Panel (a) plots the rolling emissions coefficient for the youngest quintile (Q1); Panel (b) plots the same for the oldest quintile (Q5). Solid lines show rolling point estimates in basis points per month; shaded bands are pointwise 95% confidence intervals, included as descriptive precision bands for the rolling estimates. Vertical dotted lines mark regime changes (Obama 1, Obama 2, Paris Agreement, Trump 1, Biden, Trump 2).

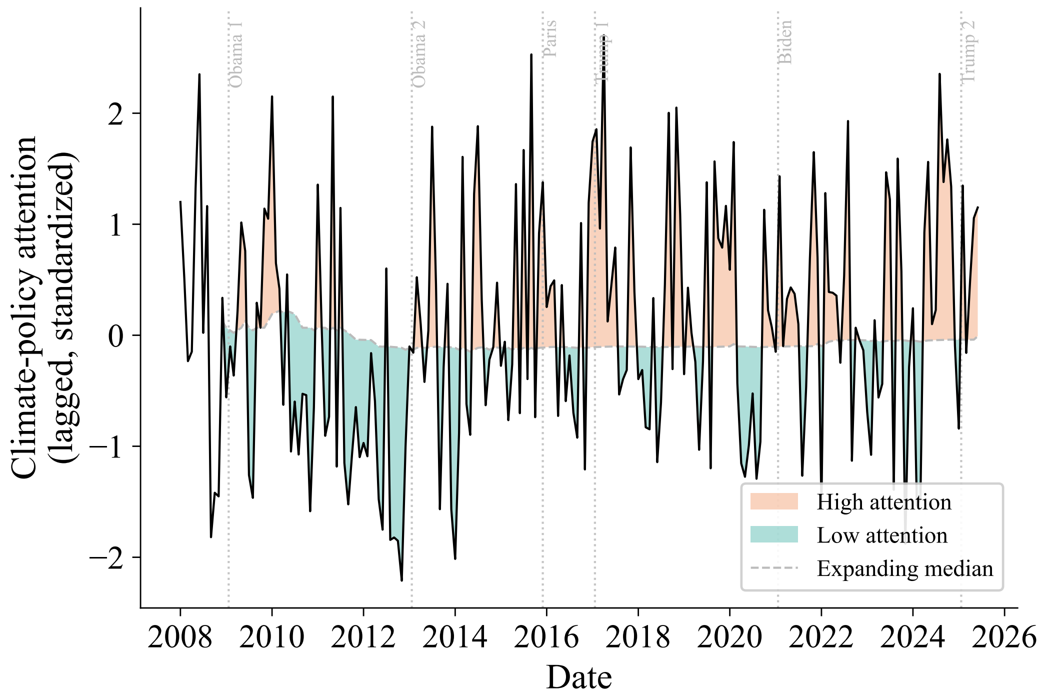

Figure 5: Climate-policy attention states.

Notes: The figure plots the one-month-lagged MCCI Legal Actions signal used to define the regime indicator in Table 5. At each month, the raw MCCI series is detrended using only observations available through the signal month and scaled by the trailing twelve-month volatility of those real-time residuals. Orange shading marks months in which the lagged signal is above its expanding-window median; teal shading marks months below that cutoff. The displayed series begins in January 2008 to omit the initial burn-in period. Vertical dotted lines mark regime changes (Obama 1, Obama 2, Paris Agreement, Trump 1, Biden, Trump 2).

and volatility normalization, with orange (teal) shading indicating when this lagged signal is above (below) its expanding-window median. Table 5 reports the emissions premium within each regime and capital-age quintile. The point estimates indicate a positive young-capital premium in months with elevated climate policy attention among capital-intensive firms, and no significant premia in low-salience months. In contrast, consistent with Table 3, the corresponding premia are smaller and less precise for capital-light firms in both regimes.

The pattern of larger carbon premia during times of elevated attention to climate policy is robust to using alternative text-based salience measures. Table E10 repeats the low- and high-state estimates for the aggregate MCCI and Climate Legislation/Regulations sub-indices, using the same one-month lag and expanding-median cutoff as the baseline Legal Actions split. Within the capital-intensive subsample, the high-state carbon premium for the youngest capital-age quintile is consistently large and exceeds the corresponding premium for the oldest quintile or for capital-light firms.

Table 5: Carbon premium by capital-age quintile and climate-policy attention regime

	Panel A: Capital-intensive industries		*Panel B: Capital-light industries*	
Age quintile	Low regime	High regime	Low regime	High regime
Q1 (youngest)	8.2	112.0	13.8	33.3
	[0.26]	[3.25]	[0.37]	[1.89]
Q2	−21.2	70.1	−5.9	−3.6
	[−0.73]	[2.10]	[−0.25]	[−0.19]
Q3	−22.2	68.5	9.5	−2.7
	[−0.89]	[2.77]	[0.43]	[−0.13]
Q4	−2.5	49.7	−4.8	7.0
	[−0.13]	[2.70]	[−0.22]	[0.40]
Q5 (oldest)	−12.9	26.7	−13.1	18.1
	[−0.41]	[1.24]	[−0.88]	[0.95]
Obs	152,823		123,833	
Low-regime obs	71,083		56,610	
High-regime obs	81,740		67,223	

Notes: The table shows full-sample estimates of the regime-dependent carbon-premium specification described in the text, shown separately for capital-intensive and capital-light industries. Coefficients are reported in basis points per month. Controls and fixed effects follow Table 3. The sample is the full available monthly panel from December 2008 to June 2025 (199 calendar months; 276,656 firm-months) with non-missing lagged MCCI Legal Actions regime, emissions, age-quintile assignment, returns, fixed effects, and controls. t-statistics based on two-way clustered standard errors (firm $\times$ calendar year) are reported in square brackets.

4.4 Conditional portfolio performance

The state dependence documented above clarifies why realized-return tests of the premium on brown-versus-green assets can be weak. Eskildsen et al. (2026) shows that green-minus-brown portfolio returns are fragile across greenness measures, portfolio constructions, countries, and risk adjustments, and argues that realized returns are too noisy to reliably identify the premium. Our evidence provides one possible explanation for this fragility. The carbon premium is not permanent and unconditional, but rather an episodic stranding premium, concentrated in young dirty capital and appearing when policy risk is salient.

To examine whether conditioning on a signal correlated with policy salience can deliver significant portfolio returns, we study the performance of a portfolio strategy that goes long stranding risk when policy salience is high and short when it is low.

Our exercise follows the factor-timing literature, which studies whether exposure to a pre-specified return spread can be conditioned on lagged public information (Moreira and Muir, 2017; Haddad et al., 2020). As in the previous section, we use a news text-based measure of climate-policy attention, in the spirit of prior work that uses text-as-data to measure climate and political risk (Engle et al., 2020; Baker et al., 2016; Hassan et al., 2019; Sautner et al., 2023; Feng et al., 2025).

First, we construct a portfolio exposed to stranding risk. In each month, within the universe of capital-intensive firms, we sort firms independently into terciles of Scope 1 emissions and recursive capital age within SIC-2 industries. We then compute NYSE-p80-capped value-weighted excess returns for the four corner portfolios: young-clean (YC), young-dirty (YD), old-clean (OC), and old-dirty (OD), where young and old denote the bottom and top age terciles and clean and dirty denote the bottom and top emissions terciles.[13] The STRAND portfolio is constructed by taking the double difference

$$\frac{1}{2}\left[(R_t^{YD} - R_t^{YC}) - (R_t^{OD} - R_t^{OC})\right],$$

which compares young dirty firms to young clean firms and old dirty firms to old clean firms, thereby capturing time variation in stranding exposure, while netting out overall exposure to capital age or emissions. We then condition exposure to this portfolio using the MCCI Legal Actions subindex, described in the previous section. We study two conditional strategies. The median strategy is long STRAND when the lagged signal is above its expanding-window median and short when it is below. The Q4–Q1 rule is long in the top quartile, short in the bottom quartile, and in cash otherwise. One caveat concerns the vintage of the data used to construct the signal. Because our signal is based on the latest available vintage of the text-based measure, rather than the contemporaneous vintages, the exercise cannot be interpreted as a real-time trading strategy.

Table 6 reports the results. The first rows compare the (always long) brown-minus-green benchmark and the static (always long) STRAND strategy; the signal-timed rows then condition STRAND exposure on the text signal. The unconditional always-on stranding portfolio performance is weak, consistent with the view that realized green-minus-brown returns are hard to detect in short, noisy samples. In contrast, the Q4–Q1 signal-timed rule has an annualized Sharpe ratio above one before trading costs and a positive pre-cost FF3 alpha. It also retains a positive gross alpha after controlling for FF5, momentum, and the industry-neutral GMB factor, which supports the interpretation that the return spread captures the age-dependent component of carbon risk rather than generic brown-minus-green exposure.[14] Panel B reports transaction-cost sensitivity for the Q4–Q1 rule: the alpha remains positive and statistically significant for 10 and 15 basis-point one-way costs, and remains positive but weaker at 25 basis points.

Figure 6 plots the cumulative post-Paris excess returns for the Signal-timed STRAND Q4-Q1 portfolio, and the brown-minus-green portfolio. Shaded regions show whether the signal-timed strategy is long, short, or in cash.

Taken together, these results provide portfolio-level evidence in favor of our maintained hypothesis that the carbon premium captures time-varying policy risk concentrated among young-dirty firms with high exposure to asset stranding.

[13] In the main industry-neutral capital-intensive construction, the average monthly counts are 168 firms in YC, 71 in YD, 93 in OC, and 195 in OD; the corresponding minima are 64, 32, 37, and 79.

[14] Table F1 reports Sharpe ratios and FF3-adjusted alphas using alternative text measures as market-timing signals.

Table 6: Signal-timed STRAND performance and factor spanning

Panel A: Gross performance

Strategy	Sharpe	FF3 α (% p.a.)	FF5+Mom+GMB_IN α (% p.a.)
Brown - green benchmark	0.21	1.4	0.5
		[0.83]	[0.45]
Static STRAND	0.24	1.7	1.2
		[0.99]	[0.75]
Signal-timed STRAND, median	0.79	3.3	3.2
		[2.03]	[2.38]
Signal-timed STRAND, Q4–Q1	1.01	3.7	3.7
		[3.11]	[3.39]

Panel B: Signal-timed STRAND, Q4–Q1 alpha by one-way overlay fee

One-way overlay fee (bp)	FF3 α (% p.a.)	FF5+Mom+GMB_IN α (% p.a.)
10	2.9	2.9
	[2.50]	[2.72]
15	2.6	2.6
	[2.19]	[2.38]
25	1.8	1.8
	[1.56]	[1.69]

Notes: The table reports post-Paris monthly portfolio performance from December 2015–December 2025. Panel A reports gross strategy performance before overlay trading costs. Sharpe ratios are annualized. Alpha estimates are annualized percentage-point intercepts from time-series regressions of monthly strategy excess returns on factor returns; HAC t-statistics are reported in square brackets below alpha estimates. FF3 α is adjusted for the Fama–French market, size, and value factors. FF5+Mom+GMB_IN α is adjusted for the Fama–French five factors, momentum, and the industry-neutral green-minus-brown factor. STRAND is $\frac{1}{2}[(R^{YD} - R^{YC}) - (R^{OD} - R^{OC})]$ within capital-intensive industries, where Y/O denote young/old capital and D/C denote dirty/clean emissions. Signal-timed rows condition STRAND exposure on the one-month-lagged MCCI Legal Actions signal. Panel B reports alpha estimates for Signal-timed Q4-Q1 STRAND after subtracting one-way overlay fees.

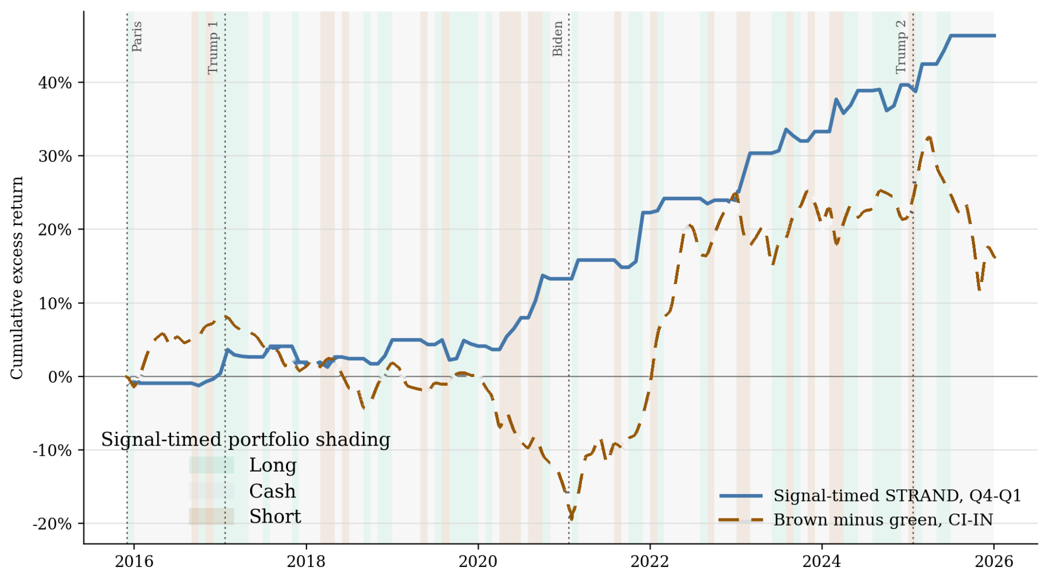

Figure 6: Post-Paris performance of Brown-minus-green and Signal-timed STRAND portfolios

Notes: The figure plots compounded cumulative monthly excess returns from December 2015-December 2025 for Signal-timed STRAND Q4-Q1 (blue solid line) and the industry-neutral Brown-minus-green portfolio (brown dashed line). The signal-timed Q4–Q1 overlay is long when the one-month-lagged MCCI Legal signal is in its top quartile (green shading), short when it is in the bottom quartile (brown shading), and in cash otherwise (unshaded).

5 Conclusion

This paper studies why the carbon premium is concentrated in some dirty firms but absent in others. In a putty-clay model, transition exposure depends on the remaining fuel commitments embedded in installed capital and on the value of the assets to which those commitments attach. This gives the carbon premium a predictable within-brown structure. Dirty firms with young, capital-intensive assets have the largest exposure to policy-induced value loss, while dirty firms with older capital have less exposure.

The evidence is consistent with this mechanism. In the utility sector, where generator-level data allow us to observe fossil-capital age directly, cost-side regulatory events reduce the value of older fossil fleets less than younger ones. In the broader U.S. equity market, the post-Paris carbon premium is concentrated in capital-intensive industries, is largest in the youngest capital-age quintile, and is close to zero in the oldest quintile and in capital-light industries. The pattern is robust to the inclusion of standard carbon-premium controls, finer industry-time fixed effects, lagged emissions, net-PP&E-scaled emissions, and controls for growth opportunities, investment intensity, financing maturity, and acquisition-related measurement concerns.

The carbon premium is also state dependent. Rolling estimates show that the young-capital premium changes sign over time and strengthens when legal and regulatory climate attention is high. A signal-timed STRAND portfolio that conditions exposure to stranding risk on lagged legal attention earns substantially larger returns in the post-Paris period than the corresponding always-on stranding spread or a standard brown-minus-green benchmark. We interpret this portfolio evidence as descriptive evidence on regime dependence. The priced object is intermittent exposure to transition-risk states, not a permanent unconditional spread between brown and green firms.

The findings have direct implications for measurement and capital allocation. For empirical asset pricing, they suggest that tests of the carbon premium should allow priced emissions exposure to vary with the physical capital stock. For investors, carbon risk is not a uniform brown characteristic; it is concentrated where dirty assets have long remaining lives. For policy, the evidence suggests that financial markets penalize new commitments to dirty long-lived capital more strongly than the continued operation of old dirty capital, whose stranding exposure has largely depreciated.

Several questions remain. The model takes the vintage distribution of capital as given and prices a one-time policy arrival. A natural extension is to embed the vintage structure in a general-equilibrium economy with Epstein–Zin preferences and solve for the dynamic path of the capital stock under an uncertain transition. Such a framework would tie the cross-sectional age gradient and the regime dependence of the premium to a single state variable, the perceived long-run path of policy, and would let transition premia feed back into investment, entry, and retirement, so that the vintage distribution that determines exposure becomes endogenous.

More work is also needed on when climate-policy news is priced as bad news for aggregate consumption and when policy rollback, delay, or abandonment makes brown assets hedges.

The source of stranding risk matters here. A carbon tax strands dirty capital in proportion to remaining fuel exposure and is contractionary on impact, so stranding losses arrive in high-marginal-utility states. Green technical change strands the same capital through relative prices and wages, and since it is good news for long-run productivity its covariance with marginal utility can differ in sign, as in the literature on investment-specific technology shocks. The two sources can therefore produce similar stranding losses with different, and possibly opposite, risk prices. A full treatment requires general equilibrium, where green technical change moves energy prices, wages, and the consumption path simultaneously. Finally, the evidence here is based on U.S. equities; extending the analysis to international markets and other asset classes would show whether the age structure of dirty capital is a general determinant of transition-risk pricing.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the author(s) used [NAME OF TOOL/SERVICE] for [REASON]. The author(s) reviewed and edited the output as needed and take full responsibility for the content of the published article.

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.

Ardia, D., K. Bluteau, K. Boudt, and K. Inghelbrecht (2023): "Climate Change Concerns and the Performance of Green vs. Brown Stocks," *Management Science*, 69, 7607–7632.

Arteaga-Garavito, M. J., R. Colacito, M. M. Croce, and B. Yang (2023): "International Climate News," Working paper, SSRN.

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

Atkeson, A. and P. J. Kehoe (2007): "Modeling the Transition to a New Economy: Lessons from Two Technological Revolutions," *American Economic Review*, 97, 64–88.

Bai, J. and P. Perron (1998): "Estimating and Testing Linear Models with Multiple Structural Changes," *Econometrica*, 66, 47–78.

——— (2003): "Computation and Analysis of Multiple Structural Change Models," *Journal of Applied Econometrics*, 18, 1–22.

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

Barroso, P. and P. Santa-Clara (2015): "Momentum Has Its Moments," *Journal of Financial Economics*, 116, 111–120.

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

Boehmer, E., J. Masumeci, and A. B. Poulsen (1991): "Event-study methodology under conditions of event-induced variance," *Journal of Financial Economics*, 30, 253–272.

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.

Chen, A. Y. and T. Zimmermann (2022): "Open Source Cross-Sectional Asset Pricing," *Critical Finance Review*, 11, 207–264.

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.

Daniel, K. and T. J. Moskowitz (2016): "Momentum Crashes," *Journal of Financial Economics*, 122, 221–247.

DeMiguel, V., J. Gil-Bazo, and M. Goel (2026): "Legislator Tweets About the Green Transition and the Returns of Green versus Brown Stocks," *Working Paper*.

Edmans, A., D. Levit, and J. Schneemeier (2024): "Socially Responsible Divestment," European Corporate Governance Institute Finance Working Paper No. 823/2022.

Engle, R. F. (2024): "Termination Risk and Sustainability," ECGI Finance Working Paper 1005/2024, European Corporate Governance Institute.

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.

Eskildsen, M., M. Ibert, T. I. Jensen, and L. H. Pedersen (2026): "Realized Returns to Green Investing: A Global Replication," Working paper.

Feng, J., J. Huang, S. Huang, and R. Shi (2025): "Beyond the Numbers: Soft Information in Conditional Asset Pricing," HKU Jockey Club Enterprise Sustainability Global Research Institute Paper No. 2025/022, available at SSRN.

Ferson, W. E. and R. W. Schadt (1996): "Measuring Fund Strategy and Performance in Changing Economic Conditions," *The Journal of Finance*, 51, 425–461.

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.

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.

Haddad, V., S. Kozak, and S. Santosh (2020): "Factor Timing," *The Review of Financial Studies*, 33, 1980–2018.

Hassan, T. A., S. Hollander, L. van Lent, and A. Tahoun (2019): "Firm-Level Political Risk: Measurement and Effects," *The Quarterly Journal of Economics*, 134, 2135–2202.

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.

Kacperczyk, M. T. (2026): "What Is the Carbon Premium a Premium On?" HKU Jockey Club Enterprise Sustainability Global Research Institute Paper No. 2026/007; FEB-RN Research Paper No. 149/2026.

Kogan, L. and D. Papanikolaou (2013): "Firm Characteristics and Stock Returns: The Role of Investment-Specific Shocks," *The Review of Financial Studies*, 26, 2718–2759.

——— (2014): "Growth Opportunities, Technology Shocks, and Asset Prices," *The Journal of Finance*, 69, 675–718.

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

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.

Martinez, J. (2026): "Putty-Clay Automation," CEPR Discussion Paper 16022, CEPR Press, Paris and London.

McFadden, D. (1974): "Conditional Logit Analysis of Qualitative Choice Behavior," in *Frontiers in Econometrics*, ed. by P. Zarembka, New York: Academic Press, 105–142.

Moreira, A. and T. Muir (2017): "Volatility-Managed Portfolios," *The Journal of Finance*, 72, 1611–1644.

Papanikolaou, D. (2011): "Investment Shocks and Asset Prices," *Journal of Political Economy*, 119, 639–685.

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

——— (2022): "Dissecting green returns," *Journal of Financial Economics*, 146, 403–424.

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

Peters, R. H. and L. A. Taylor (2017): "Intangible Capital and the Investment-q Relation," *Journal of Financial Economics*, 123, 251–272.

Ramelli, S., A. F. Wagner, R. J. Zeckhauser, and A. Ziegler (2021): "Investor Rewards to Climate Responsibility: Stock-Price Responses to the Opposite Shocks of the 2016 and 2020 U.S. Elections," *The Review of Corporate Finance Studies*, 10, 748–787.

Romano, J. P. and M. Wolf (2005): "Exact and Approximate Stepdown Methods for Multiple Hypothesis Testing," *Journal of the American Statistical Association*, 100, 94–108.

Sautner, Z., L. van Lent, G. Vilkov, and R. Zhang (2023): "Firm-Level Climate Change Exposure," *The Journal of Finance*, 78, 1449–1498.

Stroebel, J. and J. Wurgler (2021): "What do you think about climate finance?" *Journal of Financial Economics*, 142, 487–498.

Wei, C. (2003): "Energy, the Stock Market, and the Putty-Clay Investment Model," *American Economic Review*, 93, 311–323.

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

A Proofs

This appendix collects proofs omitted from the main text.

A.1 Proof of Proposition 1

Fix a compact dirtiness interval $\Phi = [\underline{\phi}, \bar{\phi}] \subset (0, 1]$. Suppose equilibrium frontier objects $P(\phi)y^*(\phi)$ and $f^*(\phi)$ are continuously differentiable on Φ, with $f^{*\prime}(\phi) \geq m_f > 0$ and both objects uniformly bounded.

Let $\rho \equiv r + \delta$. Since $g_y > 0$ and $\tilde{g} \geq 0$, the boundedness of $P(\phi)y^*(\phi)$ and $f^*(\phi)$ implies that $\pi(a, \phi)$ is uniformly bounded on $[0, \infty) \times \Phi$. Hence,

$$x(a, \phi; \kappa) \equiv \frac{\pi(a, \phi)}{\kappa}$$

converges to zero uniformly as $\kappa \to \infty$. The Taylor expansion of the utilization decision in Equation (3) around $x = 0$ is

$$\frac{1}{1 + e^{-x}} = \frac{1}{2} + \frac{x}{4} + O(x^3).$$

The Taylor expansion of the expected payoff in Equation (4) around $x = 0$ is

$$\log(1 + e^x) = \log 2 + \frac{x}{2} + \frac{x^2}{8} + O(x^4).$$

Substituting $x = \pi(a, \phi)/\kappa$ yields, uniformly on $[0, \infty) \times \Phi$,

$$u(a, \phi) = \frac{1}{2} + \frac{\pi(a, \phi)}{4\kappa} + O(\kappa^{-3}),$$

and

$$\Pi(a, \phi) = \kappa \log 2 + \frac{\pi(a, \phi)}{2} + O(\kappa^{-1}).$$

Substituting the utilization expansion into Equation (10) gives

$$I(a, \phi) = e^{-g_y a} \left[\frac{1}{2(\rho + g_y)} + O(\kappa^{-1}) \right],$$

because $\int_0^\infty e^{-\rho h} e^{-g_y(a+h)} \, dh = \frac{e^{-g_y a}}{\rho + g_y}$. The $O(\kappa^{-1})$ term inside brackets is uniform because every term in $\pi(a + h, \phi) e^{-g_y(a+h)}$ is bounded by a constant times $e^{-g_y a} e^{-g_y h}$. Substituting the log-sum expansion into the value equation in (9) gives

$$v(a, \phi) = \frac{\kappa \log 2}{\rho} + O(1),$$

since $\int_0^\infty e^{-\rho h} dh = \rho^{-1}$. The $\pi(a + h, \phi)/2$ term also enters the value, but it does not grow with κ. Because $\pi(a+h,\phi)$ is bounded, its discounted integral is finite and is included in the $O(1)$ term. Therefore

$$\kappa \mathsf{s}(a,\phi) \to L(a,\phi) \equiv \frac{f^*(\phi)}{2\log 2}\frac{\rho}{\rho+g_y}e^{-g_y a}, \tag{18}$$

uniformly in (a,ϕ). Equivalently,

$$\mathsf{s}(a,\phi) = \frac{f^*(\phi)}{2\kappa\log 2}\frac{\rho}{\rho+g_y}e^{-g_y a} + e^{-g_y a}O(\kappa^{-2}), \tag{19}$$

uniformly on $[0,\infty)\times\Phi$. In this limit, utilization is one half for every machine and value is dominated by the common option-value component $\kappa\log 2/\rho$. Thus age and dirtiness do not affect the leading term of the denominator, but they do affect the numerator. Dirtiness enters through $f^*(\phi)$, and age enters through the remaining fuel-exposure factor $e^{-g_y a}$.

Differentiating the limiting object in (18) gives

$$L_\phi(a,\phi) = \frac{f^{*\prime}(\phi)}{2\log 2}\frac{\rho}{\rho+g_y}e^{-g_y a} > 0,$$

and

$$L_{a\phi}(a,\phi) = -\frac{g_y f^{*\prime}(\phi)}{2\log 2}\frac{\rho}{\rho+g_y}e^{-g_y a} < 0.$$

The inequalities use $f^{*\prime}(\phi) \geq m_f > 0$ and $g_y > 0$. Differentiating (19) yields

$$\partial_\phi \mathsf{s}(a,\phi) = \frac{f^{*\prime}(\phi)}{2\kappa\log 2}\frac{\rho}{\rho+g_y}e^{-g_y a} + e^{-g_y a}O(\kappa^{-2}),$$

and

$$\partial_{a\phi} \mathsf{s}(a,\phi) = -\frac{g_y f^{*\prime}(\phi)}{2\kappa\log 2}\frac{\rho}{\rho+g_y}e^{-g_y a} + e^{-g_y a}O(\kappa^{-2}). \tag{20}$$

The leading terms are bounded away from zero relative to $e^{-g_y a}\kappa^{-1}$ uniformly over Φ. The remainders are $e^{-g_y a}O(\kappa^{-2})$, uniformly in age and dirtiness. Hence, there exists a finite $\bar{\kappa}$ such that, for all $\kappa \geq \bar{\kappa}$, the leading terms dominate the remainders for every $a \geq 0$ and every $\phi \in \Phi$. It follows that $\partial_\phi \mathsf{s}(a,\phi) > 0$ and $\partial_{a\phi}\mathsf{s}(a,\phi) < 0$ on $[0,\infty)\times\Phi$. $\square$

A.2 Proof of Proposition 2

Let $V_i(\tau)$ denote firm i's value after a permanent tax of size τ, and define

$$\mathsf{S}_i \equiv -\frac{1}{V_i(0)}\frac{\partial V_i(\tau)}{\partial \tau}\bigg|_{\tau=0}.$$

Then

$$\log V_i(\tau) = \log V_i(0) - \tau\,\mathsf{S}_i + o(\tau).$$

In partial equilibrium, $\mathsf{S}_i = \mathsf{S}(\phi_i)$ by construction. 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:

$$E[r_i^e] = \lambda\, \xi\, \tau\, \mathsf{S}_i + o(\tau).$$

This is equation (13). □

B Model appendix

B.1 Free entry and machine design

This appendix collects conditions underlying the equilibrium objects used in the main text. At installation, a firm of dirtiness ϕ chooses its frontier capital ratio $k^*(\phi)$ and fuel ratio $f^*(\phi)$. Let

$$A_u^{\tilde{g}}(\phi) = \int_0^\infty e^{-(r+\delta)h} u(h,\phi) e^{-\tilde{g}h}\, dh$$

$$A_u^0(\phi) = \int_0^\infty e^{-(r+\delta)h} u(h,\phi)\, dh,$$

$$A_u^{g_y}(\phi) = \int_0^\infty e^{-(r+\delta)h} u(h,\phi) e^{-g_y h}\, dh.$$

The logit operating problem admits a useful variational representation. Suppressing (a,ϕ) for notational convenience, for any deterministic payoff π the stationary expected profit satisfies

$$\Pi = \max_{u\in[0,1]} \{u\pi + \kappa\nu(u)\}, \tag{21}$$

where

$$\nu(u) \equiv -u\log u - (1-u)\log(1-u) \tag{22}$$

is binary entropy. The first term is expected operating profit. The second term is the option value created by the ability to run or idle the machine after observing the transitory shocks. At that optimum, (21) gives $\Pi = u\pi + \kappa\nu(u)$. Define the discounted option value as

$$N_u(\phi) = \int_0^\infty e^{-(r+\delta)h} \nu(u(h,\phi))\, dh.$$

Using (21) and (9), the value of a frontier (age-zero) machine can be written as

$$v(0,\phi) = P(\phi)y^*(\phi)A_u^{\tilde{g}}(\phi) - p^f f^*(\phi) A_u^{g_y}(\phi) - wA_u^0(\phi) + \kappa N_u(\phi).$$

The first-order conditions use the envelope theorem for the future utilization problem, so derivatives of the optimal utilization rule with respect to the design variables do not enter. The frontier capital ratio satisfies

$$k^*(\phi) = \alpha_k(\phi) P(\phi) y^*(\phi) A_u^{\tilde{g}}(\phi),$$

and free entry implies $v(0,\phi) = k^*(\phi)$. Combining these conditions with the fuel first-order condition yields

$$p^f f^*(\phi) = \alpha_f(\phi)\, P(\phi) y^*(\phi)\, \frac{A_u^{\tilde{g}}(\phi)}{A_u^{g_y}(\phi)}, \tag{23}$$

$$P(\phi) y^*(\phi) = \frac{w A_u^0(\phi) - \kappa N_u(\phi)}{\alpha_\ell A_u^{\tilde{g}}(\phi)}. \tag{24}$$

We work on the parameter region in which $w A_u^0(\phi) - \kappa N_u(\phi) > 0$, so frontier revenue is positive. These two equations jointly determine detrended frontier revenue $P(\phi)y^*(\phi)$ and detrended fuel commitment $f^*(\phi)$.

B.2 General equilibrium

The main text focuses on partial equilibrium to isolate the first-order machine-level stranding mechanism. General equilibrium adds four additional margins: price pass-through, demand reallocation across goods, wage adjustment, and discount-factor changes. We analyze each of these in turn.

A competitive final-good producer aggregates intermediates with CES technology:

$$Y_t = \left(\int_0^1 Y_t(\phi)^{\frac{\sigma-1}{\sigma}}\, d\phi \right)^{\frac{\sigma}{\sigma-1}}. \tag{25}$$

The implied demand system pins down the equilibrium inverse-demand schedule $P_t(\phi)$ used in the main text.

Let $\rho \equiv r + \delta$. Define three present-value kernels:

$$J_R(a,\phi) \equiv \int_0^\infty e^{-\rho h} u(a+h,\phi) P(\phi) y^*(\phi) e^{-\tilde{g}(a+h)}\, dh,$$

$$J_W(a,\phi) \equiv \int_0^\infty e^{-\rho h} u(a+h,\phi)\, dh,$$

$$L_D(a,\phi) \equiv \int_0^\infty h\, e^{-\rho h} \Pi(a+h,\phi)\, dh.$$

These generate the normalized revenue, wage, and discount-duration objects

$$\mathcal{R}(a,\phi) \equiv \frac{J_R(a,\phi)}{v(a,\phi)}, \qquad \mathcal{W}(a,\phi) \equiv \frac{J_W(a,\phi)}{v(a,\phi)}, \qquad \mathcal{D}(a,\phi) \equiv \frac{L_D(a,\phi)}{v(a,\phi)}.$$

The first two are value-weighted cash-flow exposures. The third is a valuation-duration object. Put differently, J_R, J_W, and L_D are raw discounted revenue, wage-base, and duration integrals, while $\mathcal{R}$, $\mathcal{W}$, and $\mathcal{D}$ divide each one by machine value so they can be read as GE exposures per unit of value.

Let $\hat{P}^R(\phi) \equiv \partial_\tau \log\{P(\phi)y^*(\phi)\}|_{\tau=0}$ denote the age-invariant proportional response of frontier operating revenue after the tax. This proportional response summarizes cost pass-through and demand reallocation across dirty and clean goods; for the paper's main result only the net response matters. Let $w_\tau \equiv \partial_\tau w|_{\tau=0}$ denote the common wage response, and let $\varrho_\tau(\phi)$ denote an age-invariant change in the required return for firm ϕ. Under the first-order approximation that general-equilibrium valuation effects can be summarized by a small parallel shift in the firm-specific required return, the general-equilibrium stranding semi-elasticity is

$$\mathsf{s}^{\mathrm{GE}}(a,\phi) = \mathsf{s}(a,\phi) - \hat{P}^R(\phi)\mathcal{R}(a,\phi) + w_\tau \mathcal{W}(a,\phi) + \varrho_\tau(\phi)\mathcal{D}(a,\phi). \tag{26}$$

Equation (26) shows how general equilibrium modifies the partial-equilibrium benchmark. Net revenue responses operate through the revenue kernel, wages operate through the labor-cost kernel, and discount-factor changes operate through duration.

The age loading of each channel is characterized by exact lower-limit derivative formulas. Because

$$\partial_a v(a,\phi) = \rho v(a,\phi) - \Pi(a,\phi),$$

we have

$$\begin{aligned}
\partial_a J_R(a,\phi) &= \rho J_R(a,\phi) - u(a,\phi)P(\phi)y^*(\phi)e^{-\tilde{g}a},\\
\partial_a J_W(a,\phi) &= \rho J_W(a,\phi) - u(a,\phi),\\
\partial_a L_D(a,\phi) &= \rho L_D(a,\phi) - v(a,\phi).
\end{aligned}$$

Therefore

$$\partial_a \mathcal{R}(a,\phi) = \mathcal{R}(a,\phi)\left[\frac{\Pi(a,\phi)}{v(a,\phi)} - \frac{u(a,\phi)P(\phi)y^*(\phi)e^{-\tilde{g}a}}{J_R(a,\phi)}\right], \tag{27}$$

$$\partial_a \mathcal{W}(a,\phi) = \mathcal{W}(a,\phi)\left[\frac{\Pi(a,\phi)}{v(a,\phi)} - \frac{u(a,\phi)}{J_W(a,\phi)}\right], \tag{28}$$

$$\partial_a \mathcal{D}(a,\phi) = \mathcal{D}(a,\phi)\frac{\Pi(a,\phi)}{v(a,\phi)} - 1. \tag{29}$$

These formulas show that each GE channel loads on age through the gap between the machine's current contribution to value and its current contribution to that channel's kernel.

Net revenue response. Starting with the revenue kernel, let the flow payoff per unit of current operating revenue be

$$R(a,\phi) \equiv P(\phi)y^*(\phi)e^{-\tilde{g}a}, \qquad q_R(a,\phi) \equiv \frac{\Pi(a,\phi)}{u(a,\phi)R(a,\phi)}.$$

This object summarizes how much current operating revenue translates into value-relevant flow payoff. High q_R means a given unit of current revenue carries high flow value. Then

$$\frac{v(a,\phi)}{J_R(a,\phi)} = \int_0^\infty \omega_R(h;a,\phi)\, q_R(a+h,\phi)\, dh, \qquad \omega_R(h;a,\phi) \equiv \frac{e^{-\rho h}u(a+h,\phi)R(a+h,\phi)}{J_R(a,\phi)},$$

so $\mathcal{R}(a,\phi) = 1/\mathbb{E}^R_a[q_R]$. Hence the sign of $\partial_a \mathcal{R}(a,\phi)$ is pinned down by the age profile of q_R:

$$\partial_a \mathcal{R}(a,\phi) = \frac{u(a,\phi)R(a,\phi)J_R(a,\phi)}{v(a,\phi)^2}\Big[q_R(a,\phi) - \mathbb{E}^R_a[q_R]\Big].$$

If q_R declines with age, then $\mathcal{R}$ rises with age and revenue channels load more heavily on old machines. If q_R rises with age, the reverse holds and revenue channels are front-loaded toward young machines.

An increasing q_R is the case in which utilization does not fall so sharply with age that old machines lose value too quickly relative to their current revenue base. The option value embedded in $\Pi(a,\phi)$ then keeps the value-to-revenue ratio from declining with age; this is the same benchmark studied in the main text. The revenue response $\hat{P}^R(\phi)$ captures pass-through and demand reallocation across goods. A carbon tax naturally shifts demand away from dirtier goods and toward cleaner goods, so this response may vary with ϕ. Conditional on ϕ, however, machines sell the same good and face the same proportional frontier-revenue response. Age variation in the revenue channel therefore comes from the kernel $\mathcal{R}(a,\phi)$.

Proposition 3 (Revenue channels under the maintained case)**.** *Suppose $\partial_a q_R(a,\phi) \geq 0$ for all relevant (a,ϕ), with strict inequality on a set of positive measure. Then $\partial_a \mathcal{R}(a,\phi) \leq 0$, so the normalized revenue kernel is front-loaded toward young machines. Under the maintained assumption that vintages of a given ϕ sell the same good and therefore face the same proportional frontier-revenue response $\hat{P}^R(\phi)$, the GE revenue channel $G^R(a,\phi) \equiv -\hat{P}^R(\phi)\mathcal{R}(a,\phi)$ inherits that same age loading.*

Define the revenue contribution as

$$G^R(a,\phi) \equiv -\hat{P}^R(\phi)\mathcal{R}(a,\phi).$$

The object $\hat{P}^R(\phi)$ can vary across goods, reflecting pass-through and demand reallocation away from dirty goods after a carbon tax, but it is common across vintages within a given ϕ. Its age derivative is therefore

$$\partial_a G^R(a,\phi) = -\hat{P}^R(\phi)\partial_a \mathcal{R}(a,\phi).$$

Under Proposition 3, all age variation in the revenue channel is carried by $\mathcal{R}(a,\phi)$. The natural case for dirty goods is $\hat{P}^R(\phi) < 0$. After a carbon tax, demand reallocates away from dirtier goods, and this demand effect dominates any increase in producer revenue from pass-through. In that case, $G^R(a,\phi) > 0$, so the GE revenue channel raises stranding, and it does so most strongly for young machines.

Wage adjustment. The wage channel is governed by the value-weighted labor-cost kernel rather than by revenue. Define the flow payoff per unit of operating time,

$$q_W(a,\phi) \equiv \frac{\Pi(a,\phi)}{u(a,\phi)}.$$

Then

$$\frac{v(a,\phi)}{J_W(a,\phi)} = \int_0^\infty \omega_W(h;a,\phi)\, q_W(a+h,\phi)\, dh, \qquad \omega_W(h;a,\phi) \equiv \frac{e^{-\rho h} u(a+h,\phi)}{J_W(a,\phi)},$$

so that $\mathcal{W}(a,\phi) = 1/\mathbb{E}_a^W[q_W]$ is the inverse of a discounted-utilization-weighted average of q_W. Moreover,

$$\frac{\partial q_W}{\partial \pi} = 1 - e^{-\pi/\kappa} \log\!\Big(1 + e^{\pi/\kappa}\Big) > 0,$$

because $\log(1+x) < x$ for every $x > 0$. Thus q_W is strictly increasing in operating profit.

Proposition 4 (Wage-channel monotonicity)**.** *Suppose the tax lowers wages, $w_\tau < 0$, and operating profit is weakly decreasing with age, $\partial_a \pi(a,\phi) \le 0$. Then the wage kernel is weakly increasing with age:*

$$\partial_a \mathcal{W}(a,\phi) \ge 0.$$

Equivalently, the wage-response contribution $G^W(a,\phi) \equiv w_\tau \mathcal{W}(a,\phi)$ is weakly decreasing with age, so wage adjustment attenuates stranding more for older machines than for younger machines.

Proof. Since q_W is increasing in π and π is weakly decreasing in age, $q_W(a,\phi)$ is weakly decreasing in age. Therefore the current value $q_W(a,\phi)$ is weakly larger than its future weighted average $\mathbb{E}_a^W[q_W]$. Using

$$\partial_a \mathcal{W}(a,\phi) = \frac{u(a,\phi) J_W(a,\phi)}{v(a,\phi)^2} \Big[q_W(a,\phi) - \mathbb{E}_a^W[q_W] \Big],$$

the result follows. □

Define the wage contribution as

$$G^W(a,\phi) \equiv w_\tau \mathcal{W}(a,\phi).$$

Proposition 4 gives a clean age result. When $w_\tau < 0$, a wage decline cushions old machines more than young machines, and the wage channel reinforces the paper's negative emissions-age interaction.

Discount-factor changes. Define the discount-factor contribution as

$$G^D(a,\phi) \equiv \varrho_\tau(\phi) \mathcal{D}(a,\phi).$$

This term is conceptually different from the three cash-flow channels above. It does not load on remaining fuel use, but rather on valuation duration. Equation (29) implies

$$\partial_a G^D(a,\phi) = \varrho_\tau(\phi) \left[\mathcal{D}(a,\phi) \frac{\Pi(a,\phi)}{v(a,\phi)} - 1 \right].$$

We can say more under a standard representative-household benchmark. Suppose

$$M_h = e^{-\rho h}\left(\frac{C_{t+h}}{C_t}\right)^{-\gamma},$$

so the tax changes horizon-h discount factors according to

$$\partial_\tau \log M_h\Big|_{\tau=0} = -\gamma\left[\partial_\tau \log C_{t+h} - \partial_\tau \log C_t\right]_{\tau=0}.$$

This calculation describes the post-shock term-structure effect on machine values; the sign of the pre-arrival risk premium is governed separately by the covariance of the policy jump with the stochastic discount factor. This special case is useful for age-loading intuition. If consumption falls on impact and then recovers, in the sense that $\partial_\tau \log C_t < 0$ and $\partial_\tau \log C_{t+h}$ is weakly increasing in h, then $\partial_\tau \log M_h \leq 0$ and becomes weakly more negative with horizon. The discount-factor channel therefore lowers the value of back-loaded cash flows disproportionately. Hence, whenever younger machines are longer duration than older ones, the duration channel raises stranding more for young machines than for old ones. This is the natural age implication of the transition path in Gilchrist et al. (2026), where front-loaded consumption losses imply an adverse repricing of duration. What remains model-dependent is the dirtiness dimension. A common discount-factor shock strengthens the paper's negative emissions-age interaction only if dirty firms also have larger duration exposure or experience a larger required-return response.

B.3 Alternative policy instruments

The paper's baseline mechanism is specific to policies that increase the operating cost of installed dirty capital. Other climate policies work through different cash-flow margins and therefore need not generate the same emissions-age interaction. This appendix discusses a broader policy taxonomy for the model.

Carbon taxes and emissions-fee regulation. The baseline policy changes machine cash flow by

$$d\pi(a,\phi) = -d\tau\, f^*(\phi)e^{-g_y a},$$

plus the general-equilibrium feedbacks discussed in Appendix B.2. This expression applies to carbon taxes, emissions fees, cap-and-trade systems with a marginal emissions price, and regulations whose compliance cost is approximately proportional to fuel use or emissions. This shock loads directly on remaining fuel-tax exposure. It therefore gives the cleanest mapping from policy exposure to the model's present value of fuel-exposure and, under the conditions highlighted in the main text, generates the baseline negative emissions-age interaction.

Green technology shocks and clean subsidies. Green technology improvements and clean investment subsidies do not tax installed dirty fuel use directly. Instead they work primarily through relative output prices and, in general equilibrium, wages and productivity growth. Dirty incumbent machines typically lose because clean alternatives become more competitive, but the shock is not proportional to remaining fuel exposure. The resulting emissions-age interaction is therefore generally weaker and more model-dependent than in the baseline carbon-tax mechanism.

Dirty investment taxes and entry restrictions. Policies that tax new dirty investment or restrict future dirty entry, while leaving the operating costs of installed dirty capital largely unchanged, reduce competition for existing dirty machines. Their primary effect is to raise future dirty-good prices. Existing dirty incumbents can therefore gain value even though future dirty capital is discouraged. When gains are largest for young incumbents with the most future output left to sell, the emissions-age interaction can reverse sign relative to the baseline stranding mechanism.

ESG or discount-rate wedges. A dirtiness-dependent discount-rate wedge changes valuation rather than operating costs. It can lower the value of dirty incumbents directly through discounting, but it can also raise incumbent rents by discouraging future dirty entry, as in responsible-investing/divestment channels that work through capital allocation (Edmans et al., 2024). Because the wedge does not load mechanically on remaining fuel exposure, it does not naturally produce the paper's baseline fuel-exposure interaction. It can nevertheless generate an age pattern if young dirty firms have longer-duration cash flows; that pattern is duration-driven rather than fuel-exposure-driven.

B.4 Age-invariant-profit benchmark

In this section we show why a vintage capital model is essential to our findings. Consider the model of the main text under the benchmark in which machine payoffs are age-invariant. Suppose that for every firm ϕ,

$$\pi(a, \phi) = \pi(\phi), \qquad f(a, \phi) = f^*(\phi),$$

so that both operating profits and fuel use are independent of age. Then utilization and the log-sum flow payoff are also age-invariant:

$$u(a, \phi) = u(\phi), \qquad \Pi(a, \phi) = \Pi(\phi).$$

Machine-value and fuel exposure are now

$$v(a, \phi) = \int_0^\infty e^{-(r+\delta)h}\Pi(\phi)\, dh = \frac{\Pi(\phi)}{r+\delta},$$

$$I^{\mathrm{AI}}(a, \phi) = \int_0^\infty e^{-(r+\delta)h}u(\phi)\, dh = \frac{u(\phi)}{r+\delta},$$

where the benchmark analogue of fuel exposure drops the factor $e^{-g_y(a+h)}$ because fuel use is itself assumed age-invariant. The partial-equilibrium stranding ratio then becomes

$$\mathsf{s}^{\mathrm{AI}}(a, \phi) = \frac{f^*(\phi)\, I^{\mathrm{AI}}(a, \phi)}{v(a, \phi)} = \frac{f^*(\phi)u(\phi)}{\Pi(\phi)}. \tag{30}$$

Equation (30) is the age-invariant counterpart of the main-text stranding object. It depends on dirtiness, but it is completely independent of age. Hence

$$\partial_a \mathsf{s}^{\mathrm{AI}}(a,\phi) = 0, \qquad \partial_{a\phi} \mathsf{s}^{\mathrm{AI}}(a,\phi) = 0.$$

The paper's emissions-age interaction therefore requires vintage dependence in either operating cash flows, fuel exposure, or both.

C Model calibration

Proposition 1 requires $\kappa \geq \bar{\kappa}$ for the cross-derivative signs to hold. This appendix asks whether U.S. fossil-power generation data place κ in that region. We estimate κ from the age slope of plant utilization in FERC-regulated coal and combined-cycle gas plants over 2001–2022, compute $\bar{\kappa}$ numerically over the empirically relevant range of fuel shares and ages, and find $\hat{\kappa}/\bar{\kappa} = 1.74$ with 95% CI $[1.08, 2.40]$. The lower bound exceeds one, placing the U.S. fossil-power sector inside the Proposition 1 region.

C.1 Estimating equation

Under the model's logit utilization rule, $\log[u/(1-u)] = \pi(a,\phi)/\kappa$. Differentiating in age at $a = 0$,

$$\left.\frac{d \log[u(a,\phi)/(1-u(a,\phi))]}{da}\right|_{a=0} = \frac{\pi'(0)}{\kappa}, \tag{31}$$

with

$$\pi'(0) = \frac{Py^\star}{w}\Big[g_y \alpha_f - \tilde{g}\Big], \tag{32}$$

where the second equation uses the fuel first-order condition. The age slope of logit utilization is proportional to $1/\kappa$, with the proportionality constant $\pi'(0,\phi)$ pinned down by parameters g_y and $\tilde{g}$ and the empirically observed values of $Py^\star/w$ and α_f.

We work in wage-normalized units throughout the appendix. The model's utilization rule is unchanged if we scale both the operating payoff $\pi(a,\phi)$ and the shock dispersion κ by the same positive constant, so we set the wage to one and write π and κ for the wage-normalized objects.[15] This matches the FERC plant-level data, which identify revenue and fuel costs relative to nonfuel operating costs, which we interpret as the empirical counterpart of the model wage bill.

The empirical counterpart of (31) is a capacity-weighted OLS regression

$$u_{it} = \beta a_{it} + \alpha_g + \delta_{r\times t} + \varepsilon_{it}, \tag{33}$$

where u_{it} is the annual operating fraction of plant i in year t, a_{it} is plant age, α_g is a technology fixed effect, and $\delta_{r\times t}$ is a region-by-year fixed effect. Plant-year observations are

[15] Explicitly, with wage normalization $\pi^w(a,\phi) = (Py^\star/w)\, e^{-\tilde{g}a} - (p^f f^\star/w)\, e^{-g_y a} - 1$ and $\kappa^w = \kappa/w$. Since $\pi(a,\phi)/\kappa = \pi^w(a,\phi)/\kappa^w$, the utilization rule is invariant.

weighted by capacity and standard errors are clustered by plant.

The OLS slope $\hat{\beta}$ estimates du/da, but equation (31) involves the logit slope. By the chain rule,

$$\left.\frac{d\log[u/(1-u)]}{da}\right|_{u=\bar{u}} = \frac{\hat{\beta}}{\bar{u}(1-\bar{u})}, \tag{34}$$

evaluated at the sample-mean utilization $\bar{u}$.

C.2 Data

We merge EPA emissions-monitor records (operating hours), FERC Form 1 accounts (fuel costs, nonfuel operating expenses), and EIA Form 860 (vintage, technology, regional market) at the plant-year level over 2001–2022. The sample is FERC-regulated coal and combined-cycle gas plants with $u_{it} \in (0.01, 0.99)$ and at least six months of operating data per plant-year. After deduplication, 3,293 plant-years across 255 plants remain.

C.3 Empirical moments

Two empirical ratios, computed in the 2015–2019 window, convert the estimated slope into model units.

Revenue-to-wage ratio. $Py^\star/w$ equals the capacity-weighted median of (output price / nonfuel cost per unit output). The model treats nonfuel cost as the wage bill, so this is the data counterpart of revenue per machine relative to the wage cost.

Fuel share. α_f equals the capacity-weighted median of (fuel cost per unit output / output price). The model's fuel first-order condition makes this ratio the Cobb–Douglas fuel share at the active machine.

Our use of the median is a deliberate choice. The plant-year ratio c^f_{it}/P_{it} has a heavy right tail because assigned output prices are noisy proxies for the model's revenue object. The resulting data moments are $Py^\star/w = 5.42$ and $\alpha_f = 0.642$.

C.4 Slope estimates, $\hat{\kappa}$, and the structural threshold $\bar{\kappa}$

From Equations (31) and (34) the conversion from slope to $\hat{\kappa}$ is

$$\hat{\kappa} = \frac{|\pi'(0)|}{|\hat{\beta}|/[\bar{u}(1-\bar{u})]}, \tag{35}$$

with $\pi'(0)$ given by (32), using the fuel-decay rate $g_y = 0.02$, the embodied-productivity decay rate $\tilde{g} = 0.0185$, and the sample-mean utilization $\bar{u} = 0.739$ (sources in Table C2). The standard error on $\hat{\kappa}$ follows from the delta method, $\mathrm{SE}(\hat{\kappa}) = |\pi'(0)\bar{u}(1-\bar{u})/\hat{\beta}^2| \cdot \mathrm{SE}(\hat{\beta})$.

We define $\bar{\kappa}$ as the smallest κ at which Proposition 1's two sign conditions, $\partial_\phi \mathsf{s} > 0$ and $\partial^2_{a\phi}\mathsf{s} < 0$, hold for all (a, ϕ) in the rectangle $\alpha_f \in [0.25, 0.85]$ and $a \in [0, 50]$ years. We compute it numerically by bisection. The rectangle is set by the data on the fuel-share side and by the structural ceiling on the upper end. The lower endpoint is the lower tail of the FERC fossil-plant sample, the upper endpoint is $1 - \alpha_\ell$ under constant returns, and the age range covers the lifetime of fossil plants.

Table C1: Calibration of κ and the structural threshold $\bar{\kappa}$

Range	$\hat{\beta}$ (u/yr)	SE	$\hat{\kappa}$	95% CI ($\hat{\kappa}$)	$\bar{\kappa}$	$\hat{\kappa}/\bar{\kappa}$	95% CI ($\hat{\kappa}/\bar{\kappa}$)
$\alpha_f \in [0.25, 0.85]$	-0.0052	0.0010	1.13	$[0.70, 1.56]$	0.65	1.74	$[1.08, 2.40]$

Notes. $n = 3{,}293$ plant-years across 255 FERC-regulated plants over 2001–2022. $\hat{\beta}$ is the capacity-weighted OLS slope of u_{it} on plant age, after absorbing technology and year-by-region fixed effects. $\hat{\kappa}$ converts from slope to wage-normalized model units via equation (35) at anchors $Py^\star/w = 5.42$ and $\alpha_f = 0.642$, which are 2015–2019 capacity-weighted medians. $\bar{\kappa}$ is located by geometric bisection in the same wage-normalized units. Confidence intervals on $\hat{\kappa}$ and on the ratio are computed by delta method on $\hat{\beta}$.

The headline calibration gives $\hat{\kappa}/\bar{\kappa} = 1.74$ with a 95% confidence interval of $[1.08, 2.40]$. The lower bound exceeds one, placing the calibrated economy inside the Proposition 1 region on the range of fuel shares observed in the FERC sample.

Robustness to embodied-growth assumption. The decay rate $\tilde{g}$ governs how quickly older machines fall behind the productivity frontier. Larger $\tilde{g}$ widens the gap between vintages and lowers $\bar{\kappa}$. Setting $\tilde{g}$ as if all productivity growth were disembodied, the conservative case given the embodied-growth evidence in Greenwood et al. (1997), yields $\hat{\kappa}/\bar{\kappa} = 1.59$, with 95% confidence interval $[0.99, 2.19]$.

C.5 Other calibration parameters

Table C2 summarizes the full parameter set used in Figure 1. The headline calibration uses pooled coal and combined-cycle gas anchors from the 2015–2019 accounting window with the slope estimated on the full 2001–2022 panel (Table C1).

Machine life (δ). In the stochastic-death model, $1/\delta$ is the expected machine lifetime. EIA Form 860 reports operating and retirement dates for every generator. Among fossil-fuel generators with capacity above 20 MW that retired during 2005–2024, the mean lifetime is 54 years for coal and 48 years for gas, implying $\delta \approx 0.019$–0.021. We round to $\delta = 0.02$.

Labor share (α_ℓ). The labor share governs the split between factor costs that scale with output (fuel and capital) and costs that are fixed in detrended terms (wages). From FERC Form 1, nonfuel operating expenses (labor, maintenance, administrative costs) average \$13/MWh for coal and \$15/MWh for gas, while revenue is approximately \$30–65/MWh depending on the year. The ratio of nonfuel opex to revenue ranges from 0.15 to 0.30. The calibration uses $\alpha_\ell = 0.15$.

Table C2: Calibration parameters for Figure 1

Parameter	Value	Description	Source
κ	1.13	Empirical utilization shock dispersion	Age slope of utilization, FERC-regulated coal and combined-cycle gas, 2001–2022 panel (Table C1)
$Py^{\star}/w$	5.42	Revenue-to-wage ratio	Capacity-weighted median, 2015–2019 accounting window
$\alpha_{f,\mathrm{hi}}$	0.85	Upper end of $\alpha_f(\phi)$	$1-\alpha_\ell$, the structural ceiling under constant returns
δ	0.02	Machine death rate	EIA retirement data, mean life 50 yrs
α_ℓ	0.15	Labor share	Lower end of FERC nonfuel opex / revenue range
α_f	0.25, 0.65, 0.85	Fuel-share grid for figure	5th percentile, capacity-weighted median, and structural ceiling $1-\alpha_\ell$
g_y	0.02	Output growth rate	U.S. real GDP per capita growth
r	0.05	Discount rate	Standard

Notes. Parameters are anchored to EIA, FERC, and EPA operating records for U.S. fossil-fuel plants. Details for each parameter appear below.

Fuel share (α_f). The fuel share in the Cobb–Douglas production function maps to the ratio of fuel cost to revenue through the first-order condition for machine design. Capacity-weighted on the 2015–2019 accounting window pooled across coal and combined-cycle gas plants, the median is $\alpha_f = 0.642$. The threshold computation uses $\alpha_f \in [0.25, 0.85]$. The lower endpoint is the lower tail of the fossil-fuel sample. The upper endpoint is $\alpha_{f,\mathrm{hi}} = 0.85 = 1-\alpha_\ell$, the structural ceiling under constant returns. Figure 1 shows three slices $\alpha_f \in \{0.25, 0.65, 0.85\}$. The lower value 0.25 is the capacity-weighted 5th percentile of empirical fuel shares in the calibration sample. The middle value 0.65 is the rounded capacity-weighted median ($\alpha_f = 0.642$ in the headline calibration). The upper value 0.85 is the structural ceiling $1-\alpha_\ell$, since the calibration sets $\alpha_\ell = 0.15$.

D Utility-sector event study: data, events, and robustness

This appendix describes the data, event selection, methodology, and robustness checks for the utility-sector event study in Section 3.

D.1 Sample

The sample consists of 48 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 2005–2025. Fama–French three-factor returns are from Ken French's data library via WRDS.

D.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 two technology groups:

- **Fossil:** coal, natural gas, petroleum, and other fossil-fuel generators.
- **Renewable:** solar (photovoltaic and thermal) and wind generators. We exclude 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 $k \in \{\text{fossil}, \text{renewable}\}$ 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.

D.3 Orthogonalization

The portfolio event study in Section 3 uses an orthogonalized fossil-age and renewable-age measure to isolate fuel-specific variation from total-fleet-age variation. The cross-sectional regression in Section 3.3 handles the same confound by including total fleet age as a control rather than via per-event projection. This subsection describes the orthogonalization that supports the portfolio sort.

A firm with an old overall fleet may 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 renewable) 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 0.77 with raw fossil age, showing that the residualized and raw fossil-age measures remain strongly positively related even though the orthogonalization can change individual rank orderings.

D.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 48 firms and tercile sorts, the Old and Young portfolios each contain approximately 15 firms per event.

D.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 $[-5, -1]$:

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

Event-level covariance matrix SEs. We estimate the full $T \times T$ covariance matrix of event-level abnormal returns across event days:

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

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. The calculation uses variation across events to estimate covariance across event days. It is therefore an event-level covariance adjustment for the plotted CAAR path, not a firm-level cross-sectional dependence correction.

BMP standardization. As a diagnostic for event-level volatility, we compute BMP-standardized test statistics (Boehmer et al., 1991). 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 event-level 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.

D.6 Event selection

The taxonomy of stranding mechanisms (Appendix B.3) suggests that cost-side policies, those that directly raise the operating cost of existing dirty capital, most naturally generate the strong age $\times$ dirty interaction highlighted in the paper. This distinction motivates a baseline sample restricted to events that directly govern or materially affect EPA authority over power plant emissions.

The universe of candidate events is defined *ex ante*. Table D1 lists all EPA regulatory actions and Supreme Court proceedings bearing on power plant emissions or EPA authority over them during our sample period. Events included in the baseline sample are marked with $\bullet$; 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 the set that remains after these exclusions.

D.7 Portfolio robustness

Table D2 reports the portfolio robustness checks referenced in the main text. The table varies the weighting scheme, replaces tercile sorts with median splits, reports BMP-standardized inference, and summarizes leave-one-event-out estimates.

D.8 Factor-model-adjusted CAR regressions

Table D3 replicates the cross-sectional CAR regressions of Table 1, in the same linear specification with ISO capacity-share controls, 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 is positive in every CAR window under both factor models, with point estimates close to the headline industry-adjusted regression and t-statistics between 1.47 and 2.30.

Table D4 splits the main linear specification with ISO controls by event direction. The fossil vintage coefficient is strongest in tightening events, where it is positive in all windows and has t-statistics of 2.40 and 2.20 for CAR$(0, 5)$ and CAR$(0, 10)$. In loosening events, the signed coefficient is also positive but weaker, with a t-statistic of 1.49 in the longest window. The split is not designed to impose equal-and-opposite effects across policy directions; it asks whether the signed headline result is driven entirely by deregulatory relief events or is also

Table D1: 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.

Table D2: Utility portfolio robustness

Specification	Events	Fossil vintage	Solar/wind vintage
Baseline: VW terciles	19	+0.73	+0.16
		[+3.28]	[+0.40]
Equal-weighted terciles	19	+0.85	−0.41
		[+3.20]	[−1.40]
VW median split	19	+0.51	+0.30
		[+3.24]	[+0.86]
Baseline with BMP t-statistic	19	+0.73	+0.16
		[+3.10]	[−0.08]
Leave-one-event-out range	18	+0.64–+0.83	−0.10–+0.40
		[+2.95, +3.96]	[−0.32, +1.17]

Notes: The table reports signed Old−Young cumulative abnormal return spreads at day +10, in percentage points, using the same utility event calendar and event-level covariance matrix inference as Figure 2. t-statistics are reported in square brackets beneath point estimates. The BMP row reports the same raw CAAR with a BMP-standardized t-statistic based on pre-event volatility over days −120 to −21. The leave-one-event-out row reports the range obtained by dropping one baseline event at a time.

Table D3: Cross-sectional regressions of post-event factor-adjusted CAR on generator age

	Fossil-fuel generator age			Renewable generator age		
	(0, 5)	(0, 10)	(0, 15)	(0, 5)	(0, 10)	(0, 15)
Panel A. Fama–French three-factor-adjusted CAR						
Generator age	+0.0356	+0.0409	+0.0405	+0.0359	−0.0239	+0.0599
	[2.38]	[2.06]	[1.72]	[0.91]	[−0.49]	[0.63]
N	744	744	744	445	445	445
R^2	0.558	0.494	0.410	0.636	0.645	0.507
Panel B. Fama–French five-factor-adjusted CAR						
Generator age	+0.0355	+0.0372	+0.0377	+0.0392	−0.0293	+0.0670
	[2.38]	[1.86]	[1.55]	[1.02]	[−0.58]	[0.68]
N	744	744	744	445	445	445
R^2	0.576	0.488	0.381	0.650	0.634	0.478

Notes: Sample: 19 regulatory events (see Appendix D). Panel A reports Fama-French three-factor-adjusted signed CARs; Panel B reports Fama-French five-factor-adjusted signed CARs. Dependent variables are in percentage points. All regressions include event fixed effects. Controls: total fleet age, fossil generation share, log generation capacity, log market capitalization, book-to-market ratio, leverage and ISO shares. ISO shares are the firm's capacity-weighted share in each ISO/RTO bucket. t-statistics in square brackets use heteroskedasticity-consistent (HC3) standard errors.

present when regulatory costs rise. The sharper signal comes from tightening events, consistent with news raising expected fossil operating costs providing the clearest event-time test of stranding exposure. The solar-and-wind coefficients remain statistically indistinguishable from zero in both subsamples.

Table D4: Cross-sectional regressions of post-event CAR on generator age by regulatory direction

	Fossil-fuel generator age			Renewable generator age		
	$(0,5)$	$(0,10)$	$(0,15)$	$(0,5)$	$(0,10)$	$(0,15)$
Panel A. Tightening events						
Generator age	+0.0546	+0.0719	+0.0641	+0.0097	−0.0597	+0.0486
	[2.40]	[2.20]	[1.77]	[0.15]	[−0.57]	[0.42]
Events	9	9	9	9	9	9
N	359	359	359	181	181	181
R^2	0.158	0.152	0.113	0.296	0.388	0.332
Panel B. Loosening events						
Generator age	+0.0059	+0.0157	+0.0493	+0.0419	+0.0237	+0.1336
	[0.30]	[0.58]	[1.49]	[0.71]	[0.43]	[0.80]
Events	10	10	10	10	10	10
N	385	385	385	264	264	264
R^2	0.181	0.198	0.243	0.211	0.263	0.347

Notes: Panel A restricts the sample to tightening events ($d_e = +1$); Panel B restricts the sample to loosening events ($d_e = -1$). Dependent variable: industry-adjusted signed CAR, in percentage points. All regressions include event fixed effects. Controls: total fleet age, fossil generation share, log generation capacity, log market capitalization, book-to-market ratio, leverage and ISO shares. ISO shares are the firm's capacity-weighted share in each ISO/RTO bucket. t-statistics in square brackets use heteroskedasticity-consistent (HC3) standard errors.

E Cross-sectional regressions: data, measurement, and robustness

Variable winsorization and the ln(PP&E) sample. We winsorize firm characteristics following Bolton and Kacperczyk (2021): book-to-market, leverage, investment-to-assets, and ROE at the 2.5th and 97.5th percentiles; momentum, return volatility, sales growth, and EPS growth at the 0.5th and 99.5th percentiles. In the headline cross-section regressions of Table 3, ln(PP&E) is dropped from the control set to maintain a common sample across capital-intensive and capital-light industries: it is missing for roughly 22% of capital-light firms but less than 1% of capital-intensive firms. Including ln(PP&E), as in Table 4 column 3, yields quantitatively similar estimates on a slightly smaller capital-intensive sample.

E.1 Capital age and capital intensity classification

The balance sheet-based measure of capital age follows Lin et al. (2020) and Geelen et al. (2024). Each year, the firm's capital stock consists of surviving old capital (aged by one year and depleted by depreciation) plus new gross investment (entering with age one). The average age of the stock is the weighted average of these two components, with weights proportional to their shares of the current capital stock. Formally:

$$\bar{a}_{i,t} = \underbrace{\frac{(1-\delta_{i,t})\,K_{i,t-1}}{K_{i,t}}}_{\text{surviving share}} \cdot (\bar{a}_{i,t-1}+1) \; + \; \underbrace{\frac{I^{\text{gross}}_{i,t}}{K_{i,t}}}_{\text{new-investment share}} \cdot 1, \tag{36}$$

where the capital stock evolves as $K_{i,t} = (1-\delta_{i,t})\,K_{i,t-1} + I^{\text{gross}}_{i,t}$. The firm-specific depreciation rate is $\delta_{i,t} = (\texttt{dp}_{i,t} - \texttt{am}_{i,t})/\texttt{ppent}_{i,t-1}$, where `dp` is total depreciation and amortization, `am` is amortization of intangibles, and `ppent` is net property, plant, and equipment. Subtracting intangible amortization isolates depreciation of tangible capital, the structural object in the model. Gross investment is recovered from the balance-sheet identity $I^{\text{gross}}_{i,t} = \texttt{ppent}_{i,t} - \texttt{ppent}_{i,t-1} + (\texttt{dp}_{i,t} - \texttt{am}_{i,t})$. The recursion is initialized at each firm's first Compustat observation using $\texttt{dpact}/(\texttt{dp} - \texttt{am})$. The ratio provides an estimate of the initial average age of the tangible-capital stock under straight-line depreciation. In the return regressions we winsorize $\bar{a}_{i,t}$ at the 1st and 99th percentiles of the firm-month distribution to limit the influence of extreme observations.

A firm that invests heavily pulls its average age toward one, whereas a firm with no new investment ages mechanically at one year per year. In the broad-market data this is necessarily a measure of total tangible-capital age, not fossil-capital age. A dirty firm undertaking large new clean or non-emitting investments can therefore appear young even if its legacy dirty assets are old. This limitation makes the utility-sector evidence, where fossil-generator age is observed directly, an important validation and also motivates the investment-channel robustness checks below.

Capital-intensity classification. Appendix Table E1 lists the highest- and lowest-emission SIC-2 industries separately within the capital-intensive and capital-light buckets.

E.1.1 Validation: recursive capital age vs. generator age

Figure E1 provides a validation check for the recursive capital age measure against physical asset age data from the EIA. For publicly traded electric utilities with matched EIA generator data, we plot year-month-demeaned recursive capital age (computed from Compustat PP&E using equation 36) against year-month-demeaned capacity-weighted generator age (computed from EIA Form 860 operating dates). To avoid treating M&A-related PP&E step-ups as organic investment, we apply the same persistent acquisition-contamination screen used in Table E7, excluding firm-years with $\psi^{A}_{i,t} > 30\%$.

The binscatter shows a positive cross-sectional relationship ($\beta = 1.38$, $r = 0.27$, $p < 0.001$): utilities with older capital stocks by the accounting measure also operate physically older generating fleets. The correlation is moderate, as expected. The accounting measure aggregates all tangible capital, including transmission, distribution, and other non-generation assets, while the EIA measure captures the age of generating equipment. The slope indicates that a one-year increase in accounting capital age corresponds to roughly a 1.4-year increase in physical generator age in the screened utility subsample. We therefore read the figure as a noisy but informative validation check that recursive capital age contains information about

Table E1: High- and Low-Emission Industries by Capital Intensity

Rank	SIC-2	Industry	Mean Scope 1 emissions (kt CO_2e)	Median Scope 1 emissions (kt CO_2e)	Mean Gross PP&E / assets	Median Gross PP&E / assets
Capital-intensive: top 10 average firm emissions						
1	29	Petroleum refining	19,278.8	2,759.3	1.126	1.160
2	49	Electric, gas, and sanitary services	11,380.0	2,521.4	1.087	1.114
3	45	Air transportation	10,043.2	2,773.9	0.929	0.982
4	40	Railroad transportation	4,763.5	4,247.4	1.173	1.167
5	33	Primary metals	3,458.9	340.5	0.712	0.623
6	12	Coal mining	3,203.1	2,705.0	1.044	1.024
7	26	Paper and allied products	1,633.0	721.2	0.889	0.926
8	32	Stone, clay, and glass	1,435.6	123.8	0.824	0.845
9	13	Oil and gas extraction	1,431.8	257.1	1.764	1.390
10	10	Metal mining	1,223.3	196.9	0.834	0.867
Capital-intensive: bottom 10 average firm emissions						
1	78	Motion pictures	10.4	3.1	0.610	0.776
2	65	Real estate	23.4	12.9	0.639	0.682
3	56	Apparel stores	30.0	16.7	0.541	0.504
4	27	Printing and publishing	34.0	11.7	0.495	0.463
5	25	Furniture and fixtures	40.2	23.8	0.562	0.608
6	79	Amusement and recreation	53.2	13.9	0.897	0.922
7	72	Personal services	57.2	21.3	0.483	0.368
8	58	Eating and drinking places	61.1	18.3	0.958	1.019
9	16	Heavy construction	63.7	27.6	0.646	0.596
10	80	Health services	65.2	20.3	0.571	0.552
Capital-light: top 10 average firm emissions						
1	50	Wholesale trade: durable goods	313.5	56.0	0.337	0.292
2	51	Wholesale trade: nondurable goods	296.1	124.0	0.479	0.449
3	59	Miscellaneous retail	291.7	14.5	0.316	0.265
4	55	Automotive dealers	204.9	141.6	0.320	0.287
5	01	Agricultural production: crops	178.2	103.2	0.712	0.691
6	17	Special trade contractors	171.8	73.5	0.418	0.413
7	21	Tobacco products	149.7	121.3	0.224	0.187
8	34	Fabricated metals	96.6	26.5	0.566	0.561
9	15	Building construction	86.0	57.8	0.187	0.087
10	57	Furniture stores	76.5	27.6	0.366	0.404
Capital-light: bottom 10 average firm emissions						
1	60	Depository institutions	1.9	0.2		
2	62	Security and commodity brokers	2.9	0.5	0.139	0.054
3	61	Nondepository credit institutions	3.4	1.0	0.073	0.029
4	64	Insurance agents and brokers	3.5	1.2	0.231	0.182
5	67	Holding and investment offices	6.3	1.9	0.076	0.000
6	73	Business services	8.8	2.2	0.357	0.245
7	63	Insurance carriers	10.9	1.4	0.058	0.034
8	87	Engineering and management services	17.9	5.8	0.406	0.361
9	82	Educational services	18.1	12.5	0.357	0.280
10	39	Miscellaneous manufacturing	20.3	8.3	0.341	0.298

Notes: The table ranks SIC-2 industries by average firm-year Scope 1 emissions over the post-Paris sample, January 2016 to December 2025. The unit is thousand metric tons of CO_2e. The final two columns report the mean and median gross PP&E / total assets ratio for the SIC-2 industry, measured in Compustat firm-year observations over 1985–1995. Rankings are computed separately within the patched capital-intensive and capital-light SIC-2 buckets. The capital-intensity split is the fixed SIC-2 classification that starts from this historical gross PP&E / total assets ratio and is adjusted to balance firm-month counts across the two groups. The ranking sample requires positive Scope 1 emissions, excludes nonclassifiable SIC-2 99, requires at least 5 firms, and at least 20 firm-years in an SIC-2 industry. The underlying annual firm-year panel contains 23,821 firm-years, 4,011 firms, and 69 SIC-2 industries after requiring non-missing positive Scope 1 emissions.

the vintage composition of physical assets, the key object in the putty-clay model.

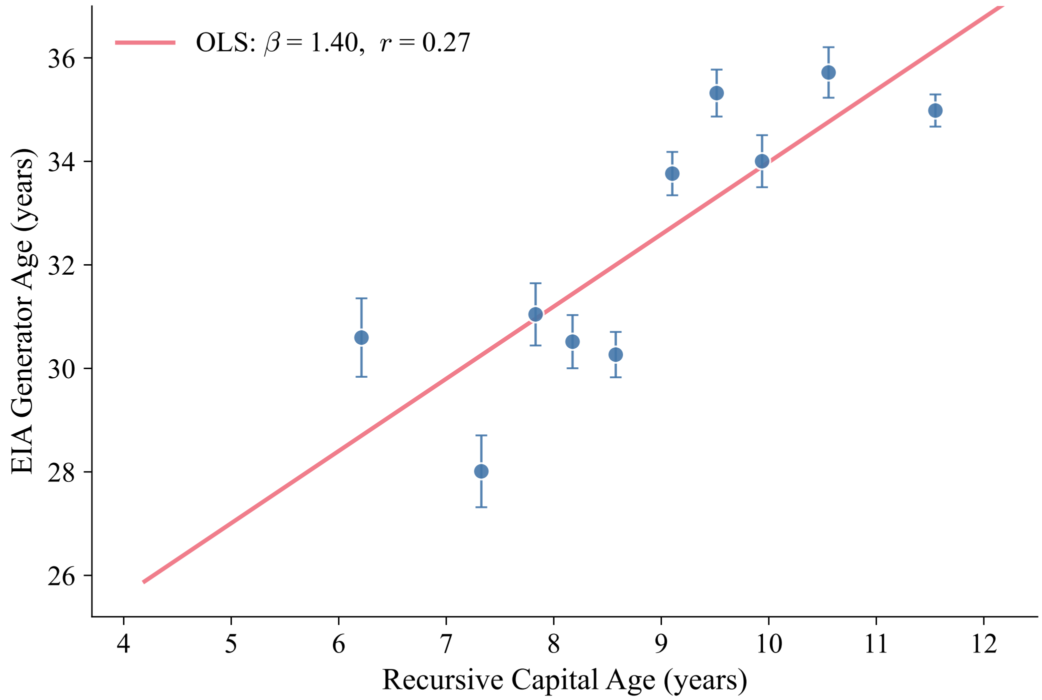

Figure E1: Recursive capital age vs. EIA generator age: electricity sector binscatter.
Notes: Each dot represents the mean of a decile bin; error bars show 95% confidence intervals. The sample covers 7,403 firm-month observations for 42 utilities with matched non-missing recursive capital age and generator age after applying the persistent acquisition-contamination screen. Firm-years are excluded when $\psi^A_{i,t} > 30\%$, which removes 1,301 firm-month observations across 113 firm-years. 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 36); generator age is the capacity-weighted mean operating age of all generators reported in EIA Form 860. The OLS slope is $\hat{\beta} = 1.38$ with Pearson $r = 0.27$ ($p < 0.001$).

E.1.2 Validation: capital age and future emissions persistence

Table E2 provides a complementary broad-panel validation of the remaining-exposure interpretation of recursive capital age. In the capital-intensive subsample, we regress the log future-to-current Scope 1 emissions multiple on standardized recursive capital age, controls, and SIC-2-by-year fixed effects. Older capital is associated with a lower future/current emissions multiple over one-, three-, and five-year horizons. The estimates remain negative when the control set is expanded to include sales growth, ROE, and leverage. We interpret this as directional evidence that recursive capital age contains information about the remaining-emissions multiple that maps current emissions into future exposure.

Table E2: Capital Age and Future Emissions Persistence

	1-year horizon		3-year horizon		5-year horizon	
	Baseline	Rich	Baseline	Rich	Baseline	Rich
Recursive capital age	−0.019	−0.012	−0.026	−0.018	−0.037	−0.030
t-statistic	[−3.62]	[−2.06]	[−2.52]	[−1.78]	[−2.47]	[−1.96]
Observations	13,356	12,629	10,062	9,607	7,449	7,174
Firms	1,737	1,689	1,386	1,348	1,136	1,101
Source years	2003–2023	2003–2023	2003–2021	2003–2021	2003–2019	2003–2019
Baseline controls	Yes					
Rich controls	No	Yes	No	Yes	No	Yes
SIC-2 × year FE	Yes					

Notes: The dependent variable is $\log\left((\sum_{\tau=1}^{H} E_{i,t+\tau})/E_{i,t}\right)$, where $E_{i,t}$ is Scope 1 emissions and H is the horizon shown in the column heading. The sample is restricted to capital-intensive firms and requires positive current emissions and non-missing future emissions for every year in the horizon. Recursive capital age is standardized within source year, so the coefficient reports the change in the log future/current emissions multiple associated with a one-standard-deviation increase in capital age. Baseline controls are log(ME), log(PPE), Tobin's q, I/K, and book-to-market; the log scale controls are also standardized within source year. Rich controls additionally include sales growth, ROE, and leverage. The outcome is clipped at the 1st and 99th percentiles. All specifications include SIC-2-by-year fixed effects. t-statistics based on firm-clustered standard errors are reported in square brackets.

E.2 Cross-sectional regression robustness

Emissions scope robustness. Table E3 repeats the baseline post-Paris capital-intensive specification using alternative Trucost emissions scopes. Since the model concerns direct fuel use by installed capital, Scope 1 is the theoretically relevant measure. Scope 2, Scope 3, and Scope 1+2 are included as measurement checks. In the post-Paris sample, standardized Scope 2 and Scope 3 emissions are highly correlated with standardized Scope 1 emissions, both overall and within recursive-capital-age quintiles.

E.3 Robustness of the carbon-premium-by-capital-age estimates

Lagged emissions and sales controls. Table E4 repeats the capital-age quintile specification after lagging Scope 1 emissions by one additional fiscal year and controlling for contemporaneous or lagged sales. The young-capital premium remains large and statistically significant.

Growth-opportunity and financing controls. As discussed in Section 4.2, two non-stranding channels could plausibly confound the age gradient. First, the Kogan and Papanikolaou (2014) and Kogan and Papanikolaou (2013) framework implies that cross-sectional returns contain a factor structure tied to investment-specific technology risk. Second, Geelen et al. (2024) argue that firms holding long-lived physical capital face a maturity-mismatch problem vis-à-vis shorter-dated debt, and that capital structure responds to this exposure.

Table E3: The Stranding Premium Holds Across Emissions Scopes

	Scope 1	Scope 2	Scope 3	Scope 1+2
Emissions	36.5	51.1	36.8	51.9
	[2.47]	[3.11]	[1.53]	[3.29]
Capital age	9.2	9.1	7.5	9.6
	[0.79]	[0.81]	[0.70]	[0.84]
Emissions × Capital age	−12.7	−16.1	−13.9	−15.2
	[−2.00]	[−2.78]	[−2.48]	[−2.37]
Obs	120,160	120,124	120,160	120,160
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. Coefficients are reported in basis points per month. All emissions and age variables are within-month z-scored; interaction terms are products of the corresponding z-scores. Sample: capital-intensive industries, defined by a fixed SIC-2 classification based on gross PP&E / total assets measured in Compustat over 1985–1995 and adjusted to balance firm-month counts across capital-intensive and capital-light groups, January 2016 to December 2025. t-statistics based on two-way clustered standard errors (firm × calendar year) are reported in square brackets.

We address both concerns by augmenting the quintile specifications with the Kogan and Papanikolaou (2013) growth-opportunities bundle (I/K, E/P, idiosyncratic volatility, IMC-factor betas, cash holdings, dividend payout, and intangible-capital share) and the Geelen et al. (2024) financing bundle (debt-maturity measures, asset maturity, and the asset-debt maturity gap). Table E5 reports the per-quintile specification on a common matched sample of firm-months with non-missing controls. The young-old slope in the carbon premium is robust to the addition of both sets of controls, separately and jointly. The coefficient on the youngest-capital firms attenuates by approximately one-quarter (68 versus 88 basis points per month) but remains statistically significant with all controls, whereas the oldest-capital premium is statistically insignificant in all specifications.

Fama-MacBeth inference. Table E6 reports the emissions coefficient for each capital-age quintile within each of the four groups. These are the numerical values underlying Figure 3. The table also shows the same quintile decomposition using Fama-MacBeth inference. Each month is first estimated using a cross-sectional regression with SIC-2 fixed effects and the full set of controls, and the table reports Newey-West means of the monthly emissions slopes. The month-level evidence confirms that the post-Paris premium is concentrated in the youngest capital-intensive firms, while pre-Paris and capital-light estimates are close to zero or imprecise.

Table E4: Carbon premium by capital-age quintile, capital-intensive post-Paris cell: lagged emissions robustness

	Baseline controls	+ sale (matched)	Lagged em baseline	Lagged em + sale	Lagged em + lagged sale
Q1 (youngest)	83.4	74.1	83.3	74.9	69.3
	[3.90]	[4.29]	[4.80]	[5.17]	[4.41]
Q2	36.9	28.9	33.3	25.2	20.4
	[2.04]	[1.84]	[1.55]	[1.24]	[1.03]
Q3	26.8	19.9	29.7	22.6	18.4
	[1.96]	[1.71]	[1.96]	[1.59]	[1.39]
Q4	12.9	6.6	16.4	9.9	6.2
	[0.51]	[0.30]	[0.64]	[0.43]	[0.26]
Q5 (oldest)	−0.1	−5.7	−9.3	−14.9	−18.3
	[−0.01]	[−0.28]	[−0.55]	[−0.93]	[−1.12]
Emissions timing	FY $t-1$	FY $t-1$	FY $t-2$	FY $t-2$	FY $t-2$
Sales control	—	FY $t-1$	—	FY $t-1$	FY $t-2$
Observations	320,972	320,853	286,609	286,526	286,573

Notes: Each cell is the coefficient on $z\_\log(\text{Scope 1}) \times Q_k$ (basis points per month per 1-SD emissions) for the capital-intensive post-Paris cell of the fully saturated quintile × emissions × capital-intensity × post-Paris regression with year-month and SIC-2 fixed effects. Capital-age quintiles are as described in Figure 3. The *Emissions timing* row records the fiscal year of ln(Scope 1) paired with each July t+1–June t+2 return window. The *Sales control* row records the fiscal year of ln(sales) added as an additional control. FY $t-2$ entries match the lagged-emissions timing, while the FY $t-1$ entry in column 4 leaves the sales control one year forward of the extra-lagged emissions. t-statistics based on two-way clustered standard errors by firm and calendar year are reported in square brackets.

Acquisition-accounting screen. Table E7 examines whether the baseline age gradient is an artifact of acquisition accounting. In the Compustat recursion, gross investment is treated as new capital; if a firm acquires old assets, the resulting PP&E increase can make the firm appear mechanically younger even when the physical assets are not new. We address this concern by constructing a persistent acquisition-contamination share, $\psi^A_{i,t} \equiv K^A_{i,t}/\text{PPE}^{\text{net}}_{i,t}$, where $K^A_{i,t}$ is an inferred acquired-PP&E stock. We initialize $K^A_{i,t}$ in AQC-positive fiscal years using the component of gross PP&E investment (net PP&E growth plus depreciation expense) not explained by capital expenditures, depreciate it forward using firm-year depreciation rates, and rerun the quintile specification after excluding observations with $\psi^A_{i,t}$ above 5%, 10%, 20%, and 30%. The young-capital premium is robust to these screens. The post-Paris capital-intensive Q1 coefficient is 105.3 bps/month in the baseline and 134.8, 131.9, 125.9, and 114.9 bps/month under the four exclusions, respectively. The corresponding Q5 coefficients remain small and statistically indistinguishable from zero. The return pattern therefore does not appear to be driven by acquisition-related PP&E step-ups mechanically classifying old acquired assets as young capital.

Table E5: Carbon premium by capital-age quintile: sensitivity to growth-options and financing controls

	B&K baseline (1)	+Growth (2)	+Financing (3)	+Both (4)
Q1 (youngest)	87.6	80.3	68.9	68.2
	[2.79]	[2.60]	[2.28]	[2.26]
Q2	45.8	40.3	32.9	34.4
	[3.04]	[3.36]	[2.72]	[2.90]
Q3	27.0	22.3	13.9	16.6
	[1.88]	[2.34]	[1.44]	[2.06]
Q4	15.6	11.9	5.5	8.8
	[0.91]	[0.83]	[0.37]	[0.61]
Q5 (oldest)	−15.3	−18.9	−22.5	−20.1
	[−0.49]	[−0.68]	[−0.75]	[−0.72]

Notes: Each cell reports the coefficient on log(Scope 1) emissions within capital-age quintiles from a pooled regression restricted to the post-Paris capital-intensive subsample. All regressions include the baseline controls, and year-month and SIC-2 fixed effects. The Growth bundle comprises I/K, E/P, cash ratio, dividend payout, intangible capital share (Peters and Taylor, 2017), idiosyncratic volatility, and two IMC-factor betas (Papanikolaou, 2011), one FF49-based and one NAICS-based via the Leontief inverse of the BEA 2017 Benchmark I-O Use Table. The Financing bundle (Geelen et al., 2024) comprises tangibility, EBITDA/AT, firm age, size squared, cash-flow volatility, debt-maturity shares (>3 and >5 years), weighted-average debt maturity, asset maturity, and the maturity gap. t-statistics based on two-way clustered standard errors by firm and calendar year are reported in square brackets.

Book-to-market valuation check. Table E8 repeats the post-Paris quintile decomposition with log(book-to-market) as the dependent variable. The exercise follows the logic of Bolton and Kacperczyk (2023), who use book-to-market ratios as a valuation-side check on return-based carbon-premium estimates. Because book-to-market is an annual characteristic, the regression uses one July Fama–French characteristic observation per firm-year rather than repeating the same annual outcome across monthly return rows.

Within-quintile emissions variation. A possible concern with the within-quintile estimation is that sorting firms on capital age may compress the emissions distribution within bins, leaving each quintile-specific slope weakly identified. Table E9 addresses this directly by decomposing the variance of ln(Scope 1) in the post-Paris capital-intensive subsample into between-quintile and within-quintile components and by cross-tabulating capital-age and emissions quintiles.

MCCI indicator robustness. Table E10 reports the low- and high-state emissions premia by capital-age quintile for two additional MCCI measures (the aggregate index and the Climate Legislation/Regulations subindex), separately for capital-intensive and capital-light firms. Across both MCCI measures, the high-state premium for the youngest capital-

Table E6: Carbon premium by capital-age quintile: pooled and Fama-MacBeth estimates

	Panel A: Pooled OLS				Panel B: Fama-MacBeth			
	Capital-intensive		Capital-light		Capital-intensive		Capital-light	
Age quintile	Post	Pre	Post	Pre	Post	Pre	Post	Pre
Q1 (youngest)	83.4	−34.8	16.4	29.5	53.8	−24.8	−10.4	21.8
	[3.90]	[−1.77]	[0.68]	[1.14]	[2.93]	[−1.35]	[−0.52]	[0.95]
Q2	36.9	−30.4	−8.9	8.2	10.6	−24.4	−33.1	11.6
	[2.04]	[−1.62]	[−0.67]	[0.67]	[0.52]	[−1.56]	[−1.89]	[0.82]
Q3	26.8	−15.4	0.6	5.6	6.0	−1.7	−7.7	11.6
	[1.96]	[−1.09]	[0.02]	[0.42]	[0.35]	[−0.16]	[−0.43]	[0.77]
Q4	12.9	−5.1	14.6	7.0	7.4	−0.9	−2.0	19.9
	[0.51]	[−0.30]	[1.08]	[0.35]	[0.45]	[−0.10]	[−0.14]	[1.34]
Q5 (oldest)	−0.1	−12.6	10.5	3.2	−7.7	−14.7	−2.2	2.6
	[−0.01]	[−0.64]	[0.80]	[0.45]	[−0.50]	[−1.18]	[−0.15]	[0.17]
Observations	320,972				223,007	97,965	223,007	97,965
Months	252				120	132	120	132
Controls	B&K full set							
Fixed effects	Year-month + SIC-2				SIC-2 within each month			

Notes: Each cell reports the coefficient on within-month standardized ln(Scope 1) interacted with capital-age quintile Q_k, in basis points per month per 1-SD emissions. Panel A reports pooled OLS estimates from a fully saturated quintile × emissions × capital-intensity × post-Paris specification with year-month and SIC-2 fixed effects; t-statistics use two-way clustered standard errors by firm and calendar year. Panel B reports time-series means of monthly cross-sectional emissions slopes from Fama-MacBeth regressions that absorb SIC-2 fixed effects within each month; t-statistics use Newey-West standard errors with six lags. Both panels use the baseline controls and the same capital-age quintiles, formed within SIC-2 industry × period.

intensive firms exceeds 100 basis points per month, while the oldest-quintile high-state premium is much smaller and less precise. The corresponding capital-light estimates are smaller and less precise.

Table E7: Persistent Acquisition-Contamination Robustness

	Baseline	Drop $\psi^A > 5\%$	Drop $\psi^A > 10\%$	Drop $\psi^A > 20\%$	Drop $\psi^A > 30\%$
Q1 (youngest)	105.3	134.8	131.9	125.9	114.9
	[4.29]	[2.95]	[3.33]	[3.72]	[3.55]
Q2	37.6	43.5	52.4	50.7	45.7
	[1.81]	[1.31]	[1.70]	[1.88]	[1.84]
Q3	32.8	38.3	45.6	48.8	42.6
	[1.89]	[1.33]	[1.66]	[1.99]	[1.87]
Q4	34.6	56.6	56.0	51.7	44.0
	[3.54]	[3.48]	[4.62]	[5.43]	[5.13]
Q5 (oldest)	2.3	−7.0	12.7	17.9	11.5
	[0.12]	[−0.24]	[0.47]	[0.79]	[0.56]
Observations	289,316	117,146	148,916	191,987	220,061
Controls			B&K full set		
Fixed effects			Year-month + SIC-3		

Notes: Entries are the post-Paris capital-intensive coefficients, in basis points per month, on standardized Scope 1 emissions interacted with recursive-capital-age quintile indicators. The regression is the same saturated quintile-by-emissions-by-capital-intensity-by-pre/post-Paris specification used in the baseline quintile table, with capital-age quintiles formed using the fallback ladder described in Figure 3: SIC-3 $\times$ period cutoffs, SIC-3 pooled cutoffs, SIC-2 $\times$ period cutoffs, SIC-2 pooled cutoffs, and global cutoffs split by capital intensity; only the displayed cells are restricted to the post-Paris capital-intensive group. ψ^A is the persistent acquired-capital share, defined as K^A/PPE^{net}, where K^A is the stock of acquired PP&E. The acquired-capital stock is initialized in Compustat AQC-positive fiscal years using the component of gross PP&E investment—net PP&E growth plus depreciation expense—not explained by capital expenditures, and is depreciated forward using the firm-year depreciation rate. t-statistics based on two-way clustered standard errors by firm and calendar year are reported in square brackets.

Table E8: Post-Paris Log Book-to-Market Slopes by Capital-Age Quintile

	Capital-age quintile				
Group	Q1 (youngest)	Q2	Q3	Q4	Q5 (oldest)
Capital-intensive	0.198	0.173	0.105	0.046	0.033
	[4.73]	[4.91]	[2.97]	[1.40]	[0.86]
Capital-light	0.067	0.042	0.022	0.034	0.003
	[1.16]	[0.76]	[0.46]	[0.66]	[0.05]
Obs	25,690				
Controls	Yes				
Year-month FE	Yes				
SIC-2 FE	Yes				

Notes: The dependent variable is log(book-to-market), where book-to-market is common equity divided by December market equity. The control set is the baseline set, excluding book-to-market (the dependent variable) and ln(market cap). t-statistics based on two-way clustered standard errors by firm and calendar year are reported in square brackets.

Table E9: Within-quintile emissions variation and firm-month counts

Panel A: Variance decomposition of ln(*Scope 1*)

Conditioning set	Between-quintile share	Within-quintile share
V1: Raw (no conditioning)	0.079	0.921
V2: Within year-month	0.090	0.910
V3: Within year-month × SIC-2	0.287	0.713

Panel B: Joint distribution of capital-age and Scope 1 emissions quintiles (firm-month counts)

Capital-age quintile	E1	E2	E3	E4	E5
Q1 (youngest)	8,432	6,413	4,538	3,351	2,404
Q2	6,297	5,337	5,921	5,020	3,974
Q3	5,533	5,414	5,606	5,015	5,050
Q4	3,857	4,936	6,014	6,341	6,365
Q5 (oldest)	2,714	4,655	4,671	7,028	9,015

Notes: The sample is the post-Paris capital-intensive subsample. Capital-age quintiles use SIC-2 × period balanced cutoffs with the fallback ladder described for Figure 3. Panel A reports a variance decomposition of ln(Scope 1) across capital-age quintiles, partialling out the conditioning set in column 1 before computing shares. Panel B reports firm-month counts in a 5 × 5 cross-tabulation of capital-age and Scope 1 emissions quintiles, with emissions quintiles formed independently within each calendar month.

Table E10: Carbon premium by capital-age quintile and MCCI climate-attention state

	Aggregate Index				*Climate Legislation/Regulation*			
	Capital-intensive		Capital-light		Capital-intensive		Capital-light	
Age quintile	Low	High	Low	High	Low	High	Low	High
Q1 (youngest)	6.7	105.2	14.5	35.6	−2.0	106.5	35.1	26.8
	[0.16]	[3.37]	[0.39]	[1.77]	[−0.07]	[3.92]	[0.99]	[1.48]
Q2	−43.6	72.3	−0.4	−5.3	−56.2	77.1	−5.4	−3.6
	[−1.59]	[2.62]	[−0.02]	[−0.38]	[−2.05]	[3.00]	[−0.21]	[−0.23]
Q3	−39.7	71.0	−7.9	10.9	−49.1	70.0	−11.7	12.1
	[−2.02]	[3.08]	[−0.48]	[0.50]	[−2.35]	[3.96]	[−0.56]	[0.48]
Q4	−11.7	53.3	4.3	−0.3	−25.6	57.2	−16.3	8.6
	[−0.55]	[3.19]	[0.15]	[−0.01]	[−1.20]	[4.04]	[−0.47]	[0.47]
Q5 (oldest)	−11.0	21.9	−0.1	4.1	−6.9	17.9	−23.1	13.8
	[−0.36]	[1.18]	[−0.01]	[0.27]	[−0.27]	[1.11]	[−1.12]	[0.89]
Obs	152,823		123,833		152,823		123,833	
Low-state obs	56,814		44,562		52,562		39,884	
High-state obs	96,009		79,271		100,261		83,949	

Notes: The table repeats the full-sample state-dependent carbon-premium specification in Table 5 for two MCCI climate-attention indicators. Within each MCCI block, the first two columns report capital-intensive industries and the next two columns report capital-light industries. Each block estimates monthly excess returns on ln(Scope 1 Emissions) interacted with capital-age-quintile dummies and with the high-state indicator for the indicated MCCI series. Coefficients are reported in basis points per month. Controls and fixed effects follow Table 3. t-statistics based on two-way clustered standard errors (firm $\times$ calendar year) are reported in square brackets.

F Conditional portfolio robustness

This appendix reports robustness checks for the conditional stranding portfolio evidence in Section 4. Table F1 compares the Legal Actions signal with related MCCI text signals.

Table F1: Text-signal comparison for signal-timed STRAND portfolios

Portfolio	Signal	Median L/S		Q4–Q1 L/S	
		Sharpe	FF3 α (% p.a.)	Sharpe	FF3 α (% p.a.)
STRAND	Legal	0.79	3.3 [2.03]	1.01	3.7 [3.11]
STRAND	Leg./reg.	0.81	4.1 [2.91]	0.67	2.0 [2.02]
STRAND	Policy composite	0.69	3.0 [1.80]	0.92	3.3 [3.33]
STRAND	Aggregate	0.87	3.6 [2.37]	0.78	2.7 [2.37]

Notes: The table reports no-cost post-Paris performance for the industry-neutral STRAND double-difference portfolio using alternative one-month-lagged MCCI text signals. All signals are detrended and volatility-standardized with the same expanding-window procedure as the main Legal Actions overlay. The signal histories are current-vintage MCCI index files, so the signal-timed rows are descriptive historical portfolio exercises rather than archived-vintage investable backtests. Sharpe ratios are annualized. FF3 alphas are annualized percentage points, with HAC t-statistics reported in square brackets below alpha estimates. Leg./reg. is the MCCI Climate Legislation/Regulations index.