Pricing Transition Risk on Spatial Networks

Jose Luis Resendiz

February 23, 2026

Abstract: How do financial markets price the credibility of the climate transition? While a growing literature documents systematic climate risk premia and network-propagated financial contagion, existing approaches have not isolated the distinct spatial channels through which physical restructuring shocks transmit to equity prices. Using plant-level geolocation data for 389 listed power utilities worldwide, we show that a single plant retirement transmits shocks through opposing channels: a negative contagion effect for technologically similar firms (stranding risk) and a positive competitive benefit for geographically proximate neighbors. Economically, a one-standard-deviation increase in fuel-similarity exposure translates to a 15-basis-point drop in abnormal returns over three months. Furthermore, we reveal a persistent credibility gap in how markets price these transitions. Idiosyncratic, voluntary plant retirements act as ambiguous signals that spike peer return volatility without moving mean prices, consistent with a cheap-talk equilibrium. In contrast, binding phase-out laws resolve this uncertainty, reducing volatility and generating directional repricing proportional to fossil exposure. Ultimately, our findings show that markets observe physical network exposure but heavily discount voluntary climate action; standardizing corporate disclosure cannot efficiently reallocate capital as long as the underlying environmental policy remains discretionary.

1. Introduction

Pricing climate transition risk has become one of the defining challenges in financial economics (Giglio et al., 2021). Meeting the 2°C target requires leaving over 80 percent of coal reserves unextracted (McGlade and Ekins, 2015), concentrating stranded asset exposure in specific geographies; under tail scenarios, up to $24 trillion in global financial assets may be at risk (Dietz et al., 2016). A broad consensus among finance academics and practitioners holds that markets underestimate this exposure (Stroebel and Wurgler, 2021), and the difficulty is compounded by deep Knightian uncertainty that propagates across carbon dynamics, temperature dynamics, and damage functions simultaneously (Barnett et al., 2020). The question is no longer whether transition risk is financially material, but whether existing measurement tools can capture its spatial and systemic dimensions.

The financial industry's dominant response has been firm-level scoring: ESG ratings, implied temperature scores, and target validation frameworks that treat transition credibility as a property of the individual firm. This approach suffers from well-documented

limitations. Major rating providers disagree on basic assessments of the same company, with pairwise correlations as low as 38 percent (Berg et al., 2022). Even mandatory sustainability reporting fails to produce consistent data across firms, because underlying incentives create strategic variation in what companies disclose and how (Christensen et al., 2021). Corporate climate targets frequently fail basic integrity criteria, and markets exhibit weak reactions when firms miss stated goals (Jiang et al., 2025). Institutional investors overwhelmingly believe climate risks are financially material yet report lacking adequate tools to measure portfolio-level exposure (Krueger et al., 2020). The problem, in short, is not a shortage of data or disclosure mandates but a conceptual gap: firm-level approaches cannot capture risks that originate in the network connecting firms.

A growing body of work suggests that transition risk is fundamentally a network property. Central banks have warned that climate change may trigger "green swan" events, systemic crises propagating through cascade effects that conventional risk models cannot capture (Bolton et al., 2020). The post-carbon transition demands systems-level thinking because nonlinear feedbacks in technology, policy, and finance can amplify small interventions into large systemic shifts (Farmer et al., 2019). Network theory has established that systemic risk exhibits phase transitions: the same interconnections that diversify small shocks amplify large ones beyond a critical threshold (Acemoglu et al., 2015). Applied to climate policy, network-based stress tests reveal that direct and indirect exposures to policy-relevant sectors represent nearly half of Euro Area equity portfolios (Battiston et al., 2017), and omitting the financial system from climate transition models creates a dangerous blind spot for scenario analysis (Battiston et al., 2021). These findings point toward measuring transition exposure through the spatial economic networks connecting firms, rather than through subjective ratings of individual companies.

This paper develops and tests such an approach. We construct a spatial network framework for pricing transition risk that bridges the macro-asset pricing literature documenting the scale of the problem, the network finance literature providing the theoretical architecture, and the firm-level evidence exposing the limits of isolated scoring. By measuring exposure through hard physical and economic linkages, specifically fuel-mix similarity, geographic proximity, and regulatory proximity, the framework sidesteps the noise that plagues ESG-based approaches. It also yields a testable distinction: when the direction of a network shock is ambiguous, markets should respond first through volatility rather than directional returns; when binding policy resolves the ambiguity, prices should adjust.

We document three empirical facts. First, spatial transmission exhibits a short-horizon channel split. Over a three-month horizon, fuel-mix similarity predicts negative repricing ($w^{\text{fuel}}_{[-1,+3]} = -1.497$) while geographic proximity generates positive repricing ($w^{\text{geo}}_{[-1,+3]} = 0.354$); regulatory proximity is weak. To contextualize these magnitudes, a 0.10 increase in fuel similarity translates to approximately 15 basis points of negative

three-month abnormal returns, while geographic proximity yields a modest 3.5 basis-point competitive benefit. At longer horizons (6–12 months), these estimates lose statistical precision, consistent with the gradual incorporation of complex spatial information by informed investors. Second, plant-level retirement announcements (EIA-860 filings) raise neighbouring volatility with no directional return response, consistent with mean-preserving uncertainty. Third, binding phase-out laws reduce volatility and are associated with negative abnormal returns for coal-exposed firms, consistent with uncertainty resolution and stranding. The channel decomposition carries the signal.

To interpret these patterns, we develop a networked asset-pricing framework under policy uncertainty. The framework defines a "credibility gap" between physical exposure and what markets can price in real time. Because firms' interests in maintaining fossil operations diverge from investors' climate-risk assessments, voluntary commitments function as a cheap-talk equilibrium (Crawford and Sobel, 1982), leaving policy uncertainty unresolved. Idiosyncratic shocks (plant retirements) merely inject ambiguous noise, inflating volatility. Only when a binding policy removes managerial discretion does uncertainty collapse, forcing directional asset repricing.

Ultimately, this paper makes three broad contributions to climate finance and policy. First, by documenting the short-horizon channel split, we reveal a paradox in isolated climate action: because geographic proximity provides a competitive benefit ($\gamma_G > 0$), firm-by-firm divestment campaigns risk inadvertently subsidizing neighboring fossil assets. Second, by showing that ambiguous commitments transmit as volatility rather than directional returns, we highlight the structural limits of financial disclosure: standardizing ESG reporting cannot efficiently reallocate capital if the underlying environmental policy remains discretionary. Third, by introducing a plant-level spatial network empirical design, we provide a framework that overcomes ESG-rating noise and bypasses the "brown-spinning" illusion, wherein dirty assets are transferred to private markets rather than genuinely retired. The energy transition is a uniquely clean laboratory for these mechanisms, but the implications extend to any networked sector facing systemic regulatory change.

The paper proceeds as follows. Section 2 develops the empirical framework and lays out the empirical strategy. Section 3 presents results. Section 4 discusses implications for transition risk assessment and policy. Section 5 concludes. Appendices provide channel definitions, derivations, robustness tables, and data construction details.

2. Empirical Framework: Networked Pricing under Policy Uncertainty

2.1. Overview

Industry-standard assessment tools — ESG ratings, implied temperature scores, TCFD-aligned disclosures — treat each firm in isolation, assuming markets can continuously

and accurately price firm-level commitments. We propose a different framework: the **credibility gap** is fundamentally an information friction. Investors observe the physical network of power utilities and each firm's legacy asset exposure, but they cannot reliably assign a directional price to transition commitments until binding policy removes managerial and political discretion.

We frame transition risk as an asset-pricing problem on a multi-layered spatial network under regulatory uncertainty. This framework yields two core short-horizon insights. First, spatial transmission operates through distinct channels with opposing signs. Second, when the regulatory regime is uncertain, idiosyncratic shocks (e.g., plant-level retirements) act as noisy signals that increase return volatility without a clear directional mean. Conversely, binding phase-out laws partially resolve uncertainty, compress volatility, and generate negative returns proportional to firms' own fossil exposure.

2.2. Environment and Networked Cash Flows

Consider a continuum of listed power utilities $i \in [0, 1]$. Firm i's fundamental value is the expected present discounted value of its future cash flows, $V_i = \mathbb{E}[\Pi_i]$. Each firm is characterized by its **legacy asset intensity**, $\alpha_i \in [0, 1]$, which measures reliance on fossil generation (e.g., coal share).

Firms are interconnected via a spatial network decomposed into three row-normalized adjacency matrices: geographic proximity ($\mathbf{W}^{\text{geo}}$), fuel-mix similarity ($\mathbf{W}^{\text{fuel}}$), and regulatory proximity ($\mathbf{W}^{\text{reg}}$). Let $s_j \geq 0$ denote a physical restructuring shock originating at neighbor j (e.g., a capacity retirement). We specify firm i's realized cash flows as a function of own exposure, neighbor shocks, and an unobserved regulatory state $\omega \in \{0, 1\}$:

$$\Pi_i(\omega) = \bar{\Pi}_i - \omega\beta\alpha_i + \gamma_G \sum_j w_{ij}^{\text{geo}} s_j + \gamma_R \sum_j w_{ij}^{\text{reg}} s_j - \omega\gamma_F \sum_j w_{ij}^{\text{fuel}} s_j, \tag{1}$$

where $\bar{\Pi}_i$ is baseline profitability. The channels operate as follows:

- $\gamma_G > 0$ captures **local competitive benefits**: when a geographically proximate neighbor retires capacity, local grid congestion eases and wholesale margins improve.
- $\gamma_R \geq 0$ captures **regulatory spillovers**: shared jurisdiction effects, which we allow to be positive but noisy at short horizons.
- $\gamma_F > 0$ captures **competitive stranding**: when a fuel-similar peer retires capacity, it signals shared technological obsolescence, but this penalty binds only if the transition regime is enforced ($\omega = 1$). The sign is ultimately an empirical question — technology-similar peers could in principle benefit from shared learning ($\gamma_F < 0$) — and the data adjudicate.

- $\beta > 0$ is the direct penalty on the firm's own legacy assets (α_i) under a binding transition ($\omega = 1$).

The opposing-sign spatial structure formalizes the distinction between contagion and competitive effects identified by Lang and Stulz (1992): the same event (peer exit) produces opposite-signed spillovers depending on the dimension of proximity. Jorion and Zhang (2007) confirm this dual-sign pattern in credit default swap markets, where the nature of distress determines whether peer spreads widen (contagion) or tighten (competitive benefit). In electricity markets specifically, Davis and Hausman (2016) demonstrate that plant closure in a geographically constrained grid increases market power for nearby generators, providing the microeconomic mechanism for $\gamma_G > 0$. The network propagation structure follows Acemoglu et al. (2012), who show that idiosyncratic shocks propagate through weighted adjacency matrices with magnitudes determined by network topology; Herskovic (2018) establishes that such production network linkages are priced in equilibrium.

2.3. The Credibility Gap and Asset Pricing

The market does not perfectly observe the true regulatory state ω. Let $p_t = \Pr_t(\omega = 1)$ denote the market's prior probability at time t that the strict transition regime will be enforced.

Definition 1 (The Credibility Gap)**.** *The credibility gap is the variance in market beliefs regarding policy enforcement, proportional to $p_t(1-p_t)$. When p_t is near 0.5, uncertainty is maximized, and markets heavily discount stated transition targets because discretionary reversal remains highly probable.*

Following Pástor and Veronesi (2012), who demonstrate that uncertainty about government policy depresses stock prices and generates a risk premium proportional to $p_t(1-p_t)$, we treat this regulatory uncertainty as the core friction driving transition risk. The carbon premium documented by Bolton and Kacperczyk (2021) is consistent with markets partially pricing this risk ($p_t > 0$) while remaining uncertain about ultimate enforcement ($p_t < 1$). Ilhan et al. (2021) confirm that option-market pricing of carbon tail risk responds dynamically to such credibility signals.

The market valuation of firm i at time t is the expectation of Equation (1):

$$V_{i,t} = \bar{\Pi}_i - p_t \beta \alpha_i + \gamma_G \sum_j w_{ij}^{\text{geo}} s_j + \gamma_R \sum_j w_{ij}^{\text{reg}} s_j - p_t \gamma_F \sum_j w_{ij}^{\text{fuel}} s_j. \tag{2}$$

The variance of firm i's value, driven by regulatory uncertainty, is

$$\text{Var}_t(V_{i,t}) = p_t(1-p_t) \left(\beta \alpha_i + \gamma_F \sum_j w_{ij}^{\text{fuel}} s_j \right)^2. \tag{3}$$

2.4. Empirical Predictions

This framework maps different shock types to observable short-horizon returns ($R_i = \Delta\mathbb{E}_t[V_{i,t}]$) and return volatility.

Empirical Prediction 1 (Short-Horizon Channel Split)**.** *Holding policy uncertainty p_t constant, a generic vector of physical restructuring shocks* $\mathbf{s}$ *transmits positively through the geographic network ($\gamma_G > 0$). It transmits negatively through the fuel-similarity network ($-p_t\gamma_F < 0$).*

Derivation: Immediate from the partial derivatives of Equation (2) with respect to neighbor shocks. *Empirical mapping:* This formalizes the channel decomposition design. Rather than a single composite peer effect, spatial transmission splits by sign at the 3-month horizon (documented in Table 2).

We next evaluate how the market responds to variations in signal clarity.

Empirical Prediction 2 (Ambiguous Shocks: Volatility Transmission)**.** *Suppose an idiosyncratic plant retirement announcement (e.g., EIA-860) acts as a noisy signal of ω. While Bayesian updating may shift p_t marginally, the first-order observable effect for exposed neighbors is a spike in return volatility, with a theoretically ambiguous (or weak) directional mean return due to offsetting geographic and fuel channels.*

Derivation: With $p_t \in (0, 1)$, Equation (3) shows that an increase in s_j strictly increases $\text{Var}_t(V_{i,t})$. Because the signal is noisy, the update to p_t is small, and the opposing γ_G and γ_F channels mute the directional mean return. *Literature grounding:* This prediction is consistent with Epstein and Schneider (2008), who show that ambiguity-averse investors evaluate uncertain-quality signals under a worst-case assessment of precision, inflating perceived variance even when the mean impact is zero. Veronesi (1999) provides the general-equilibrium mechanism: return volatility peaks at maximum posterior uncertainty about the hidden state, precisely where $p_t(1 - p_t)$ is maximized.

Empirical Prediction 3 (Resolving Shocks: Directional Repricing)**.** *Suppose a binding phase-out law is enacted. This systemic shock partially resolves policy uncertainty, causing a discrete jump in beliefs $p_{post} \gg p_{pre}$. For fossil-exposed firms, this strictly reduces return volatility and generates negative abnormal returns proportional to their own legacy intensity (α_i).*

Derivation: As p_{post} approaches 1, the uncertainty term $p(1 - p)$ in Equation (3) drops, compressing volatility. From Equation (2), the change in value is strictly negative and scales with own coal-share α_i. *Empirical mapping:* Predictions 2 and 3 distinguish the market reaction to ambiguous filings (volatility increases, no CAR) versus Tier-1 binding laws (volatility falls, negative CAR scales with α_i), tested in Table 3. *Literature*

grounding: The volatility-collapse prediction draws on Kelly et al. (2016), who show across 20 countries that implied volatility drops sharply once political uncertainty resolves. Bloom (2009) documents the canonical pattern: uncertainty episodes approximately double volatility, with rapid normalization upon resolution. The α_i-proportional negative repricing is validated by Hsu et al. (2023), who show that a pollution premium of 4.42% per annum is strictly proportional to within-industry emission intensity and driven by environmental policy regime-shift risk.

Remark (Valuation wedge / credibility gap)**.** *If firm-level commitments were reliably priced, directional returns would separate credible from non-credible firms around restructuring news. Instead, we observe volatility responses with weak mean returns for ambiguous shocks. We interpret this as a credibility gap: markets observe exposure (fuel mix, geography) but cannot price commitment quality until binding policy resolves the direction. Formally, this friction is captured by the belief variance $p_t(1-p_t)$ in Definition 1. Our model parsimoniously collapses two distinct sources of uncertainty into the single parameter p_t: uncertainty about whether* governments *will enforce the transition regime, and uncertainty about whether* firms *will honour voluntary commitments. Both sources leave p_t unresolved because voluntary pledges function as cheap talk in the sense of Crawford and Sobel (1982); only binding regulation acts as a costly signal.*

2.5. Empirical Strategy

We test the framework on global listed power utilities — a sector that offers three advantages over alternatives. First, the physical asset base is directly observable: power plants have geolocated coordinates, rated capacities, fuel types, and commission/retirement dates recorded in publicly available registries. This means the legacy intensity parameter α and all spatial exposure measures are built from physical data rather than self-reported disclosures. Second, the spatial network is well-defined: utilities are connected through transmission grids and shared wholesale markets, shared regulatory jurisdictions, and local factor markets for engineering labour and equipment — connections that are directly documented in public registries. Third, the sector contains large variation in financial fundamentals (from unconstrained Nordic hydro operators to heavily leveraged Southeast Asian coal incumbents) and network density (from isolated island systems to the deeply interconnected European grid), providing the cross-sectional variation needed to test the networked transmission mechanism through return and volatility responses. This single-sector focus follows a distinguished empirical tradition: Cicala (2022) and Fabrizio et al. (2007) demonstrate that electricity generation is a uniquely transparent laboratory where plant-level inputs, outputs, and costs are directly observable, eliminating the measurement noise that plagues cross-sector studies reliant on ESG scores or self-reported disclosures.

Sample Construction

The sample comprises global listed power utilities identified by SIC codes 4911, 4931, and 4932 (electric services) and NAICS 2211 in Compustat Global/NA, cross-referenced with plant ownership records in the Global Energy Monitor (GEM) Global Power Plant Tracker (and GeoAsset where available). We require firms to have (i) at least one operating power plant identifiable in GEM or GeoAsset, (ii) financial statements in Compustat Global/NA, and (iii) equity return data in CRSP (US firms) or Datastream (non-US firms). Appendix C documents an interim replication that uses Compustat Security returns when CRSP/Datastream coverage is incomplete. The final sample comprises 389 firms across 50+ countries, with plant-level data covering the bulk of global installed capacity. The event set includes 179 first-mover retirements with exact or announcement dates, 92 EIA-860 retirement announcements, and 27 national or subnational phase-out laws (14 Tier-1 binding shocks). Summary statistics for the analysis sample are reported in Table 1, and regional composition is reported in Appendix C.

Data Sources and Variable Construction

Panel A: Firm Fundamentals ($\boldsymbol{\theta}$). The four fundamental parameters $(\lambda, \rho, \kappa, \alpha)$ are constructed from three core databases. Financial statements from Compustat Global/NA provide leverage, an operating-ROA proxy for return spread, and an interest-coverage proxy for cash-flow adequacy. Plant-level asset data from GEM/GeoAsset provide legacy intensity. Equity returns from CRSP (US) and Datastream (non-US) are used for return-based specifications and market adjustment. Full variable definitions and supplementary sources are reported in Appendix C.

Panel B: Transformation Metrics (Physical Exposure). We compile physical transformation metrics to construct exposure and validate shock classification. These include: (i) capacity mix shifts (changes in clean MW share over 3- and 5-year horizons), (ii) emissions intensity changes, and (iii) binary asset retirement events. Capacity mix shifts are constructed from GEM/GeoAsset plant-level records of additions and retirements; when panel coverage is limited, we use the cross-sectional clean share $(1-\alpha)$ to summarise legacy intensity. Emissions intensity uses LSEG ESG Scope 1 data supplemented by implied emissions from capacity mix. These measures are long-horizon outcomes and are reported in Appendix C; the main outcomes in this paper are returns and volatility around restructuring shocks.

Panel C: Spatial Weight Matrix. The spatial weight matrix $\mathbf{W}$ is constructed from geographic proximity of plant locations (GEM/GeoAsset GPS coordinates) combined with shared regulatory jurisdiction (ETS membership from the World Bank Carbon Pric-

ing Dashboard; ICAP is used as a cross-check). The geographic component uses inverse-distance weights with exponential decay; the regulatory component is binary (same ETS = 1). The fuel-similarity component $\mathbf{W}^{\text{fuel}}$ uses the cosine similarity of firm-level generation portfolio vectors (capacity shares across coal, gas, oil, wind, solar, hydro, nuclear), following the text-based network methodology of Hoberg and Phillips (2016); geographic proximity generates return comovement among co-located firms (Pirinsky and Wang, 2006). Following Grieser et al. (2022), we also estimate under each component separately. Full construction details are in Appendix C. Sample sizes vary across specifications because return-window availability and exposure coverage drop different event–firm pairs; each table reports its own N.

Identification Strategy

The central identification challenge is the reflection problem (Manski, 1993): spatial correlation in outcomes may reflect (a) network interaction effects, (b) correlated fundamentals, or (c) common shocks. Bramoullé et al. (2009) show that peer effects are generically identified when network structure is non-uniform. Our three independently constructed weight matrices provide cross-network exclusion restrictions: firm A's geographic neighbor B need not be fuel-similar to A's fuel-linked neighbor C, generating the variation required for identification. We therefore emphasize the event-study strategy tied to the channel-split and uncertainty predictions, and report robustness and ancillary diagnostics in Appendix D.

A primary identification challenge in spatial asset pricing is the endogeneity of the restructuring shocks themselves. Plant retirements are not randomly assigned; they are endogenous to local market conditions, input prices, and regulatory pressures. Consequently, a naive spatial regression risks conflating true network spillovers with correlated effects (Manski, 1993) — for instance, a local macroeconomic downturn could simultaneously trigger a plant retirement and depress the equity returns of neighbouring utilities.

We address this endogeneity through the asymmetric structure of our spatial decomposition. While an omitted common shock would typically induce a uniform directional bias (e.g., depressing all local returns), our empirical design tests for opposing-sign responses across non-nested networks. Specifically, an omitted negative local shock cannot readily explain why geographic proximity yields a *positive* competitive benefit ($\gamma_G > 0$) while fuel-mix similarity yields a *negative* contagion effect ($\gamma_F < 0$). By isolating these channel-differentiated responses, and adjusting returns for market-wide movements, our identification strategy filters out generic macroeconomic shocks to isolate the targeted transmission of transition risk. Furthermore, we leverage the distinct network topology formalized by Bramoullé et al. (2009): because firm A's geographic neighbour B is not necessarily its fuel-similar peer C, the cross-network exclusion restrictions break the re-

flection problem inherent in uniform peer groups. We note that a structured confounder — one that simultaneously benefits nearby firms and penalises fuel-similar ones — could in principle mimic the opposing-sign pattern. Such confounders are difficult to construct in practice, because they require a single omitted variable to operate with opposite signs along two independent network dimensions; the low pairwise correlation among our weight matrices ($r < 0.16$, Table 5) makes this implausible. Appendix D validates this strategy empirically, and a firm fixed-effects specification confirms that the channel split survives the absorption of all time-invariant firm heterogeneity (Table 7).

Spatial Transmission in Returns and Volatility. We study two types of shocks with different information content.

- **Plant-level retirement announcements (EIA-860).** These events convey new information about individual plants but do not resolve the direction of regional adjustment. We test whether exposure affects *volatility* rather than mean returns.
- **Binding phase-out laws.** These events resolve the direction of adjustment at the jurisdiction level. We test whether coal-exposed firms experience negative abnormal returns and lower volatility.

To our knowledge, the use of EIA-860 retirement filings as information events for financial markets is novel; we are not aware of prior studies that have employed regulatory capacity registrations as event dates. The volatility-based approach follows Beaver (1968), who established that return variance around information events is informative independent of mean price changes, and Hann et al. (2019), who demonstrate that corporate announcements transmit volatility to industry peers as a distinct second-moment information channel. For climate policy events, Koch et al. (2016) validate the event-study methodology applied to discrete regulatory announcements. The coal-share×post-event interaction for binding laws follows the DID specification of Seltzer et al. (2022), who interact continuous environmental exposure with the Paris Agreement date; Bolton and Kacperczyk (2023) establish that carbon premia intensify in stricter policy environments, supporting the hypothesis that binding legislation produces measurable repricing.

For returns, we estimate spatial event-study regressions using monthly CARs and channel-specific exposure:

$$\mathrm{CAR}_j = \alpha + \beta_{\mathrm{geo}} w_{ij}^{\mathrm{geo}} + \beta_{\mathrm{fuel}} w_{ij}^{\mathrm{fuel}} + \beta_{\mathrm{reg}} w_{ij}^{\mathrm{reg}} + \beta_s \mathrm{SameSector}_j + \varepsilon_j.$$

For volatility, we compute event-time changes in realized volatility and estimate the corresponding exposure coefficients. The primary specification uses vwretd-adjusted returns and event-clustered standard errors; two-way clustering is reported as conservative robustness.

Coal-share intensity for phase-out laws. For binding regulatory shocks we use a continuous treatment intensity: firm coal capacity share (coal MW / total MW) from GEM. This avoids the zero-inflation of spatial weights and aligns treatment with the jurisdictional scope of the law.

Robustness, Placebos, and Falsifiability

The design is falsifiable: if the channel split is absent at short horizons, spatial transmission is not operating through identifiable mechanisms; if retirement announcements do not raise volatility, the uncertainty mechanism is unsupported; if binding laws do not generate negative returns for coal-exposed firms, stranding is not priced. Appendix D reports window sensitivity and clustering variants.

3. Results

The core empirical finding is a short-horizon channel split: the same physical event, a neighbor's capacity retirement, transmits positively through the geographic network and negatively through fuel-mix similarity. We first document this pattern in cross-sectional return regressions, then show that the nature of the shock determines whether markets respond through volatility or through prices. All specifications use market-adjusted returns (vwretd) with event-clustered standard errors; Appendix D reports robustness across windows, clustering, and alternative exposure definitions.

Table 1: Summary Statistics (Analysis Sample)

Variable	N	Mean	SD	Min	Median	Max
Leverage (λ)	389	0.397	0.261	0.000	0.397	3.704
Return spread (ρ)	389	0.068	0.144	-1.754	0.076	0.305
Cash flow adequacy (κ)	389	13.073	84.352	-41.909	3.521	1211.712
Legacy intensity (α)	389	0.382	0.426	0.000	0.140	1.000
Delivery ($1-\alpha$)	389	0.618	0.426	0.000	0.860	1.000

Notes: Cross-sectional summary statistics for the 389-firm analysis sample (firms with complete $\boldsymbol{\theta}$ vector). κ is winsorized at the 1st/99th percentiles in regressions; this table reports raw values. Variable definitions are in Appendix C.

Table 1 summarises the cross-sectional distribution of fundamentals. The most striking feature is the skewness of delivery: most firms in the sample already operate predominantly clean portfolios (median clean share 0.86), while a heavy left tail of fossil-intensive incumbents pulls the mean down to 0.62. This distributional shape matters because the

channel-split and volatility results are identified primarily from variation among the fossil-intensive firms that populate that left tail. Leverage is moderate on average ($\bar{\lambda} = 0.40$) but varies considerably, from unlevered Nordic hydro operators to Southeast Asian coal generators carrying debt-to-asset ratios above three.

3.1. Channels and Uncertainty Transmission

The central question is whether the spatial network carries differentiated information about transition risk, or whether a single composite exposure measure suffices. We address it by regressing firm-level CARs on the three row-normalized spatial weights simultaneously with a same-sector control, as specified in Section 2.5, using event-clustered standard errors.

Table 2 reports the channel decomposition. The three-month window reveals a clean split, consistent with the opposing partial derivatives of the valuation equation (Equation 2). The fuel-similarity coefficient is negative and highly significant (-1.497, $t = -3.16$), indicating that firms whose generation portfolios resemble the event firm's experience negative repricing when a neighbor retires capacity, as predicted by $\partial V_i / \partial w^{\text{fuel}} < 0$ (Prediction 1). The geographic-proximity coefficient is positive and significant (0.354, $t = 2.97$), consistent with the local competitive benefit ($\gamma_G > 0$) that arises when a nearby generator exits and wholesale margins improve for the survivors. The regulatory-proximity coefficient is large but imprecise; we retain it to guard against omitted-variable bias and report a two-channel specification in Appendix D (Table 11), where the qualitative pattern is unchanged. At the six- and twelve-month horizons, the coefficients lose statistical precision under event clustering; the fuel-channel point estimates reverse sign but are not distinguishable from zero, and we regard the longer-horizon estimates as uninformative rather than as evidence of genuine reversal.

Table 2: Channel Decomposition of Spatial Exposure

	$[-1, +3]$	$[-1, +6]$	$[-1, +12]$
w^{geo}	0.354*** (0.119)	0.027 (0.169)	-0.308 (0.283)
w^{reg}	1.127 (0.691)	1.314 (1.045)	3.138 (2.018)
w^{fuel}	-1.497*** (0.474)	0.917 (0.704)	1.364 (1.118)
Same sector	0.021*** (0.005)	0.026*** (0.008)	0.036*** (0.013)
N	72,398	72,398	72,398
R^2	0.0012	0.0006	0.0005

Notes: Monthly CARs with alpha-only classification and vwretd returns. Standard errors are event-clustered (the primary inference standard for the decomposition); two-way clustered versions are reported in Appendix D. Coefficients are estimated jointly in the decomposition regression and are not directly comparable to a composite exposure term. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

To put the magnitudes in context: a 0.10-unit increase in bilateral fuel-similarity exposure corresponds to approximately 15 basis points of differential three-month CARs (-1.497×0.10), while the same increase in geographic proximity contributes 3.5 basis points in the opposite direction (0.354×0.10). The fuel channel is more than four times as large in absolute value. These are not large effects for individual event-firm pairs, and the decomposition explains a modest share of cross-sectional return variation ($R^2 = 0.0012$). While low in absolute terms, this explanatory power is standard for event-time peer-effect regressions where idiosyncratic noise heavily dominates short-horizon returns (e.g., Hou 2007; Cohen and Frazzini 2008). The empirical contribution lies not in the model fit, but in the qualitative channel split: documenting that the same physical event produces opposite-signed repricing along different dimensions of proximity, cumulated across 179 first-mover events.

A natural concern is whether these spatial exposures merely proxy for unobserved firm fundamentals, such as incumbent fossil utilities being structurally larger and more levered. As discussed in Section 2.5, our identification relies on the opposing-sign channel split to filter out common shocks; we further address this concern with firm fixed effects. As we detail in Appendix D (Table 7), introducing cross-sectional controls for size and leverage naturally attenuates the pooled estimates because these variables capture static structural differences between firms. However, when we estimate the model using firm fixed effects to absorb all time-invariant unobserved heterogeneity, the fuel-similarity channel re-emerges with even greater economic magnitude and statistical significance (-2.023, $p < 0.01$), and geographic proximity remains robust (0.472, $p < 0.01$). This confirms that our baseline 15-basis-point magnitude is, in fact, a conservative lower bound; the true spatial contagion operates dynamically through within-firm variation around the restructuring

event, effectively isolating the network shock from static firm characteristics.

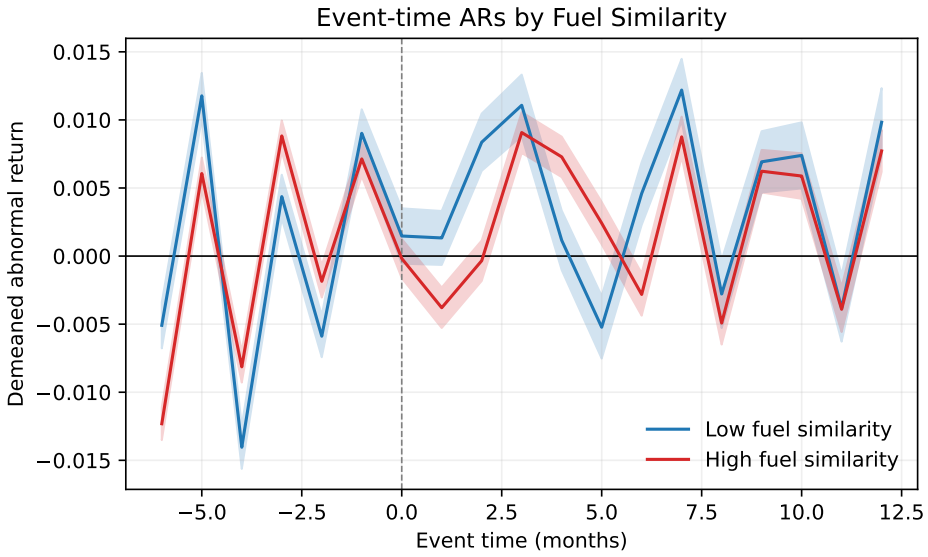


Figure 1: Event-Time Abnormal Returns by Fuel-Similarity Exposure

Notes: Demeaned monthly abnormal returns around first-mover events ($\tau \in [-6, +12]$) for high versus low fuel-similarity neighbours. Returns are market-adjusted using vwretd and demeaned by the pre-event mean for each event–firm pair.

Figure 1 makes the channel split visible in event time. In the six months preceding each retirement event, demeaned CARs for high- and low-fuel-similarity neighbors move in parallel, a key supporting condition for causal inference in this design. The divergence begins at the event month ($\tau = 0$) and widens progressively through the three-month post-event window, with high-fuel-similarity neighbors accumulating negative differential returns relative to their low-similarity counterparts. By the sixth post-event month the gap stabilises and the estimates become imprecise, consistent with the loss of statistical precision reported in Table 2. Two features of the figure deserve emphasis. First, the pre-event parallelism is not an artifact of demeaning: raw CARs for high-exposure neighbors are mildly positive before the event, reflecting selection into the high-exposure group (cleaner firms with higher baseline returns) rather than anticipation of the shock. Second, the gradual widening of the gap, rather than an instantaneous jump, is consistent with the information-diffusion interpretation developed in Section 4: spatial signals are incorporated progressively by segmented investor bases, not absorbed all at once.

We now turn from the first moment to the second. The variance equation (Equation 3) predicts that ambiguous shocks should increase volatility without moving mean returns (Prediction 2), while binding policy shocks should compress the $p_t(1 - p_t)$ uncertainty term, compressing volatility and shifting prices in a predictable direction (Predic-

tion 3). To test these predictions, we construct a measure of realized volatility change: the standard deviation of daily abnormal returns in the twenty trading days following each announcement ($[+1, +20]$), minus the corresponding pre-announcement standard deviation ($[-21, -1]$). We then regress this volatility change on spatial exposure for EIA-860 events and on coal-share intensity for binding phase-out laws, with event-clustered standard errors.

The results, summarised in Table 3, confirm the predicted asymmetry. For EIA-860 retirement filings, the volatility coefficient is positive and highly significant (0.056, $t = 3.89$): firms with greater spatial exposure to the retiring plant experience a measurable increase in return variability in the twenty days following the announcement. The corresponding return coefficient is economically small and statistically indistinguishable from zero (-0.222, $t = -0.17$). This combination, higher volatility with no directional price movement, is the signature of mean-preserving uncertainty. The announcement conveys that a neighbor's asset base is changing, which raises the perceived variance of transition-relevant cash flows for nearby firms, but it does not resolve whether the net effect is positive (through the geographic channel) or negative (through fuel similarity).

Binding phase-out laws show the mirror image: coal-share intensity reduces volatility (-0.130, $t = -2.53$) and is associated with negative abnormal returns ($t = -1.80$, marginally significant). The volatility decline confirms that binding legislation resolves a component of the regulatory uncertainty captured by $p_t(1 - p_t)$ in the framework; the negative returns are suggestive that, once uncertainty is resolved, the market prices the stranding channel.

Table 3: Uncertainty and Resolution: Volatility vs. Returns

Shock type	Volatility change $\sim$ exposure	CAR $\sim$ exposure
EIA-860 retirements	0.056*** (0.014)	-0.222 (1.306)
Binding phase-out laws	-0.130** (0.051)	-18.118* (10.066)

Notes: Daily abnormal returns are computed using vwretd; CAR window is $[-1, +20]$. Volatility change is measured as the post- versus pre-announcement standard deviation of daily abnormal returns. Standard errors (in parentheses) are event-clustered. Binding laws are Tier-1, legally binding phase-out shocks; exposure is measured using continuous firm-level coal-share intensity. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

As a complementary specification, we estimate a panel difference-in-differences model using coal-share intensity as continuous treatment for Tier-1 binding laws. The exposure-by-post coefficient is negative (event-clustered $t = -1.85$, two-way clustered $t = -1.62$), consistent in sign with stranding but below conventional significance thresholds. The attenuation under two-way clustering reflects the small number of independent binding-law events (14 Tier-1 shocks), which limits the effective degrees of freedom for inference. We

report this estimate as suggestive evidence that complements the volatility result rather than as a standalone causal claim; the volatility channel, which is identified with considerably greater statistical precision, carries the primary inferential weight for Prediction 3.

In sum, the empirical record traces a coherent arc from the framework's predictions to the data. At the three-month horizon, spatial transmission splits cleanly by channel: fuel similarity carries a negative signal, geographic proximity a positive one. When the information event is ambiguous, markets respond through volatility rather than prices. When the event removes discretion, volatility falls and prices adjust in the direction predicted by the stranding channel. Robustness checks on window length, clustering, and alternative exposure definitions are reported in Appendix D; the qualitative pattern is stable across specifications.

4. Discussion

The core message of this paper is that transition credibility is a network property subject to severe information frictions. Markets observe physical exposure but do not consistently price commitment quality in real time. Instead, transition shocks transmit through distinct spatial channels and through volatility. This section contextualizes these findings within the literature and outlines three specific, counterintuitive implications for investors and policymakers.

4.1. Information Diffusion and the Limits of Disclosure

The three-month channel split (fuel similarity $= -1.497$, geographic proximity $= +0.354$) and subsequent loss of statistical precision at 6 and 12 months map directly onto the gradual-information-diffusion literature. Hong and Stein (1999) show that when fundamental information diffuses slowly across investors, prices underreact in the short run and correct as information percolates. Our 3-month spatial effects are consistent with the attention-driven return predictability documented by Cohen and Frazzini (2008) for economically linked firms, which concentrates at short horizons and then decays. The larger absolute magnitude of the negative fuel channel aligns with Hou (2007), who finds that negative information diffuses more slowly across linked firms, producing stronger predictability.

More fundamentally, the contrast between volatility spikes following ambiguous retirements and directional repricing following binding laws highlights the structural limits of voluntary disclosure. Under the cheap-talk framework of Crawford and Sobel (1982), because firms' incentives to maintain profitable fossil operations diverge from investors' climate-risk assessments, costless voluntary pledges cannot credibly transmit precise information. Kim and Lyon (2011) document exactly this strategic manipulation in voluntary greenhouse gas registries.

This generates a critical implication for current regulatory efforts. Recent initiatives by financial authorities to standardize ESG disclosures aim to resolve information asymmetry. However, our findings suggest that the primary friction preventing efficient climate risk pricing is not a lack of corporate data, but pervasive political uncertainty. Because markets respond to ambiguous signals with volatility rather than directional repricing (Epstein and Schneider, 2008), even perfectly standardized emissions data will fail to shift the cost of capital if the underlying transition policy remains discretionary. Financial regulation (disclosure) cannot substitute for environmental regulation (binding phase-outs) (Downar et al., 2021; Greenstone et al., 2006).

4.2. The Paradox of Isolated Climate Action

The opposing signs of our spatial channels, mirroring the contagion-versus-competition dichotomy documented by Lang and Stulz (1992), reveal a paradox in targeted climate activism. While fuel similarity transmits competitive stranding (contagion), geographic proximity provides a competitive local benefit ($\gamma_G > 0$). In electricity markets, the retirement of a local generator reduces grid congestion and increases the market power and margins of proximate surviving plants (Davis and Hausman, 2016).

The counterintuitive implication is that isolated, firm-by-firm climate activism may be self-defeating. If an activist campaign successfully forces a utility to retire a coal plant, the resulting reduction in local energy supply inadvertently increases the profitability of neighboring fossil-fuel plants that were not targeted. Because transition risk is a spatial network property, climate interventions that fail to account for geographic grid spillovers risk subsidizing the very emissions they seek to eliminate. Effective transition policy must operate at the network or jurisdictional level to prevent this spatial leakage.

4.3. External Validity and the Illusion of Green Portfolios

The spatial transmission mechanism we document is a specific instance of the general network propagation framework formalized by Acemoglu et al. (2012). Any networked infrastructure sector with hub-and-spoke characteristics, such as shipping, pipelines, or heavy industrial supply chains, is predicted to exhibit analogous spatial transmission. Barrot and Sauvagnat (2016) demonstrate that localized shocks propagate through supply networks, imposing significant sales declines on distant customers. Furthermore, spatial climate risk is demonstrably priced in non-utility asset classes, such as real estate (Bernstein et al., 2019).

However, analyzing spatial risk exclusively through public equity markets exposes a critical limitation: the illusion of the "green" public portfolio. Duchin et al. (2025) document the phenomenon of "brown-spinning": publicly traded firms increasingly divest their most pollutive assets to private, non-ESG-rated buyers. Following divestment, the

public seller's ESG rating improves, yet the total pollution of the underlying physical asset actually increases.

Because transition risk is tied to the physical asset and its spatial network (as our results show), transferring ownership from a public to a private entity does not reduce systemic risk; it merely hides it from public equity markets. Consequently, climate stress tests that rely solely on the balance sheets of publicly listed firms will systematically underestimate systemic vulnerability. Furthermore, derivative markets (Ilhan et al., 2021) and long-term municipal bond markets (Painter, 2020) capture tail risks and spatial immobility dynamics that average equity returns may miss.

5. Conclusion

This paper establishes that transition credibility is a network property. By mapping the spatial transmission of restructuring shocks across global power utilities, we document a short-horizon channel split: fuel similarity drives negative repricing, while geographic proximity produces a positive competitive benefit. The loss of statistical precision at longer horizons is consistent with gradual information diffusion through spatially segmented investor bases.

Crucially, we show that markets process transition credibility through uncertainty rather than immediate repricing. Idiosyncratic retirement announcements raise volatility without moving mean returns. Only when binding phase-out laws remove political discretion does volatility compress, generating directionally negative returns proportional to a firm's fossil exposure.

Three specific implications emerge for policymakers and investors. First, the limits of disclosure: standardizing ESG reporting cannot overcome the credibility gap if the underlying climate policy remains discretionary. Second, the paradox of isolated action: because plant retirements benefit geographically proximate competitors, firm-by-firm divestment campaigns risk subsidizing neighboring fossil assets. Third, the brown-spinning risk: because transition exposure is embedded in physical spatial networks, transferring dirty assets from public markets to private equity hides systemic risk without reducing it. Ultimately, evaluating transition credibility requires looking beyond the isolated firm to the physical network, and looking beyond corporate pledges to the commitment force of binding law.

Future work should extend the spatial framework to credit spreads, sovereign bond markets, and derivative pricing, where tail-risk dynamics may amplify the channels documented here. Theoretically, while our parsimonious framework captures the first-order pricing of policy uncertainty, extending the model to incorporate strategic interactions, dynamic Bayesian learning, and heterogeneous investor beliefs (Baldauf et al., 2020) would further illuminate how spatial signals propagate.

Acknowledgements. The author is affiliated with the Smith School of Enterprise and the Environment, University of Oxford. This research received no external funding. The author declares no financial relationships or conflicts of interest.

References

Daron Acemoglu, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016, 2012. doi: 10.3982/ECTA9623.

Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608, 2015. doi: 10.1257/aer.20130456.

Markus Baldauf, Lorenzo Garlappi, and Constantine Yannelis. Does climate change affect real estate prices? only if you believe in it. *The Review of Financial Studies*, 33(3): 1256–1295, 2020. doi: 10.1093/rfs/hhz073.

Michael Barnett, William Brock, and Lars Peter Hansen. Pricing uncertainty induced by climate change. *The Review of Financial Studies*, 33(3):1024–1066, 2020. doi: 10.1093/rfs/hhz144.

Jean-Noël Barrot and Julien Sauvagnat. Input specificity and the propagation of idiosyncratic shocks in production networks. *The Quarterly Journal of Economics*, 131(3): 1543–1592, 2016. doi: 10.1093/qje/qjw018.

Stefano Battiston, Antoine Mandel, Irene Monasterolo, Franziska Schütze, and Gabriele Visentin. A climate stress-test of the financial system. *Nature Climate Change*, 7(4): 283–288, 2017. doi: 10.1038/nclimate3255.

Stefano Battiston, Irene Monasterolo, Keywan Riahi, and Bas J. van Ruijven. Accounting for finance is key for climate mitigation pathways. *Science*, 372(6545):918–920, 2021. doi: 10.1126/science.abf3877.

William H. Beaver. The information content of annual earnings announcements. *Journal of Accounting Research*, 6:67–92, 1968. doi: 10.2307/2490070. Supplement.

Florian Berg, Julian F. Kölbel, and Roberto Rigobon. Aggregate confusion: The divergence of ESG ratings. *Review of Finance*, 26(6):1315–1344, 2022. doi: 10.1093/rof/rfac033.

Asaf Bernstein, Matthew T. Gustafson, and Ryan Lewis. Disaster on the horizon: The price effect of sea level rise. *Journal of Financial Economics*, 134(2):253–272, 2019. doi: 10.1016/j.jfineco.2019.03.013.

Nicholas Bloom. The impact of uncertainty shocks. *Econometrica*, 77(3):623–685, 2009. doi: 10.3982/ECTA6248.

Patrick Bolton and Marcin Kacperczyk. Do investors care about carbon risk? *Journal of Financial Economics*, 142(2):517–549, 2021. doi: 10.1016/j.jfineco.2021.05.008.

Patrick Bolton and Marcin Kacperczyk. Global pricing of carbon-transition risk. *The Journal of Finance*, 78(6):3677–3754, 2023. doi: 10.1111/jofi.13272.

Patrick Bolton, Morgan Després, Luiz Awazu Pereira da Silva, Frédéric Samama, and Romain Svartzman. *The Green Swan: Central Banking and Financial Stability in the Age of Climate Change*. Bank for International Settlements, 2020.

Kirill Borusyak, Xavier Jaravel, and Jorn Spiess. Revisiting event-study designs: Robust and efficient estimation. *Review of Economic Studies*, 91(6):3253–3285, 2024. doi: 10.1093/restud/rdae007.

Yann Bramoullé, Habiba Djebbari, and Bernard Fortin. Identification of peer effects through social networks. *Journal of Econometrics*, 150(1):41–55, 2009. doi: 10.1016/j.jeconom.2008.12.021.

A. Colin Cameron, Jonah B. Gelbach, and Douglas L. Miller. Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, 29(2):238–249, 2011. doi: 10.1198/jbes.2010.07136.

Hans Bonde Christensen, Luzi Hail, and Christian Leuz. Mandatory CSR and sustainability reporting: Economic analysis and literature review. *Review of Accounting Studies*, 26(3):1176–1248, 2021. doi: 10.1007/s11142-021-09609-5.

Steve Cicala. Imperfect markets versus imperfect regulation in US electricity generation. *American Economic Review*, 112(2):409–441, 2022. doi: 10.1257/aer.20172034.

Lauren Cohen and Andrea Frazzini. Economic links and predictable returns. *The Journal of Finance*, 63(4):1977–2011, 2008. doi: 10.1111/j.1540-6261.2008.01379.x.

Timothy G. Conley. GMM estimation with cross sectional dependence. *Journal of Econometrics*, 92(1):1–45, 1999. doi: 10.1016/S0304-4076(98)00084-0.

Vincent P. Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, 1982. doi: 10.2307/1913390.

Lucas W. Davis and Catherine Hausman. Market impacts of a nuclear power plant closure. *American Economic Journal: Applied Economics*, 8(2):92–122, 2016. doi: 10.1257/app.20140473.

Simon Dietz, Alex Bowen, Charlie Dixon, and Philip Gradwell. 'climate value at risk' of global financial assets. *Nature Climate Change*, 6(7):676–679, 2016. doi: 10.1038/nclimate2972.

Benedikt Downar, Jürgen Ernstberger, Stefan Reichelstein, Sebastian Schwenen, and Aleksandar Zaklan. The impact of carbon disclosure mandates on emissions and financial operating performance. *Review of Accounting Studies*, 26(3):1137–1175, 2021. doi: 10.1007/s11142-021-09611-x.

Ran Duchin, Janet Gao, and Qiping Xu. Sustainability or greenwashing: Evidence from the asset market for industrial pollution. *The Journal of Finance*, 2025. doi: 10.1111/jofi.13416. Forthcoming.

Larry G. Epstein and Martin Schneider. Ambiguity, information quality, and asset pricing. *The Journal of Finance*, 63(1):197–228, 2008. doi: 10.1111/j.1540-6261.2008.01314.x.

Kira R. Fabrizio, Nancy L. Rose, and Catherine D. Wolfram. Do markets reduce costs? assessing the impact of regulatory restructuring on US electric generation efficiency. *American Economic Review*, 97(4):1250–1277, 2007. doi: 10.1257/aer.97.4.1250.

Eugene F. Fama and Kenneth R. French. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56, 1993. doi: 10.1016/0304-405X(93)90023-5.

J. Doyne Farmer, Cameron Hepburn, Matthew C. Ives, Thomas Hale, Thom Wetzer, Penny Mealy, Ryan Rafaty, Sugandha Srivastav, and Rupert Way. Sensitive intervention points in the post-carbon transition. *Science*, 364(6436):132–134, 2019. doi: 10.1126/science.aaw7287.

Stefano Giglio, Bryan Kelly, and Johannes Stroebel. Climate finance. *Annual Review of Financial Economics*, 13(1):15–36, 2021. doi: 10.1146/annurev-financial-102620-103311.

Michael Greenstone, Paul Oyer, and Annette Vissing-Jorgensen. Mandated disclosure, stock returns, and the 1964 securities acts amendments. *The Quarterly Journal of Economics*, 121(2):399–460, 2006. doi: 10.1162/qjec.2006.121.2.399.

William Grieser, Charles J. Hadlock, James LeSage, and Morad Zekhnini. Network effects in corporate financial policies. *Journal of Financial Economics*, 144(1):247–272, 2022. doi: 10.1016/j.jfineco.2021.05.060.

Rebecca N. Hann, Heedong Kim, and Yue Zheng. Intra-industry information transfers: Evidence from changes in implied volatility around earnings announcements. *Review of Accounting Studies*, 24(3):927–971, 2019. doi: 10.1007/s11142-019-09493-2.

Bernard Herskovic. Networks in production: Asset pricing implications. *The Journal of Finance*, 73(4):1785–1818, 2018. doi: 10.1111/jofi.12684.

Gerard Hoberg and Gordon Phillips. Text-based network industries and endogenous product differentiation. *Journal of Political Economy*, 124(5):1423–1465, 2016. doi: 10.1086/688176.

Harrison Hong and Jeremy C. Stein. A unified theory of underreaction, momentum trading, and overreaction in asset markets. *The Journal of Finance*, 54(6):2143–2184, 1999. doi: 10.1111/0022-1082.00184.

Kewei Hou. Industry information diffusion and the lead-lag effect in stock returns. *The Review of Financial Studies*, 20(4):1113–1138, 2007. doi: 10.1093/revfin/hhm003.

Po-Hsuan Hsu, Kai Li, and Chi-Yang Tsou. The pollution premium. *The Journal of Finance*, 78(3):1343–1392, 2023. doi: 10.1111/jofi.13217.

Emirhan Ilhan, Zacharias Sautner, and Grigory Vilkov. Carbon tail risk. *The Review of Financial Studies*, 34(3):1540–1571, 2021. doi: 10.1093/rfs/hhaa071.

Xiangyuan Jiang, Sangwon Kim, and Shuo Lu. Limited accountability and awareness of corporate emissions target outcomes. *Nature Climate Change*, 15:81–87, 2025. doi: 10.1038/s41558-025-02273-8.

Philippe Jorion and Gaiyan Zhang. Good and bad credit contagion: Evidence from credit default swaps. *Journal of Financial Economics*, 84(3):860–883, 2007. doi: 10.1016/j.jfineco.2006.06.001.

Bryan Kelly, Ľuboš Pástor, and Pietro Veronesi. The price of political uncertainty: Theory and evidence from the option market. *The Journal of Finance*, 71(5):2417–2480, 2016. doi: 10.1111/jofi.12406.

Eun-Hee Kim and Thomas P. Lyon. Strategic environmental disclosure: Evidence from the DOE's voluntary greenhouse gas registry. *Journal of Environmental Economics and Management*, 61(3):311–326, 2011. doi: 10.1016/j.jeem.2010.11.001.

Nicolas Koch, Godefroy Grosjean, Sabine Fuss, and Ottmar Edenhofer. Politics matters: Regulatory events as catalysts for price formation under cap-and-trade. *Journal of Environmental Economics and Management*, 78:121–139, 2016. doi: 10.1016/j.jeem.2016.03.004.

Philipp Krueger, Zacharias Sautner, and Laura T. Starks. The importance of climate risks for institutional investors. *The Review of Financial Studies*, 33(3):1067–1111, 2020. doi: 10.1093/rfs/hhz137.

Larry H. P. Lang and René M. Stulz. Contagion and competitive intra-industry effects of bankruptcy announcements: An empirical analysis. *Journal of Financial Economics*, 32(1):45–60, 1992. doi: 10.1016/0304-405X(92)90024-R.

Charles F. Manski. Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, 60(3):531–542, 1993. doi: 10.2307/2298123.

Christophe McGlade and Paul Ekins. The geographical distribution of fossil fuels unused when limiting global warming to 2°C. *Nature*, 517(7533):187–190, 2015. doi: 10.1038/nature14016.

Marcus Painter. An inconvenient cost: The effects of climate change on municipal bonds. *Journal of Financial Economics*, 135(2):468–482, 2020. doi: 10.1016/j.jfineco.2019.06.006.

Ľuboš Pástor and Pietro Veronesi. Uncertainty about government policy and stock prices. *The Journal of Finance*, 67(4):1219–1264, 2012. doi: 10.1111/j.1540-6261.2012.01746.x.

Mitchell A. Petersen. Estimating standard errors in finance panel data sets: Comparing approaches. *The Review of Financial Studies*, 22(1):435–480, 2009. doi: 10.1093/rfs/hhn053.

Christo Pirinsky and Qinghai Wang. Does corporate headquarters location matter for stock returns? *The Journal of Finance*, 61(4):1991–2015, 2006. doi: 10.1111/j.1540-6261.2006.00895.x.

Jonathan Roth. Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights*, 4(3):305–322, 2022. doi: 10.1257/aeri.20210236.

Lee Seltzer, Laura Starks, and Qifei Zhu. Climate regulatory risk and corporate bonds. NBER Working Paper No. 29994, 2022.

Johannes Stroebel and Jeffrey Wurgler. What do you think about climate finance? *Journal of Financial Economics*, 142(2):487–498, 2021. doi: 10.1016/j.jfineco.2021.08.004.

Pietro Veronesi. Stock market overreaction to bad news in good times: A rational expectations equilibrium model. *The Review of Financial Studies*, 12(5):975–1007, 1999. doi: 10.1093/rfs/12.5.975.

A. Channel Decomposition and Primitives

This appendix describes how the three exposure matrices in Equation (1) are constructed.

Geographic proximity (w^{geo}). We compute pairwise distances between firms using plant-level GPS coordinates and define

$$w_{ij}^{\text{geo}} = \frac{\exp(-d_{ij}/h)}{\sum_{k \neq i} \exp(-d_{ik}/h)},$$

where h is the geographic half-life. We correct standard errors for spatial dependence following Conley (1999). This captures shared grids, factor markets, and local competition.

Fuel similarity (w^{fuel}). We build a fuel-mix vector for each firm from GEM plant records (coal, gas, oil, wind, solar, hydro, nuclear). Fuel similarity is measured as the cosine similarity of these vectors and row-normalised across neighbours. This captures competitive stranding among firms with similar fossil exposure.

Regulatory proximity (w^{reg}). We set $w_{ij}^{\text{reg}} = 1$ if firms share a carbon-pricing jurisdiction (ETS membership from the World Bank Carbon Pricing Dashboard; ICAP used as cross-check), and 0 otherwise, then row-normalise. This captures policy spillovers across common regulatory regimes.

These matrices form the empirical channel decomposition used in Section 2.5. The composite exposure is not interpreted as a structural parameter; instead, we focus on channel-specific responses and their horizon-dependence.

B. Derivations

Derivation of Prediction 1

From Equation (2),

$$\frac{\partial V_i}{\partial(\sum_j w_{ij}^{\text{geo}} s_j)} = \gamma_G > 0, \qquad \frac{\partial V_i}{\partial(\sum_j w_{ij}^{\text{fuel}} s_j)} = -p_t \gamma_F < 0,$$

for $p_t \in (0, 1)$. The regulatory term enters additively and does not alter the sign of the geographic or fuel channels. □

Derivation of Prediction 2

With $p_t \in (0, 1)$, Equation (3) implies

$$\frac{\partial \text{Var}_t(V_{i,t})}{\partial s_j} = 2p_t(1 - p_t)\left(\beta\alpha_i + \gamma_F \sum_k w_{ik}^{\text{fuel}} s_k\right)\gamma_F w_{ij}^{\text{fuel}} > 0$$

for exposed firms. Let y denote the announcement signal with likelihoods $L_1 = \Pr(y \mid \omega = 1)$ and $L_0 = \Pr(y \mid \omega = 0)$. Bayes' rule implies

$$p_t' = \Pr(\omega = 1 \mid y) = \frac{p_t L_1}{p_t L_1 + (1 - p_t) L_0}.$$

If the signal is noisy ($L_1 \approx L_0$), then $p_t' - p_t$ is small. The mean effect in Equation (2) remains ambiguous because the positive γ_G and negative γ_F channels offset, while the variance effect above is strictly positive. □

Derivation of Prediction 3

Let $p_{post} > p_{pre}$ and suppose the prior is near maximum uncertainty. Then

$$p_{post}(1 - p_{post}) < p_{pre}(1 - p_{pre}),$$

so volatility in Equation (3) falls. From Equation (2),

$$\Delta V_i = -(p_{post} - p_{pre}) \left(\beta \alpha_i + \gamma_F \sum_j w_{ij}^{\text{fuel}} s_j \right) < 0,$$

which yields negative abnormal returns for coal-exposed firms. □

C. Data Construction

Panel A: Firm Fundamental Parameters (Full Detail)

Table 4: Construction of Firm Fundamental Parameters

Parameter	Measure	Construction and Source
λ (leverage)	Debt / Assets	(DLTT + DLC) / AT. Source: Compustat Global/NA.
ρ (return spread)	Operating ROA (proxy)	OIBDP / AT. Source: Compustat Global/NA.
κ (cash flow adequacy)	Interest coverage (proxy)	OANCF / XINT. Source: Compustat Global/NA.
α (legacy intensity)	Fossil MW / Total MW	Coal + gas + oil capacity as share of total installed capacity. Source: GEM (GeoAsset where available).

Baseline replication uses the Compustat-only proxies shown in Table 4. Preferred definitions use ROIC−WACC and required-capex benchmarks; these extensions rely on CRSP/Datastream, Fitch/Mergent, and IEA benchmarks and are documented below.

All financial variables are measured annually. Following standard practice in the empirical finance literature (Fama and French, 1993), we winsorize all continuous financial ratios at the 1st and 99th percentiles to mitigate the influence of outliers; Table 1 reports raw values. The legacy intensity parameter α is updated when new GEM tracker vintages are available; the current replication uses the latest snapshot, which records plant-level status changes (operating, announced, under construction, retired) with dates.

Panel B: Transformation Delivery (Full Detail)

We also construct physical transformation outcomes for validation and auxiliary analyses. These are not the main dependent variables in the return/volatility tests, but they anchor the exposure measures and provide long-horizon context. We measure:

(i) **Capacity mix shift** for horizons $k \in \{3, 5\}$ years:

$$\Delta\text{Clean}_{i,t} = \frac{(\text{Renewable} + \text{Nuclear MW})_{i,t}}{\text{Total MW}_{i,t}} - \frac{(\text{Renewable} + \text{Nuclear MW})_{i,t-k}}{\text{Total MW}_{i,t-k}}.$$

Source: GEM Global Power Plant Tracker and GeoAsset, which record capacity additions and retirements with dates, fuel types, and ownership. This is our preferred measure because it captures physical transformation that is verifiable and not subject to accounting discretion. In the baseline replication we use the cross-sectional clean share $(1-\alpha)$ as a proxy when panel coverage is incomplete.

(ii) **Emissions intensity change**:

$$\Delta\text{EI}_{i,t} = \frac{\text{Emissions}_{i,t}}{\text{Generation}_{i,t}} - \frac{\text{Emissions}_{i,t-k}}{\text{Generation}_{i,t-k}}.$$

Source: LSEG ESG (via Workspace) provides Scope 1 emissions for listed firms. For firms not covered by LSEG, we construct implied emissions from capacity mix (GEM) $\times$ technology-specific emission factors (IPCC). Generation data from national regulators and IEA World Energy Statistics.

(iii) **Asset retirement events**: Binary indicator for whether firm i retired fossil capacity exceeding 10% of its installed base in period t. Source: GEM Coal Plant Tracker (retirement year; exact dates obtained from GEM Wiki or news where available) and GEM Gas Plant Tracker.

Baseline replication focuses on capacity mix shifts and retirement events; emissions-intensity changes are incorporated when LSEG coverage is sufficient and are reported as a robustness check.

Panel C: Spatial Weight Matrix Construction (Full Detail)

Baseline replication uses a composite of geographic proximity and shared ETS jurisdiction (the $\mathbf{W}^{\text{geo}}$ geographic component and $\mathbf{W}^{\text{reg}}$). Channel decomposition in the main text uses geographic, regulatory, and fuel-similarity components.

$\mathbf{W}^{\text{reg}}$ (Regulatory Proximity) is constructed from two layers:

(i) *Carbon pricing regime*: $w_{ij}^{\text{reg}} = 1$ if firms i and j operate under the same emissions trading system or carbon tax regime. Source: World Bank Carbon Pricing Dashboard (open access); ICAP Emissions Trading Worldwide Status Report used as a cross-check. This creates discrete clusters: EU ETS, UK ETS, RGGI, China ETS, Korean ETS, and unpriced firms.

(ii) *Regulatory framework similarity*: For firms not sharing a carbon pricing regime, a continuous similarity measure based on jurisdiction-level policy indicators in the Climate Change Laws of the World database (Grantham Research Institute/LSE, open access), which codes climate legislation, targets, and policy instruments across 200+ jurisdictions. Pairwise similarity is the cosine distance between jurisdiction policy vectors.

The composite is $w_{ij}^{\text{reg}} = \omega_1 \cdot \text{SharedETS}_{ij} + \omega_2 \cdot \text{PolicySimilarity}_{ij}$, where the ω coefficients are estimated from the data. Row-normalised.

$\mathbf{W}^{\text{geo}}$ (Geographic / Factor Market Overlap) is constructed from plant-level geolocation:

(i) *Geographic proximity*: Using GPS coordinates from GeoAsset and GEM, compute great-circle distance d_{ij} between each pair of firms' nearest plants. Apply inverse-distance weights with exponential decay:

$$w_{ij}^{\text{geo}} = \exp(-d_{ij}/\bar{d})/d_{ij}, \quad \bar{d} = 500/\ln 2.$$

This uses a 500 km half-life and imposes no hard cutoff. It captures competition for local engineering labour, construction capacity, and grid access.

(ii) *Technology overlap*: Firms building the same technology (e.g., both commissioning offshore wind in the same maritime region) compete for technology-specific bottleneck resources (installation vessels, specialised turbines, grid interconnection permits). Constructed from GEM project pipeline data using the same geographic kernel:

$$w_{ij}^{\text{geo,tech}} = \frac{\text{shared technology projects within the kernel}}{\text{total projects of firm } i}.$$

The composite is

$$w_{ij}^{\text{geo}} = (1-\chi)\cdot \text{GeographicProximity}_{ij} + \chi \cdot \text{TechnologyOverlap}_{ij},$$

where $\chi \in [0,1]$ is the technology overlap weight, estimated from the data. Row-normalised.

D. Robustness Checks

We report robustness across return windows and clustering in Table 10 and a two-channel decomposition (geo + fuel) in Table 11. Standard errors are two-way clustered by firm and time following Petersen (2009); Cameron et al. (2011) provide the formal asymptotic justification for multiway clustering when the number of clusters is small. Event-time pre-trends are summarized in the paragraph below.

Orthogonality of the networks. The three spatial weight matrices are not collinear. Table 5 reports pairwise correlations between w^{geo}, w^{fuel}, and w^{reg}. All correlations are below 0.16, validating the exclusion restriction required for network identification (Bramoullé et al., 2009): geographic proximity does not imply fuel similarity, and regulatory proximity does not proxy for either physical channel.

Table 5: Correlation Matrix of Spatial Weights

	w^{geo}	w^{fuel}	w^{reg}
w^{geo}	1.000	0.083	0.155
w^{fuel}	0.083	1.000	0.134
w^{reg}	0.155	0.134	1.000

Strong placebo (shuffled networks). We perform a falsification test by randomly permuting the network topology (assigning each firm a false set of neighbors) 200 times while keeping event dates fixed. Table 6 shows that exposure coefficients collapse toward zero under this permutation, demonstrating that our results are driven by the true spatial topology rather than an unconditional drift in fossil-sector returns.

Table 6: Strong Placebo: Shuffled Exposure Networks (3-Month CAR)

	Mean coefficient	SD across permutations
w^{geo}	-0.031	(0.314)
w^{fuel}	-0.168	(2.309)
w^{reg}	-0.021	(0.805)

Controls and firm fixed effects. Introducing cross-sectional controls (size and leverage) attenuates the fuel-similarity coefficient because these variables absorb structural differences between incumbent fossil utilities (larger, more levered) and non-incumbents. Table 7 shows that when we estimate a within-firm specification with firm fixed effects, the fuel channel re-emerges strongly (-2.023^{***}) and the geographic channel remains robust (0.472^{***}). This confirms that spatial contagion operates through within-firm variation around events once time-invariant firm heterogeneity is removed.[1]

Table 7: Channel Decomposition: Controls Sensitivity (3-Month CARs)

	Size only	Leverage only	Firm FE
w^{geo}	0.276**	0.281**	0.472***
	(0.117)	(0.117)	(0.103)
w^{fuel}	-0.619	-0.295	-2.023***
	(0.473)	(0.501)	(0.513)
w^{reg}	1.014	0.983	0.829
	(0.675)	(0.656)	(0.670)
Same sector	0.019***	0.019***	0.030***
	(0.005)	(0.005)	(0.006)
Log assets	0.003***		
	(0.001)		
Leverage (λ)		0.038**	
		(0.016)	
N	76279	76088	76818
Firm FE	No	No	Yes

Specification progression. To ensure transparency in our spatial decomposition, Table 8 reports the step-by-step inclusion of our spatial weights. The geographic benefit emerges strongly in isolation and remains highly stable, while the fuel-similarity contagion effect reveals its full magnitude only when estimated jointly with geographic proximity and appropriate firm-level controls. Notably, when estimated in isolation, the fuel-similarity coefficient is attenuated and statistically indistinguishable from zero. This reflects classical omitted variable bias: because geographic proximity and fuel similarity are positively correlated, failing to control for the positive geographic competitive benefit mathematically masks the negative contagion effect of fuel similarity.

Bandwidth sensitivity. Our baseline geographic weight matrix uses an exponential decay with a 500 km half-life, calibrated to regional wholesale electricity markets. Table 9 demonstrates that our results are robust to alternative network definitions. Shrinking the

[1]The sample size fluctuates slightly across these specifications (e.g., $N = 76{,}818$ in the fixed-effects model versus $N = 72{,}398$ in the baseline) due to differing data availability requirements for cross-sectional control variables versus within-firm estimation.

Table 8: Specification Progression: 3-Month CARs

	Geo only	Fuel only	Geo + Fuel	Full + controls
w^{geo}	0.333***		0.339***	0.282**
	(0.107)		(0.111)	(0.117)
w^{fuel}		-0.129	-0.235	-0.565
		(0.439)	(0.449)	(0.497)
w^{reg}				0.991
				(0.665)
Same sector				0.018***
				(0.005)
Log assets				0.003***
				(0.001)
Leverage (λ)				0.036**
				(0.015)
Return spread (ρ)				-0.017
				(0.014)
N	76819	76819	76819	75956

half-life to 250 km attenuates the effect as the spatial kernel becomes too restrictive to capture grid-wide wholesale dynamics. Expanding the half-life to 1,000 km strengthens the geographic coefficient (0.481***), consistent with the massive geographic footprint of interconnected transmission systems (e.g., PJM or the European synchronous grid).

Table 9: Bandwidth Sensitivity (3-Month CARs)

	Half-life 250 km	Half-life 1000 km
w^{geo}	0.140	0.481***
	(0.086)	(0.164)
w^{fuel}	-0.200	-0.272
	(0.472)	(0.470)
w^{reg}	1.056	0.942
	(0.671)	(0.657)
Same sector	0.020***	0.020***
	(0.005)	(0.005)
N	76819	76819

Event-time pre-trends. Event-time analysis of neighbour CARs around first-mover retirements shows that demeaned cumulative abnormal returns are approximately flat in the pre-event window ($\tau \in [-6, -1]$) for both high-exposure and low-exposure neighbour groups, supporting the parallel-trends assumption. We verify the absence of pre-trends following the recommendations of Roth (2022), who demonstrates that conventional pre-trend tests have low power against relevant violations and can induce distortions in post-

treatment estimates; the staggered-adoption estimator of Borusyak et al. (2024) is reported as an additional robustness check. Raw (undemeaned) CARs for high-exposure neighbours are mildly positive pre-event, which reflects selection into the high-exposure group (cleaner firms with higher baseline returns) rather than anticipation of the retirement shock.

Table 10: Robustness: Exposure Coefficient Across Windows (VWRETD)

Window	Event-clustered	Two-way clustered	N	R^2
Daily $[-1, +20]$	-2.081*** (0.652)	-2.081 (1.832)	62,022	0.0010
Monthly $[-1, +3]$	0.442*** (0.140)	—	72,398	0.0009
Monthly $[-1, +6]$	0.218 (0.187)	—	72,398	0.0004
Monthly $[-1, +12]$	-0.140 (0.283)	-0.140 (0.600)	72,398	0.0003
Monthly $[-1, +12]$ (bivariate)	-0.063 (0.289)	—	72,398	0.0000

Notes: Exposure coefficients using market-adjusted returns (vwretd) and alpha-only classification. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Two-way clustering is only computed for the primary $[-1, +12]$ window.

Table 11: Channel Decomposition Robustness: Geo+Fuel Only (3-Month Window)

	Event-clustered	N
w^{geo}	0.430*** (0.126)	
w^{fuel}	-1.281*** (0.458)	
Same sector	0.021*** (0.005)	
N	72,398	
R^2	0.0009	

Notes: Monthly CARs in the $[-1, +3]$ window using the two-channel specification (geo + fuel + same-sector). Standard errors clustered by event. This specification omits regulatory proximity to assess sensitivity of the short-horizon split.