Silent Costs: Firm Biodiversity Pressure and Earnings Expectations

Atreya Dey

University of Cambridge, Judge Business School

Abstract

I construct a biodiversity pressure index from citizen science data matched to industrial facility locations and show that facility openings reduce local bird abundance by 6–9%. Among small and mid-cap firms, long-short portfolios sorted on biodiversity pressure earn factor-adjusted alphas of 79 and 56 basis points per month, respectively. The premium reflects forecast errors because investors misinterpret production costs tied to ecologically sensitive locations as evidence of operational decline. Prices correct when earnings reveal that these costs are stable rather than increasing, and 40% of the small-firm premium materializes around earnings announcements. These errors persist because firms exerting the greatest biodiversity pressure are the least likely to disclose their footprints, and ESG raters assign them the lowest biodiversity exposure scores.

Keywords: biodiversity, equity returns, forecast errors, information frictions, sustainable finance
JEL Codes: G12, G14, G41, Q54, Q57

The author's email address is a.dey@jbs.cam.ac.uk. I am grateful for valuable discussions and feedback from Sarah Kroechert, Raghu Rau, Diego Guisande (discussant), Luca Taschini, Caroline Flammer, Glen Gostlow, and Tim Newbold. I also thank seminar participants at the University of Cambridge, Judge Business School Brown Bag Seminar. I acknowledge funding from the Cambridge Endowment for Research in Finance.

1 Introduction

Firms both depend on functioning ecosystems and impose ecological pressure on the places where they operate, linking local environmental conditions to firm value (Gjerde et al., 2026). Yet the biodiversity lost from firm operations is difficult to attribute to individual firms, lacking the informational infrastructure that has developed around carbon emissions—a comparable environmental externality. Carbon is fungible and aggregates into standardized corporate disclosures, but the pressure firms exert on biodiversity depends on local ecosystem conditions and resists similar standardization (Giglio et al., 2024).[1] This tension between financial materiality and informational frictions raises the question of whether firms' local ecological footprints are reflected in the cross-section of stock returns.

One possibility is that biodiversity pressure carries a risk premium. Firms with larger local ecological footprints may face higher expected costs from ecosystem disruption, regulatory intervention, litigation, and remediation (Akey and Appel, 2021; He et al., 2020). Akbari et al. (2026) document that protected-area designations cause production cuts at affected facilities, and emerging nature-related disclosure frameworks signal increasing scrutiny of firms operating in sensitive ecosystems (Taskforce on Nature-Related Financial Disclosures, 2025). If investors recognize these exposures, they may demand compensation for holding firms with larger ecological footprints, generating a positive relation between biodiversity pressure and expected returns. I refer to this as the biodiversity risk premium hypothesis.

An alternative possibility is that markets form biased expectations about biodiversity-driven production costs. Firms with greater dependence on ecosystem services may face location-bound production costs that are difficult to identify (Garel et al., 2026). Investors who lack information regarding the ecological context may rely on simplified models that omit these location-specific drivers of profitability (Hong et al., 2007; Schwartzstein, 2014), generating biased earnings forecasts. Comparable patterns appear when investors form ex-

[1] Biodiversity spans several dimensions (species richness, genetic diversity, ecosystem integrity, functional diversity), each requiring distinct measurement approaches. Carbon emissions, by contrast, aggregate into a single scalar such as CO_2 equivalents (IPBES, 2019).

pectations about other unobserved firm attributes. Employee satisfaction (Edmans, 2011) and physical climate exposure (Pankratz et al., 2023) both generate earnings surprises that analysts fail to anticipate. Under this forecast-error hypothesis, firms with larger ecological footprints generate positive earnings surprises, and returns concentrate around earnings announcements rather than accruing smoothly.

To test these hypotheses, I construct a biodiversity pressure index that combines bird abundance—bird counts from eBird citizen science data, adjusted for observer effort—around each firm's industrial facilities with financial data, covering 1,873 US firms from 1988 to 2024. The index reflects each firm's ongoing ecological footprint per dollar of assets and varies continuously across firms and over time. Using satellite-implied construction dates in a staggered difference-in-differences design, I show that new industrial facilities reduce local bird abundance by 6–9%. The effect attenuates with distance from the facility and is stronger in years with higher toxic releases, two patterns that establish the index as a measure of real ecological damage.[2]

Using this index, I document four sets of results. First, I find a biodiversity premium concentrated in small and mid-cap firms. In conditional double sorts on size and biodiversity pressure, the high-minus-low spread earns a Fama-French three-factor alpha of 79 basis points per month in small firms and 56 basis points in mid-caps, before disappearing in large firms. The premium is distinct from that associated with pollution intensity, as biodiversity pressure retains a large and significant coefficient after controlling for toxic releases in a joint regression.[3] The premium also survives controls for institutional ownership, analyst coverage, board governance, liquidity, idiosyncratic volatility, and short interest. I also show that unscaled bird abundance predicts returns independently of firm size, and randomization

[2]Dose-response tests confirm that years with higher toxic releases are associated with lower subsequent bird abundance. Placebo tests using randomized construction dates confirm that facility timing drives the effect, displacement checks rule out spatial reshuffling of bird populations, and elevated lead concentrations near facilities provide a chemical pathway linking toxic releases to the ecological damage.

[3]Although the pressure index and toxic releases intensity share the same firms and source data from the EPA Toxics Release Inventory (TRI), they capture different dimensions of environmental risk. Releases measures a standardized quantity that investors can compare; the pressure index measures the environmental context of the release (Hsu et al., 2023).

inference confirms the biodiversity numerator drives the effect.

Second, firms with larger biodiversity footprints have higher production costs. A one-standard-deviation increase in the biodiversity pressure index raises cost of goods sold by 0.40 percentage points of sales, with no corresponding effect on selling and administrative expenses. The cost penalty is stable year over year, consistent with a persistent location-bound feature of the firm's operating environment rather than a transitory shock. The cost rises only in production inputs, ruling out a demand-side channel in which consumers cut spending in response to environmental news and firms raise selling expenses to retain them (Derrien et al., 2025). The penalty likely reflects the operational difficulties of producing in ecologically sensitive areas rather than regulatory enforcement.[4] Conditional on emissions, Environmental Protection Agency inspections, informal actions, and penalties are all comparable across high- and low-footprint firms. Unlike the regulatory channels documented by Chen (2025) and Akbari et al. (2026), where discrete policy events trigger production adjustments, the cost penalty is independent of enforcement action.

Third, prices adjust when earnings reveal information about these operational costs. Analysts are systematically pessimistic about firms with larger ecological footprints: a one-standard-deviation increase in the index is associated with a 6.3 percentage point lower 12-month implied return from analyst price targets. Yet these firms repeatedly beat pessimistic forecasts, generating positive standardized earnings surprises against both seasonal random-walk and analyst-consensus benchmarks. The return correction is concentrated around the same information events. Four quarterly announcement windows account for 40% of the small-firm alpha, in line with the concentration documented for other anomalies (Engelberg et al., 2018). This concentration around scheduled earnings news is difficult to reconcile with compensation for a recognized risk, but follows naturally from systematic errors in earnings expectations. The correction begins at earnings announcements and continues in small firms

[4]Potential frictions include constraints on water and energy inputs, longer supply routes around environmentally protected zones, and site-specific compliance costs (wetlands, endangered-species protections). Although I cannot isolate which channel drives the elevated cost of goods sold, the production-side concentration of the response is consistent with each.

over the following weeks (Hong and Stein, 1999). Among small firms, the return spread between high- and low-footprint portfolios grows from 0.75% in the announcement window to 2.47% over 60 trading days, with no comparable drift among mid-cap and large firms.

Fourth, the premium persists because sources of financial information treat biodiversity-driven production costs as indistinguishable from other operational costs. Standard environmental proxies work for carbon because emissions can be standardized across firms. However, they fail for biodiversity because ecological pressure depends on local ecosystem conditions, and firms themselves lack tools to assess nature-risk impacts on cash flows (Gjerde et al., 2026). Fewer than 2% of 10-K filings mention biodiversity (Giglio et al., 2026), and firms with the largest biodiversity footprints are significantly less likely to voluntarily report biodiversity risks to the Carbon Disclosure Project. MSCI ratings capture environmental risk unrelated to the biodiversity footprint, and biodiversity-related media coverage is unrelated to actual biodiversity pressure. MSCI's dedicated biodiversity sub-score is driven by voluntary CDP disclosure and is negatively correlated with actual ecological exposure, so the most exposed firms receive the highest safety ratings. When enforcement actions do reveal a firm's biodiversity exposure, the market absorbs the information slowly. Among penalized firms, the high-minus-low BPI spread reaches -7.78% over 60 trading days.

The results support the forecast-error hypothesis. Lower margins for high-footprint firms appear as operational decline to investors, who set earnings forecasts too low. These firms then beat the forecasts, similar to the value-stock anomaly of La Porta et al. (1997). Here, however, the error stems from an unobserved environmental attribute rather than the extrapolation of past growth. The premium persists across varying levels of analyst coverage, suggesting that the issue lies in the absence of a framework connecting environmental conditions to production costs rather than in a lack of intermediary resources.

I contribute to the literature on biodiversity and financial markets. Dimson et al. (2026) and Giglio et al. (2026) document the financial materiality of biodiversity risk. Existing corporate-level measures draw on supply-chain footprints (Garel et al., 2024), dependence

on ecosystem services (Garel et al., 2026), and country-level biodiversity indices (Creti et al., 2024). The supply-chain measure generates a return premium after the 2021 Kunming Declaration, and the country-level index predicts firm profitability and returns across 35 countries. My measure proxies for local ecological conditions around individual facilities, validated causally against field observations, with a premium that exists throughout the 1988–2024 sample rather than only after a salient policy event. I also trace the full economic mechanism from ecology through operating costs to asset prices, and identify the information intermediation failure that sustains the premium. The premium is priced separately from the emissions intensity studied by Hsu et al. (2023) and Bolton and Kacperczyk (2021), with the strongest incremental content in firms with low aggregate emissions.

I also link environmental exposures to corporate outcomes. Akey and Appel (2021) show that pollution liability affects firm investment, Chen (2025) documents the reallocation costs of environmental regulation, and Akbari et al. (2026) find that protected area designations prompt production cuts. I identify a cost channel that is independent of regulatory designation. Firms bear elevated production costs from their biodiversity footprint whether or not the surrounding area receives formal protection. Without a discrete policy event to observe, the cost enters the financial statements as elevated cost-of-goods-sold, indistinguishable from any other production expense. Unlike the channel documented by Derrien et al. (2025), in which salient ESG news reduces consumer demand, the cost is production-driven rather than demand-driven. Active pollution prevention can partially offset the cost burden, but non-remediating firms bear the full pricing penalty.

Finally, I bridge biodiversity measurement (Dimson et al., 2026) and the literature on financial information production. Berg et al. (2022) show that ESG rating divergence stems from measurement differences across raters, and Avramov et al. (2022) demonstrate that this disagreement generates its own pricing effect. For biodiversity, the problem is more fundamental than disagreement. Existing frameworks capture emissions and governance without extending to the biodiversity context of production. ESG raters, corporate disclosures, and

financial analysts each assign scores or forecasts that are unrelated to actual biodiversity exposure, and surveys confirm that companies themselves lack the data infrastructure to assess nature risks (Gjerde et al., 2026). Analyst coverage does not moderate the premium (Hirshleifer et al., 2013), confirming that more intermediary resources cannot resolve the information friction when the missing information is ecological rather than financial.

Section 2 details the data and index construction. Section 3 validates the index against biological field data using a staggered difference-in-differences design. Section 4 presents the biodiversity premium and its relationship to emissions intensity. Section 5 traces the cash-flow expectations channel, from operating costs through earnings surprises to the role of abatement. Section 6 documents the information gap that sustains the premium and tests what happens when enforcement relaxes these frictions. Section 7 concludes.

2 Data

The Biodiversity Pressure Index combines citizen science bird abundance observations around industrial facilities with financial data to measure the extensiveness of a firm's ecological footprint per dollar of assets. I match this measure to monthly stock returns, quarterly financial statements, analyst forecasts, and enforcement records. Table 1 defines all variables and Table 2 reports summary statistics.

2.1 Measuring biodiversity pressure

I introduce the Biodiversity Pressure Index (BPI), which scales bird abundance around a firm's industrial facilities by total assets to measure its ecological footprint per dollar of assets. Its construction begins with the precise geographic location of industrial facilities, which I obtain from the Toxic Release Inventory (TRI). Established by the Emergency Planning and Community Right-to-Know Act of 1986, the TRI provides the most comprehensive public registry of facility-level industrial activity in the United States, containing annual re-

ports from approximately 21,000 facilities on over 800 listed chemicals (U.S. Environmental Protection Agency, 2023).

To quantify the ecological context of these facilities, I draw on bird abundance data from eBird, a citizen science platform operated by the Cornell Lab of Ornithology (Fink et al., 2021). I use structured checklists from the eBird Basic Dataset, submitted by volunteer observers recording species identities and counts, covering 1985 through 2024. Coverage grows from approximately 7,600 checklists per year in the mid-1980s to over 6.5 million by 2024, expanding after eBird's online launch in 2002. The numerator of BPI is time-varying throughout the sample period, constructed from checklist data available in each year rather than backfilled from a later cross-sectional surface. I treat facility-years with no nearby checklists as missing and exclude them from the firm-level sum. Because the denominator (total assets) includes all facilities regardless of checklist coverage, firms with a higher fraction of unobserved facilities receive mechanically lower BPI scores. The Fama-MacBeth cross-sectional standardization within each vintage year absorbs the resulting cross-sectional variation, and the post-2002 subsample test (Internet Appendix Table IA.8) confirms that results strengthen when coverage is densest. Early coverage is sparser, which adds noise to early-sample estimates, but using only contemporaneous data avoids backfilling bird counts from later years.

Raw checklist totals reflect observer effort and timing as much as underlying bird populations, so I implement an explicit effort-adjustment procedure following Hogan et al. (2025). At the individual checklist level, I partial out the number of observation hours, the number of observers, and a full set of hour-of-day and state-by-month fixed effects from log total bird counts. The residual captures within-state relative abundance after removing predictable variation in detectability. Because the state-by-month fixed effects absorb absolute cross-state differences in biological richness, BPI measures local deviations from each state's seasonal average rather than absolute ecological levels. For each facility, ring, and year, I average these effort-adjusted log counts across all checklists falling within the ring and ex-

ponentiate back to levels to obtain an effort-adjusted abundance measure. Appendix IA.1 provides the full specification and fixed-effect structure. This removes predictable variation in how many birds observers detect, leaving differences that reflect local abundance.

I aggregate bird abundance within a system of concentric rings around each TRI facility, distinguishing between a "treatment" ring (1 to 14 kilometers) where industrial activity has a direct effect and a "control" ring (15 to 30 kilometers) that captures the biodiversity potential of the surrounding landscape. In Section 3, I validate this structure, showing that a staggered difference-in-differences design confirms a 6% to 9% decline in bird abundance within the treatment ring relative to the control ring after a facility is constructed.

Although the outer control ring appears in both the ecological validation and the index construction, the two uses are conceptually consistent. In the difference-in-differences design of Section 3, it serves as a counterfactual for what bird populations would have looked like absent the facility because at 15 to 30 kilometers it lies far enough from operations to remain unaffected by direct industrial impact. That same insulation from damage makes it the appropriate measure of the ecological endowment a firm has chosen to locate near, and therefore the potential that industrial activity places at risk.

The final index uses bird abundance from the outer control ring as its numerator, capturing the ecological richness of the landscape surrounding a firm's facilities, scaled by total assets to produce an intensity measure of exposure per dollar of capital. The Biodiversity Pressure Index is

$$\text{BPI}_{i,t} = \log\left(\frac{\text{Bird Abundance}_{i,t}}{\text{Total Assets}_{i,t}}\right), \tag{1}$$

where Bird Abundance sums across all facilities operated by firm i in year t the effort-adjusted abundance in each facility's 15–30 km control ring, and Total Assets is book assets from Compustat. Each facility's contribution is the exponentiated mean of effort-adjusted log bird counts across checklists within the ring, so the firm-level numerator aggregates these facility-level measures additively. BPI is high when a firm's facilities sit in an ecologically rich landscape relative to its asset base, and low when the surrounding landscape is degraded

or the asset base is large, capturing biodiversity at risk per dollar of capital deployed.

BPI scales with a firm's operational footprint in the same way that emissions intensity scales with total pollution. When facilities are geographically clustered, their control rings overlap and nearby checklists contribute to multiple facility measures. Clustered operations place cumulative pressure on the same ecosystem, so the additive structure is appropriate. BPI correlates -0.606 with log market capitalization because smaller firms carry higher ecological exposure per dollar of assets, a pattern that reinforces the information gap documented in this paper. All return tests condition on size, and a denominator decomposition confirms that both log(birds) and log(assets) independently predict returns (Appendix A.2).

To move from the facility to the firm level, I sum these effort-adjusted abundance measures across all facilities operated by a firm in a given year. I then match this annual BPI measure to monthly returns using a single timing rule applied uniformly across all tests in the paper. TRI data for reporting year Y are published in October of year $Y+1$, so the BPI assigned to any return, earnings announcement, or operating outcome occurring in calendar month m uses vintage $Y = \text{year}(m) - 2$ for months January through October and vintage $Y = \text{year}(m) - 1$ for months November and December. This ensures that only publicly available data enter each test. The same vintage rule applies to Releases. Portfolio sorts, Fama-MacBeth regressions, quarterly profitability tests, and enforcement event studies all use this convention, so the assigned BPI is between one and two years stale by construction. In the final two months of the calendar year, the current vintage is already public, so this rule is deliberately conservative rather than a strict informational constraint. Zhang (2025) shows that apparent carbon return predictability vanishes once emissions data release lags are properly accounted for; the deliberate staleness of BPI addresses this concern by construction. The BPI premium is robust to alternative construction choices and to conditioning on potential confounds. Alternative denominators and unscaled numerators (Appendix A.2), sorts within industry-size bins (Appendix IA.4), and joint control for emissions intensity (Appendix A.3) all yield qualitatively similar return spreads.

2.2 Emissions intensity

To distinguish ecological context from pollution output and separate BPI from the pollution premium documented by Hsu et al. (2023), I construct a parallel measure of emissions intensity following the standard proxy in the literature (Hsu et al., 2023; Chen, 2025; Akey and Appel, 2021). This measure, which I call Releases, is defined as

$$\text{Releases}_{i,t} = \log\left(\frac{\text{Total TRI Releases}_{i,t}}{\text{Total Assets}_{i,t}}\right), \tag{2}$$

where Total TRI Releases is the sum of all reported air, water, and ground releases in pounds for firm i in year t. I deliberately mirror BPI's construction in data source, denominator, and vintage timing, which makes their moderate pairwise correlation of $r = 0.311$ notable. The two measures overlap in data source but capture different aspects of a facility's environmental footprint. While some heavily polluting facilities are located in rich ecosystems, many high-BPI firms emit little, and many high-Releases firms operate in ecologically degraded areas.

2.3 Financial data and controls

I merge the BPI and Releases measures into a comprehensive monthly panel of U.S. public equities. I obtain stock returns for all common shares (share codes 10 and 11) on major exchanges from CRSP and all balance sheet information from the Compustat quarterly files, linking the two through the CCM linking table. The final sample contains 235,586 firm-month observations for 1,873 unique firms from October 1988 through December 2024. Firm coverage peaks at 732 firms around 2000 and declines to 502 firms by 2024, reflecting consolidation in TRI-reporting industries. The Fama-MacBeth estimation sample, which requires non-missing values for all controls, contains 194,036 firm-months spanning 428 calendar months with an average of 453 firms per cross-section. I classify firms into Fama-French 12 industries using SIC codes.

The baseline set of control variables includes log market equity, log book-to-market equity,

the prior 12-month return (momentum, skipping the most recent month), return on assets, leverage, and idiosyncratic volatility from a Fama-French three-factor model estimated over the prior 60 trading days. These controls are added progressively across specifications to show the sensitivity of the BPI coefficient to each addition.

2.4 Enforcement, disclosure, and intermediary data

I draw on five additional datasets to probe the mechanisms through which firms realize these ecological costs and investors learn about them.

First, to capture the direct realization of legal liability, I obtain the universe of federal enforcement actions from the EPA's Enforcement and Compliance History Online (ECHO) database.[5] These records detail case timing, linked facilities, and assessed penalties, allowing me to identify 4,328 distinct enforcement events affecting 682 publicly traded firms in my sample, of which 3,543 events affecting 583 firms occur after 2000 and form the primary event study sample.

Second, to examine the role of public attention, I source 12,092 biodiversity-related news events from RepRisk, a commercial provider that screens media, regulatory filings, and NGO reports for environmental, social, and governance incidents. RepRisk assigns a severity score (1 to 3) to each incident based on the novelty and consequences of the event.

Third, I investigate corporate disclosure using firm-year scores from Giglio et al. (2026), who count biodiversity keywords in 10-K filings. 1.8% of firm-years in the sample contain any mention of biodiversity. I also obtain voluntary climate disclosure participation data from CDP (formerly the Carbon Disclosure Project) to distinguish general environmental transparency from biodiversity-specific disclosure.

Fourth, to assess the role of specialized information intermediaries, I use ESG ratings from MSCI. The matched sample covers 762 firms from 2007 through 2024 and includes the firm's overall industry-adjusted score, its environmental pillar score, and its biodiversity and

[5] ECHO data are publicly available at https://echo.epa.gov. I use the ICIS-FE&C bulk download files.

land use sub-score (available for approximately 110 firms).

Fifth, I measure shifts in aggregate public attention using the New York Times daily article index constructed by Giglio et al. (2026), which counts articles mentioning biodiversity-related terms.

3 Ecological Validation

For the Biodiversity Pressure Index to represent a priced financial risk, it must first correspond to a *real* ecological footprint. I test whether the industrial facilities in my sample generate measurable damage to their local ecosystems using a staggered difference-in-differences design that measures the evolution of bird populations around the construction of new emitting facilities.

A central challenge in this exercise is that a facility's construction date often precedes its first regulatory filing by several years. I overcome this measurement problem by extracting construction dates directly from satellite imagery. Using annual data from the National Land Cover Database (NLCD), I pinpoint the year of construction by detecting discrete jumps in impervious surfaces (new rooftops and pavement) within a 500-meter radius of each facility's reported coordinates. This approach anchors the analysis to the physical construction date and Figure A.1 in the Appendix validates this procedure. Average impervious surface growth spikes at the detected construction year and is near zero in surrounding years. Visual inspection of high-resolution aerial imagery confirms that these statistical jumps correspond to physical land clearing and infrastructure development (Appendix Figure A.2).

To isolate a facility's causal footprint from broader regional trends, my design compares changes in bird populations across two concentric rings. The inner, "treated" ring, spanning from one to X kilometers, captures the area of direct impact. An outer, "control" ring, from $X + 1$ to 30 kilometers, serves as a contemporaneous benchmark for what would have happened in the absence of the facility. I vary the treatment threshold X from 5 to 24

kilometers, which simultaneously tests for spatial attenuation and ensures that any contamination of the control ring by spillovers at the boundary is directly observable as a function of X. I use the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), which is robust to biases of traditional two-way fixed effects models when treatment timing is heterogeneous. For each facility f, I define an outcome Y_{ut} as log effort-adjusted bird abundance (or richness) in ring u and year t. The inner ring receives treatment at the facility's satellite-inferred construction year $\hat{T}_f$; the outer ring is never treated. I estimate the average treatment effect on the treated. The estimand is

$$\mathrm{ATT}(e) \;=\; \mathbb{E}\Big[\, \Delta Y_{f,t}^{\mathrm{Inner}} - \Delta Y_{f,t}^{\mathrm{Outer}} \;\Big|\; t - \hat{T}_f = e \,\Big], \tag{3}$$

where ΔY denotes the change in log effort-adjusted bird abundance relative to the pre-construction period, superscripts identify the ring, and e is the event-time relative to construction. The Callaway and Sant'Anna (2021) estimator recovers this quantity through cohort-time-specific comparisons against not-yet-treated units, aggregated using cohort-size weighting. I cluster standard errors at the facility level.

Figure 1 presents the estimated DiD at every integer threshold from 5 to 24 kilometers and plots the ATT. Both abundance and richness show negative effects that attenuate monotonically as the treatment ring expands to encompass less affected areas. The abundance effect is largest at short radii and becomes statistically insignificant around 19–20 kilometers. Richness follows the same spatial pattern at roughly half the magnitude. This monotonic decay rules out a simple regional trend and supports a site-specific externality that dissipates with distance from the facility.

I select 14 kilometers as the treatment threshold for the main specification, as it falls in the center of the significant range, avoids the noisier short-radius estimates where few checklists drive large standard errors, and leaves a buffer between the treatment and control rings. Table 3 reports the results for this threshold across four progressively controlled

specifications. For each specification, I report the pre-construction average ($e < 0$) as a diagnostic for parallel trends and the post-construction average ($e \geq 0$) as the main causal estimate. Bird abundance decreases by 6% to 9% relative to the outer control ring (Panel A). Species richness falls by 3% to 5%, roughly half the abundance decline (Panel B). Figure 2 traces the event-time path, confirming that the decline begins at construction and persists through the post-treatment window.[6]

Within a given facility over time, years with higher reported TRI emissions are associated with lower proximate bird abundance, even after absorbing all time-invariant facility characteristics with fixed effects (Internet Appendix Figure IA.1), indicating that ecological degradation persists as an ongoing byproduct of industrial activity rather than a one-time consequence of initial land use. The urbanization control in all specifications uses impervious surface growth in the outer (15–30 km) control ring, which proxies for regional development trends without absorbing the facility's own physical footprint. Results are invariant to using inner-ring, full-circle, or decomposed urbanization controls.

Does this result reflect the ring design rather than real ecological damage, do facilities simply displace birds rather than reduce their populations, or could differential pre-treatment trends inflate the estimates? I find no evidence for these alternative explanations. A randomization inference exercise permutes facility construction dates 1,000 times at each distance threshold. Figure 3 overlays the actual estimates on these placebo distributions. The placebo ATTs cluster around zero at every threshold, while the actual abundance estimates diverge from approximately 12 km onward. The actual abundance effect crosses the 5% significance threshold at 12–13 km and remains significant through 17 km, outside the 5th percentile of the placebo distribution. Richness follows the same pattern at smaller magnitudes, crossing the threshold around 13 km and remaining significant through 21 km. Facility construction

[6] A concern with citizen science data is that facility construction may alter where birders choose to survey, independently of ecological change. I test this by replacing the abundance outcome with log(checklist count + 1) and re-estimating the primary specification. The ATT on checklist density is -0.024 ($z = -1.23$), economically negligible relative to the 6–9% abundance decline. Observer spatial sorting does not drive the result.

timing rather than ring geometry drives the result, and a displacement test finds no corresponding increase in bird populations in the areas immediately surrounding the impact zone (Appendix Figure A.3). The losses in the near-field are not offset by gains farther out, ruling out spatial reshuffling as an alternative explanation. The concern raised by Sager and Singer (2025), that high-abundance areas may converge toward a lower mean regardless of treatment, is mitigated here because inner and outer rings show statistically indistinguishable pre-construction abundance ($t = 0.88$). Column 4 of Table 3 adds observer effort controls (log observers, log hours) to the doubly-robust estimator, and the post-construction ATT strengthens to -9.2%, confirming that the result reflects facility-driven ecological damage rather than observer sorting or mean reversion.

A direct environmental pathway links these ecological losses to industrial chemical releases. TRI facilities measurably increase local lead concentrations in ambient water by 1.7 to 2.3 μg/L within 5 to 10 kilometers, an effect that also attenuates with distance (Internet Appendix Figure IA.2). A similar randomization inference exercise confirms that this lead contamination effect lies outside the placebo distribution (Figure IA.3).

4 The Biodiversity Premium

4.1 Baseline return premium

I test the relationship between ecological pressure and cross-sectional stock returns using conditional double sorts on size and BPI, complemented by Fama-MacBeth regressions (Fama and MacBeth, 1973). Table 4 reports returns for portfolios formed on conditional double sorts of market capitalization and BPI, using the vintage-year timing from Section 2 so that all portfolio assignments rely on data publicly available before the return month. In the universe of small firms, the high-minus-low BPI spread earns a monthly Fama-French three-factor alpha of 0.79% (roughly 9% per year). The premium remains economically large in mid-cap firms at 0.56% per month before falling to zero in large firms, where the spread is

statistically and economically indistinguishable from zero.

The result is robust to augmented factor models including liquidity. Adding the Pástor et al. (2021) green factor to the Fama and French (2015) five-factor model leaves the premium unchanged, confirming that the ecological footprint premium is distinct from the return spread associated with ESG tilts. Fama-MacBeth regressions reinforce this portfolio evidence (Table 5). A specification with only industry fixed effects yields a BPI coefficient of 0.13% per month. Adding book-to-market and momentum controls nearly doubles the coefficient to 0.25%, as these characteristics correlate with BPI and partially offset the return spread in the univariate specification. With the full set of controls, the premium is roughly 0.17% per month (roughly 2% per year), confirming that the biodiversity signal remains distinct from known return predictors. Adding the modified Altman (1968) Z-score as a distress control, following the approach in Akey and Appel (2021), leaves the BPI coefficient unchanged at 0.17%, ruling out the possibility that the premium compensates for financial distress risk.

As noted in Section 2, BPI correlates with firm size, raising the concern that the premium reflects size rather than ecology. Bird abundance alone, without asset scaling, produces the same return pattern (Appendix A.2), confirming that the ecological numerator rather than the asset denominator drives the premium. Replacing total assets with sales or PP&E as the scaling denominator yields comparable premiums (Table A.1), and conditional double sorts on size $\times$ BPI_{sales} terciles produce pooled FF3 alphas of 33 bps/month. A randomization test shuffles bird counts across firms within the same industry and vintage year, leaving size, industry, and distributional shape unchanged. The predictability survives, suggesting that each firm's specific ecological context drives the premium (Appendix A.2).

Restricting the sample to post-2002 eBird vintages, when citizen science coverage is densest, strengthens rather than weakens the results. The small-firm FF3 alpha rises to 81 basis points per month, and the Fama-MacBeth BPI coefficient remains significant across all specifications (Internet Appendix Table IA.8).

4.2 BPI and emissions intensity as complementary measures

Biodiversity exposure and toxic emissions map to related but distinct economic concepts. While BPI and total toxic releases share an underlying data source, emissions intensity measures the standardized quantity of pollution, whereas BPI captures the ecological footprint of those operations. The question is whether a firm's ecological footprint contains pricing information beyond what emissions alone convey.

Adding standardized Releases as a control in the Fama-MacBeth specification (Table 5, columns 5–9) leaves the BPI coefficient unchanged, while Releases itself enters with a small and often insignificant coefficient. Prior work documents a pollution premium tied to toxic releases (Hsu et al., 2023), but this premium does not reflect expectational errors (Atilgan et al., 2024b). Section 5 tests whether the premium for biodiversity pressure follows the same channel or is distinct. Internet Appendix Table A.2 extends this analysis with size and idiosyncratic-volatility interactions and finds that the BPI and Releases premiums do not differ along either dimension.

Portfolio sorts in Internet Appendix Table IA.1 confirm that the two measures capture different return variation. The BPI spread is positive across all releases terciles under FF3 controls. The spread is largest among low-releases firms under three-factor benchmarks, but this ranking shifts across augmented factor models.

4.3 The size gradient

The biodiversity premium decays monotonically with firm size. Internet Appendix Table IA.2 reports a granular sort into ten size deciles crossed with BPI quintiles. The return spread is a continuous gradient rather than a discrete small-firm effect, statistically significant in eight of the ten deciles and falling from nearly 2% per month in the smallest stocks to an economically negligible spread in the largest. Whether benchmarked against the three-factor model, the five-factor model, or augmented models including liquidity, the small-firm alpha remains economically and statistically significant. An industry-conditioned triple sort

produces nearly identical results (Internet Appendix Table IA.4). The gradient is robust to NYSE size breakpoints that guard against microcap concentration (Internet Appendix Table IA.3) and survives restriction to the post-2002 period when eBird coverage is densest (Internet Appendix Table IA.8).

The portfolio sorts cannot by themselves distinguish whether this spread represents rational compensation for systematic exposure (the biodiversity risk premium hypothesis) or the correction of biased cash-flow expectations. A risk premium for aggregate ecological shocks should load broadly across the size distribution, since large firms retain wide cross-sectional dispersion in bird abundance ($\sigma = 1.39$ versus 1.13 for small firms). Instead, the premium vanishes in large firms. The next sections tests the risk premium hypothesis against the cash-flow expectations channel.

5 The Cash-Flow Expectations Channel

5.1 Operating costs

To trace the return premium back to firm fundamentals, I estimate a quarterly panel specification linking biodiversity pressure to operating costs,

$$\frac{\text{COGS}_{j,q}}{\text{Sales}_{j,q}} = \mu_j + \delta_y + \rho \frac{\text{COGS}_{j,q-1}}{\text{Sales}_{j,q-1}} + \beta \, \text{BPI}_{j,y} + \boldsymbol{\gamma}' \mathbf{X}_{j,q-1} + \varepsilon_{j,q}, \tag{4}$$

where μ_j and δ_y are firm and year fixed effects. The lagged dependent variable isolates within-firm changes beyond mean reversion. The subscript y on $\text{BPI}_{j,y}$ denotes the TRI vintage year, not the calendar year of the quarter.[7] The control vector $\mathbf{X}$ includes log assets, log book-to-market, ROA, leverage, and tangibility. I cluster standard errors by firm.[8]

[7] Following the timing rule in Section 2.1, the vintage assigned to quarter q is the most recent TRI release publicly available before q, which lags the calendar year by one to two years.

[8] Combining firm fixed effects with a lagged dependent variable introduces Nickell (1981) bias of order $1/T$. The median firm in the panel contributes 81 quarterly observations, so the bias is approximately 1.2%, negligible relative to the reported standard errors. BPI is also predetermined by construction (TRI data enter with a two-year lag), limiting contamination of the BPI coefficient from the dynamic panel correlation.

Firms with larger ecological footprints have lower operating performance across multiple measures (Table 6). A one-standard-deviation increase in BPI lowers quarterly income-to-assets by 0.07 percentage points and year-over-year sales growth by 1.3 percentage points. Gross margin declines by 0.40 percentage points. Decomposing operating costs into components, a one-standard-deviation increase in BPI raises cost of goods sold by 0.40 percentage points of sales, with no corresponding effect on selling, general, or administrative expenses (Table 7).

Although Derrien et al. (2025) document consumer demand effects for high-profile ESG events, the low public salience of the externality I study points to a production-side cost penalty. Table 7 shows that COGS rises significantly while SGA remains flat in both levels and year-over-year changes, ruling out a demand-side response in which firms increase spending to retain customers. Evidence from regulatory data indicates that the cost penalty is independent of the regulatory pipeline. In a panel of 270,889 facility-year observations with EPA region, industry, and year fixed effects, ecological richness is unrelated to inspections, informal enforcement actions, penalty probability, and penalty amounts (Appendix Table A.3). Emissions intensity, by contrast, strongly predicts all four outcomes. Regulators appear to target emissions levels rather than the ecological sensitivity of the surrounding area. The ecological cost channel is therefore distinct from the regulatory channels documented by Chen (2025) and Akbari et al. (2026), where discrete policy events trigger production adjustments.

5.2 Earnings expectations

If the return premium reflects errors in expectations rather than compensation for a recognized risk, three patterns should follow. First, analysts should hold ex-ante pessimistic views of high-footprint firms and embed that pessimism in their published forecasts. Second, BPI should predict positive earnings surprises, because the market underestimates future earnings for firms with large ecological footprints. Third, returns should concentrate around the

dates when hard information corrects those expectations (La Porta et al., 1997; Engelberg et al., 2018). Pankratz et al. (2023) document a similar pattern for physical climate risk, finding that firm-level heat exposure reduces operating income and generates negative earnings surprises that analysts fail to anticipate. Investors observe elevated costs but interpret them too pessimistically, so firms consistently beat expectations.

Analyst price targets show the ex-ante pessimism directly. I measure the analyst-implied 12-month return as the natural log of the median analyst price target divided by the closing stock price, following Engelberg et al. (2020). Table 8 reports panel regressions of this implied return on standardized BPI, with industry and year fixed effects, controls for size, book-to-market, analyst coverage, and price-target dispersion, and standard errors double-clustered by firm and year-month. A one-standard-deviation increase in BPI corresponds to a 6.3 percentage point lower implied return, and the coefficient is unchanged when toxic releases enter as an additional control. Analysts observe the depressed margins documented in Section 5.1 but lack a framework that ties them to ecological geography, so they extrapolate further operational decline into their price targets.

Next, I test whether firms beat these pessimistic forecasts. I measure earnings surprises two ways: against the analyst consensus forecast (SUE2), which captures what the market's intermediaries expected, and against the seasonal random-walk benchmark (SUE1), which compares a firm's quarterly earnings to the same quarter one year earlier (Chan et al., 1996). BPI predicts positive surprises against both benchmarks (Table 9, Panel A). For analyst consensus surprises, the coefficient attenuates when industry fixed effects are added, indicating that analysts partially capture biodiversity-related costs through industry peer benchmarking. I/B/E/S coverage also skews the SUE2 sample toward larger firms, where the BPI premium is weakest, so this specification covers the part of the cross-section where BPI matters least. For seasonal surprises, the coefficient remains stable whether or not industry fixed effects are included. Releases enters negatively for analyst-based surprises but does not predict seasonal surprises. The divergence indicates that the market partially

incorporates an observable pollutant signal into analyst forecasts (Derrien et al., 2025).[9]

Panel B of Table 9 tests whether these surprises explain the announcement return. In the baseline, BPI predicts announcement returns of 0.12 percentage points per standard deviation. Adding the analyst consensus surprise absorbs most of this effect and renders the BPI coefficient statistically insignificant, though the SUE2 specification mechanically restricts the sample to analyst-covered firms and therefore excludes the zero-coverage small firms that carry the largest return spread. Adding the seasonal surprise alone, which remains measurable on the broader sample, absorbs roughly one-third of the baseline return. BPI becomes insignificant in the I/B/E/S universe with both surprise controls (column 8). Analyst expectations absorb more of the announcement return than the seasonal benchmark in the covered subsample, because the market reacts to human forecasts rather than last year's earnings. BPI predicts which firms will beat analyst expectations, and the market reprices when those expectations are corrected at the earnings date.

Earnings announcements account for a substantial but incomplete share of the premium. In small firms, four quarterly announcement windows capture 40% of the annualized five-factor alpha, falling to 20% in mid-caps and negligible in large firms (Table 10, Panel A). The remaining share accrues between announcements as information diffuses gradually among heterogeneous investors (Hong and Stein, 1999). Resolution extends well beyond the three-day window: the High−Low BPI spread grows from 0.42% at $[-1, +1]$ to 2.18% at $[-1, +60]$ trading days, and this drift concentrates entirely in small firms (Table 10, Panel B). Sequential earnings reports therefore slowly force a repricing over multiple quarters. Atilgan et al. (2024a) document a parallel pattern for carbon emissions, with 30–50% of the emission premium concentrating at earnings announcements.

The earnings date sets the timing of correction, but it does not explain why expectations were too low to begin with. The two surprise benchmarks each capture a different source of the forecast error. The analyst consensus surprise (SUE2) directly tests whether hu-

[9]The effect on SUE1 is independent of abatement status (Internet Appendix Table IA.6).

man forecasters are too pessimistic about high-BPI firms. The ecological costs documented in Section 5.1 are location-bound rather than escalating. Year-over-year changes in both COGS/Sales and SGA/Sales are indistinguishable from zero (Table 7). Without observing the ecological source, investors see depressed margins and expect further deterioration (Schwartzstein, 2014). When the firm simply maintains its established cost structure, it beats the pessimistic forecast. This explanation applies to SUE2, which measures human expectations, and the attenuation of the BPI coefficient when industry fixed effects are added confirms that analysts partially learn ecological costs through industry peer benchmarking.

The seasonal random-walk surprise (SUE1) represents a different pattern. SUE1 is a mechanical benchmark that compares current earnings to the same quarter one year earlier, with no human expectations involved. High-BPI firms beat this benchmark within industries because their depressed but stable margins create a low base from which even modest revenue growth produces positive year-over-year earnings changes.

Greater analyst coverage does not resolve these expectational errors. Adding analyst coverage to the Panel A consensus surprise specification has no effect on the BPI coefficient, and the interaction of analyst coverage with BPI is similarly null for announcement returns (Internet Appendix Table IA.12). Ecological costs affect production but fall outside the analytical frameworks that intermediaries use to forecast earnings (Hong et al., 2007). Adding more analysts adds more forecasters applying the same framework. The premium persists because high-BPI firms are disproportionately outside the analyst-covered universe. 28% of high-BPI firm-quarters have zero analyst coverage, compared to 11% for low-BPI firms, so the seasonal forecast error goes uncorrected for the firms that generate the largest return spread.

BPI predicts positive earnings surprises through two complementary channels. Against analyst consensus forecasts (SUE2), the positive coefficient directly confirms that human forecasters are systematically too pessimistic about high-BPI firms. Against the seasonal random walk (SUE1), the positive coefficient reflects the positive year-over-year earnings

changes that depressed but stable margins produce off a low base. Analysts within an industry narrow their consensus forecast errors through peer comparison, but they attribute the cost structure to operational factors rather than ecology. Most high-BPI firms fall outside the analyst-covered universe, so this partial learning reaches only a fraction of the firms that generate the return spread.

5.3 Abatement and cost persistence

Unlike the production cuts that follow protected-area designations (Akbari et al., 2026), the ecological cost penalty requires no discrete regulatory trigger. Table 11 shows that elevated COGS persists at least four quarters beyond the BPI measurement date, consistent with a location-bound rather than transitory shock. Firms bear elevated production costs simply by operating in ecologically sensitive areas, and the absence of a triggering event sustains the information gap.

Although the regulator does not differentially target ecologically rich areas, firms may still mitigate the BPI cost penalty through voluntary action. I therefore examine whether voluntary investments in pollution prevention can offset the penalty. Following Chen (2025), I use data from the EPA's Toxic Release Inventory to identify firms that actively engage in source reduction activities through the voluntary Pollution Prevention (P2) program, which involves process changes, input substitutions, and waste-reduction investments at the facility level. I then interact this abatement status with BPI in the panel regressions of operating costs.

In the contemporaneous specification with firm and year fixed effects (Table 11, column 2), BPI raises COGS by 0.48 percentage points, but the BPI$\times$Abatement interaction is -0.21, offsetting roughly 43% of the penalty. At average ecological exposure, abatement has no independent effect on costs, indicating that remediation specifically offsets the ecological component of the cost burden rather than shifting operating expenses broadly. The market rewards firms that reduce their ecological cost burden, with the return premium concentrated

in non-abaters (52 bps/month) and attenuating for firms that actively remediate (Internet Appendix, Table IA.7). Because abaters mitigate the persistent cost drag, they avoid the shocks that generate the premium in non-remediating firms. Abating firms are also more likely to participate in voluntary CDP disclosure, a pattern examined in Section 6, suggesting that firms managing their ecological exposure also tend to make it visible. Voluntary disclosure may therefore reduce the forecast error by giving market participants a concrete ecological signal to incorporate into their models.

Forward-looking specifications in Table 11 show that the BPI cost penalty persists at four quarters out, with the BPI×Abatement interaction directionally aligned with the contemporaneous offset though less precisely estimated at longer horizons.

In unreported tabulations, roughly two-thirds of firm-years involve no abatement investment. Among firms that have faced enforcement, abatement rates are higher (53% of enforced firm-years versus 35% of non-enforced), though this association is largely a cross-sectional level difference rather than a within-firm timing response. An interaction of BPI with a post-penalty indicator shows that following enforcement, the cost gap between high- and low-BPI firms narrows, while the Post_Penalty main effect is indistinguishable from zero (Internet Appendix Table IA.14). The adjustment is firm-specific, as penalized high-BPI firms begin moving toward the cost structure of low-BPI peers, consistent with compliance-driven operational adjustments of the kind the abatement results document.

6 The Information Environment

The information environment surrounding biodiversity is sparse, and the intermediaries that produce pricing information are silent on the ecological footprint. Corporate disclosures rarely mention it, ESG ratings assign lower biodiversity risk scores to firms with larger footprints, and additional analyst coverage produces no incremental price discovery. The market only reprices after a regulatory penalty, and even that adjustment unfolds over 60

trading days.

6.1 The information gap

Both mandatory and voluntary disclosure channels are sparse, and voluntary disclosure is negatively related to the underlying footprint. Fewer than 2% of 10-K filings contain biodiversity-related language (Giglio et al., 2026), and fewer than one in five firms participates in the Carbon Disclosure Project's (CDP) voluntary biodiversity questionnaire. For mandatory 10-K disclosure, BPI does not predict whether a firm uses biodiversity language in any specification (Internet Appendix IA.13). For voluntary CDP participation, the relationship is inverted: Table 12 shows that firms facing greater ecological pressure are less likely to report their actions across five CDP biodiversity indicators.

Self-reported penalties also strategically understate true fines. When firms voluntarily disclose regulatory penalties, the amounts average $286,000, whereas violations later discovered by the EPA carry fines averaging $3.3 million, over an order of magnitude larger (Internet Appendix IA.13). Voluntary disclosure programs reward early reporting with reduced penalties, biasing the self-reported figures downward.

6.2 Information intermediary failure

Beyond disclosure, information intermediaries do not compensate for the sparse reporting. Double sorts on BPI interacted with institutional ownership, analyst coverage, board governance, and business complexity show no independent moderating effect (Internet Appendix Tables IA.10, IA.11, IA.13, and IA.16). The pooled BPI premium in these subsamples is often statistically zero because the data requirements (13-F filings, IBES estimates) filter out the small firms that drive the anomaly, but within each subsample the moderating variables themselves add no incremental explanatory power. Liquidity, idiosyncratic volatility, and short interest ratio interactions are similarly null (Internet Appendix Table IA.18). Analyst scrutiny and board oversight target financial complexity, leaving ecological externalities

outside their measurement scope.

Greater analyst coverage attenuates anomalies driven by limited attention to publicly available data, but it does not attenuate the BPI premium. Hirshleifer et al. (2013) show that innovative efficiency—a measure constructed from publicly available patent data—predicts returns among small firms with a similar size gradient. However, analyst coverage partially attenuates that premium because the underlying data exists and can be processed with sufficient effort. By contrast, additional analyst coverage of the BPI premium merely adds more forecasters applying the same framework that excludes ecology. Broader media attention similarly fails to prompt price discovery. RepRisk biodiversity news coverage rises with aggregate public attention for high-BPI firms, but the BPI premium itself does not vary with this attention (Internet Appendix Table IA.9).

Specialized ESG raters reproduce the disclosure failure rather than correcting it. I match annual MSCI scores and CDP participation to same-calendar-year BPI.[10] Table 13 shows that BPI is unrelated to a firm's overall MSCI environmental score, indicating that the rater's general assessment is uninformative about a firm's specific ecological footprint. Because MSCI coverage skews toward large firms where the BPI premium is already absent, this test asks whether the rater incorporates ecological footprint at all, rather than whether it moderates the premium.

The dedicated biodiversity sub-scores show a sharper failure. MSCI's biodiversity score is primarily a function of whether a firm voluntarily reports to the CDP. The coefficient on CDP disclosure is 1.07 and statistically significant, while the coefficient on BPI is not. Because firms with the largest ecological footprints are the least likely to report to CDP (Table 12), MSCI's biodiversity ratings systematically miss the most exposed firms. The marginal effect of BPI on the probability of receiving any MSCI biodiversity score is -0.011, strengthening to -0.027 with year and industry fixed effects (Table 14). Both estimates are significant at the 1% level. CDP voluntary disclosure increases coverage in the baseline but

[10] Both are measured with a lag, so the comparison is contemporaneous within the public information set.

loses significance once fixed effects absorb time trends. High-BPI firms are therefore absent from the sample that Table 13 covers.

In sum, MSCI's reliance on voluntary disclosure inverts the relationship between ecological footprint and ratings. MSCI rates high-BPI firms as *less* exposed to biodiversity risk (coefficient -0.23) and assigns *less* weight to the biodiversity issue where it is most material (coefficient -0.88). Firms that face the most ecological risk but disclose the least are rated as the least exposed.[11]

6.3 Enforcement and price discovery

Enforcement actions are an additional channel that brings ecological costs into prices. Regulatory penalties provide a discrete, firm-specific signal that forces repricing even for firms outside the analyst-covered sample, complementing the continuous earnings-time repricing documented in Section 5.

The market absorbs this information slowly, and the reaction is concentrated in enforcement events that carry monetary penalties. I compute Fama-French three-factor CARs around EPA enforcement conclusions, splitting events by whether a monetary penalty was recorded. Table 15 shows that the immediate three-day CAR is statistically zero, but a significant negative drift emerges over the following weeks. Over the $[-20, +60]$ window around a penalty enforcement action, the spread between high- and low-BPI firms reaches -7.78%. The spread represents an equity value loss of approximately \$25 million for the median penalized firm. Enforcement of ecological standards is spatial, discrete, and regime-dependent (He et al., 2020), making it difficult for investors to price the risk of regulatory action until a penalty materializes. The drift reflects the same information gap that prevented ex ante pricing.

Panel B of Table 15 shows that only enforcement actions carrying monetary penalties generate the drift. Non-penalty actions produce a spread of only -1.14%, less than one-

[11] In unreported distributed-lag regressions, MSCI scores show no significant response to BPI at horizons of one to four years.

sixth the magnitude of penalty events. Within the enforcement channel, investors respond to concrete cost realizations rather than to regulatory attention alone.

The enforcement drift and the continuous earnings-surprise mechanism are complementary rather than contradictory. The positive earnings surprises in Section 5 arise because investors misread existing ecological costs. The enforcement drift instead reflects new information about future compliance costs, remediation obligations, and heightened regulatory scrutiny that a penalty reveals for the first time. Enforcement is rare, with roughly 400 penalty events affecting fewer than 200 firms over the sample period, so this correction does not offset the broad-based premium that accrues across thousands of firm-quarters.

The premium is difficult to reconcile with rational ambiguity pricing (Epstein and Schneider, 2008). Under recognized ambiguity, sophisticated investors identify the uncertain attribute and demand compensation for bearing it. Three features of the data argue against this channel. First, enforcement triggers a prolonged 60-day drift rather than an immediate price jump, whereas ambiguity-averse investors who recognize the risk should reprice immediately upon revelation. Second, analyst coverage does not attenuate the premium, whereas if investors recognized biodiversity as an ambiguous attribute, greater coverage should reduce the uncertainty and lower the required compensation. Institutional ownership similarly shows no moderating effect. Third, positive earnings surprises recur quarter after quarter rather than dissipating as investors update their priors, suggesting the ecological attribute is not recognized as a source of uncertainty in the first place. These patterns are more consistent with partial observability. Investors observe aggregate costs but cannot identify the ecological source, so the errors in expectations that generate the premium persist regardless of intermediary sophistication.

7 Conclusion

I document an equity premium driven by unmeasured ecological footprints. Financial markets have developed effective tools for pricing environmental risks that reduce to a single number: carbon emissions aggregate into CO_2 equivalents, regulators mandate their disclosure, and intermediaries incorporate the resulting data into forecasts and ratings. Biodiversity pressure resists this treatment because it depends on spatial heterogeneity and local ecosystem sensitivity—characteristics that vary across facilities and cannot be easily summarized in a standardized metric. Analysts, ESG raters, and corporate disclosures are each silent on localized ecological degradation, so investors observe the higher production costs that ecological exposure generates but cannot identify their source. Investors misattribute stable, location-bound expenses to deteriorating operational efficiency, are pessimistic in their price targets, and are surprised when margins hold steady.

The evidence points to measurement as the central obstacle. Emerging frameworks such as the Taskforce on Nature-related Financial Disclosures require location-specific ecological assessments, but adoption remains limited to firms that voluntarily opt in, which are the firms that already manage their ecological exposure. For the small and mid-cap firms where the premium is largest, the cost of specialized ecological measurement continues to exceed the revenue it would generate for any individual intermediary. Until measurement technologies lower this cost or mandatory disclosure extends beyond voluntary adopters, localized biodiversity risks will continue to generate persistent forecast errors in equity markets.

The analysis has several limitations. I rely on bird abundance as a proxy for overall ecosystem health. Birds serve as a reliable indicator species given dense citizen science tracking, but they represent only one component of a complex biological network. I restrict the sample to industrial facilities with toxic releases, leaving the mechanism untested in service-sector firms with less physical infrastructure. Finally, the evidence linking ecological exposure to operating costs depends on panel regressions with firm fixed effects rather than a strictly exogenous shock to the local environment.

References

Akbari, Amir, Lilian Ng, Man Duy Pham, and Jing Yu, "The Real Effects of Protecting Biodiversity," *Review of Finance*, 2026, *30* (1), 193–230.

Akey, Pat and Ian Appel, "The Limits of Limited Liability: Evidence from Industrial Pollution," *Journal of Finance*, 2021, *76* (1), 5–55.

Altman, Edward I., "Financial Ratios, Discriminant Analysis and the Prediction of Corporate Bankruptcy," *Journal of Finance*, 1968, *23* (4), 589–609.

Amihud, Yakov, "Illiquidity and Stock Returns: Cross-Section and Time-Series Effects," *Journal of Financial Markets*, 2002, *5* (1), 31–56.

Atilgan, Yigit, K. Ozgur Demirtas, Alex Edmans, and A. Doruk Gunaydin, "Does the Carbon Premium Reflect Risk or Outperformance?," Finance Working Paper 940/2023, ECGI 2024. Available at SSRN: https://ssrn.com/abstract=4573622.

_ , _ , and A. Doruk Gunaydin, "Pollution Premium: Further Evidence," Technical Report, Sabanci University 2024. Available at SSRN: https://ssrn.com/abstract=4974375.

Avramov, Doron, Si Cheng, Abraham Lioui, and Andrea Tarelli, "Sustainable Investing with ESG Rating Uncertainty," *Journal of Financial Economics*, 2022, *145* (2), 642–664.

Berg, Florian, Julian F Kölbel, and Roberto Rigobon, "Aggregate Confusion: The Divergence of ESG Ratings," *Review of Finance*, November 2022, *26* (6), 1315–1344.

Bolton, Patrick and Marcin Kacperczyk, "Do investors care about carbon risk?," *Journal of Financial Economics*, November 2021, *142* (2), 517–549.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Carhart, Mark M., "On Persistence in Mutual Fund Performance," *Journal of Finance*, 1997, *52* (1), 57–82.

Chan, Louis K. C., Narasimhan Jegadeesh, and Josef Lakonishok, "Momentum Strategies," *The Journal of Finance*, 1996, *51* (5), 1681–1713.

Chen, Jiaqi, "Redeploying Dirty Assets: The Impact of Environmental Regulation on Asset Allocation," *Journal of Financial Economics*, 2025, *163*, 103959.

Chen, Wei, Paul Hribar, and Samuel Melessa, "Incorrect Inferences When Using Residuals as Dependent Variables," *Journal of Accounting Research*, 2018, *56* (3), 751–796.

Creti, Anna, Yujun Huang, Binghan Jiang, and Maria-Eugenia Sanin, "Biodiversity Risk, Firm Performance, and Market Mispricing," Technical Report, University Paris Dauphine – PSL 2024. Available at SSRN: https://ssrn.com/abstract=4765039.

Derrien, François, Philipp Krueger, Augustin Landier, and Tianhao Yao, "ESG News, Future Cash Flows, and Firm Value," *Journal of Finance*, 2025, *80* (6), 3499–3554.

Dimson, Elroy, Marcin Kacperczyk, and Laura Starks, "Biodiversity and Natural Resource Finance," *Review of Finance*, 2026.

Edmans, Alex, "Does the Stock Market Fully Value Intangibles? Employee Satisfaction and Equity Prices," *Journal of Financial Economics*, 2011, *101* (3), 621–640.

Engelberg, Joseph, R. David McLean, and Jeffrey Pontiff, "Anomalies and News," *The Journal of Finance*, 2018, *73* (5), 1971–2001.

_ , _ , and _ , "Analysts and Anomalies," *Journal of Accounting and Economics*, 2020, *69* (1), 101249.

Epstein, Larry G. and Martin Schneider, "Ambiguity, Information Quality, and Asset Pricing," *The Journal of Finance*, 2008, *63* (1), 197–228.

Fama, Eugene F. and James D. MacBeth, "Risk, Return, and Equilibrium: Empirical Tests," *Journal of Political Economy*, 1973, *81* (3), 607–636.

_ **and Kenneth R. French**, "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics*, 1993, *33* (1), 3–56.

_ **and** _ , "A Five-Factor Asset Pricing Model," *Journal of Financial Economics*, 2015, *116* (1), 1–22.

Fink, Daniel, Tom Auer, Alison Johnston, Matthew Strimas-Mackey, Shawn Ligocki, Orin Robinson, Wesley Hochachka, Steve Jaromír Kelling, Amanda Rodewald et al., "eBird Status and Trends, Data Version: 2021," *Cornell Lab of Ornithology*, 2021.

Frazzini, Andrea and Lasse Heje Pedersen, "Betting Against Beta," *Journal of Financial Economics*, 2014, *111* (1), 1–25.

Garel, Alexandre, Arthur Romec, Zacharias Sautner, and Alexander F. Wagner, "Do Investors Care About Biodiversity?," *Review of Financial Studies*, 2024, *37* (5), 1468–1504.

_ **,** _ **,** _ **, and** _ , "Firm-Level Nature Dependence," *Review of Finance*, 2026, *30* (1), 231–272.

Giglio, Stefano, Theresa Kuchler, Johannes Stroebel, and Olivier Wang, "The Economics of Biodiversity Loss," Working Paper 32678, National Bureau of Economic Research July 2024.

_ **,** _ **,** _ **, and Xuran Zeng**, "Biodiversity Risk," *Review of Finance*, 2026, *30* (1), 131–161.

Gjerde, Snorre, Zacharias Sautner, Alexander F. Wagner, and Alexis Wegerich, "Corporate Nature Risk Perceptions," *Review of Finance*, 2026, *30* (1), 11–56.

He, Guojun, Shaoda Wang, and Bing Zhang, "Watering Down Environmental Regulation in China," *The Quarterly Journal of Economics*, 2020, *135* (4), 2135–2185.

Hirshleifer, David, Po-Hsuan Hsu, and Dongmei Li, "Innovative Efficiency and Stock Returns," *Journal of Financial Economics*, 2013, *107* (3), 632–654.

Hogan, Dylan, Eyal G. Frank, Josephine Gantois, and Anouch Missirian, "The Effectiveness of Local Conservation Ballots," *American Economic Review: Papers & Proceedings*, 2025, *115*, 403–408.

Hong, Harrison and Jeremy C. Stein, "A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets," *Journal of Finance*, 1999, *54* (6), 2143–2184.

_ , _ , and Jialin Yu, "Simple Forecasts and Paradigm Shifts," *Journal of Finance*, 2007, *62* (3), 1207–1242.

Hsu, Po-Hsuan, Kai Li, and Chi-Yang Tsou, "The Pollution Premium," *The Journal of Finance*, 2023, *78* (3), 1343–1392. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.13217.

IPBES, "Global Assessment Report on Biodiversity and Ecosystem Services," Technical Report, Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services, Bonn, Germany 2019. E. S. Brondizio, J. Settele, S. Díaz, and H. T. Ngo (editors).

Nagel, Stefan, "Short Sales, Institutional Investors and the Cross-Section of Stock Returns," *Journal of Financial Economics*, 2005, *78* (2), 277–309.

Pankratz, Nora M.C., Rob Bauer, and Jeroen Derwall, "Climate Change, Firm Performance, and Investor Surprises," *Management Science*, 2023, *69* (12), 7352–7398.

Porta, Rafael La, Josef Lakonishok, Andrei Shleifer, and Robert Vishny, "Good News for Value Stocks: Further Evidence on Market Efficiency," *Journal of Finance*, 1997, *52* (2), 859–874.

Sager, Lutz and Gregor Singer, "Clean Identification? The Effects of the Clean Air Act on Air Pollution, Exposure Disparities, and House Prices," *American Economic Journal: Economic Policy*, 2025, *17* (1), 292–330.

Schwartzstein, Joshua, "Selective Attention and Learning," *Journal of the European Economic Association*, 2014, *12* (6), 1423–1452.

Taskforce on Nature-Related Financial Disclosures, "Recommendations of the Taskforce on Nature-Related Financial Disclosures," https://tnfd.global/publication/recommendations-of-the-taskforce-on-nature-related-financial-disclosures/ 2025.

U.S. Environmental Protection Agency, "2022 TRI National Analysis," Technical Report, EPA 2023.

Ľuboš Pástor and Robert F. Stambaugh, "Liquidity Risk and Expected Stock Returns," *Journal of Political Economy*, 2003, *111* (3), 642–685.

_ , _ , and Lucian A. Taylor, "Sustainable Investing in Equilibrium," *Journal of Financial Economics*, 2021, *142* (2), 550–571.

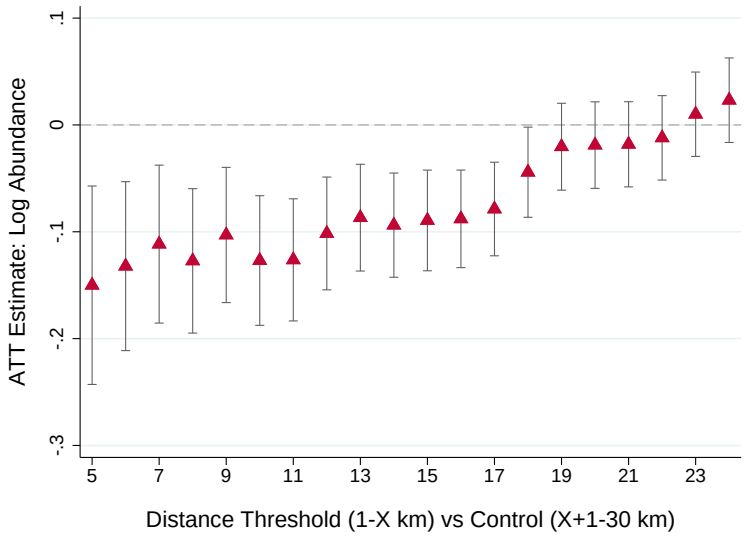
Zhang, Shaojun, "Carbon Returns across the Globe," *Journal of Finance*, 2025, *80* (1), 615–645.

Figures

Figure 1: Ecological impact across distance thresholds

Callaway–Sant'Anna doubly robust ATT estimates of the effect of TRI facility construction on nearby bird populations, varying the treatment-control boundary from 5 to 24 km. For each threshold X, the treated ring spans $(1, X]$ km and the control ring spans $(X+1, 30]$ km. Bars show 95% confidence intervals clustered at the facility level. Controls include seasonal coverage, NLCD impervious surface growth, total observer-hours, and total observers. Panel (a): log bird abundance. Panel (b): log species richness. Both outcomes show significant negative effects that attenuate monotonically as the treatment ring expands to encompass less affected areas, becoming statistically insignificant around 19–20 km.

(a) Abundance



(b) Richness

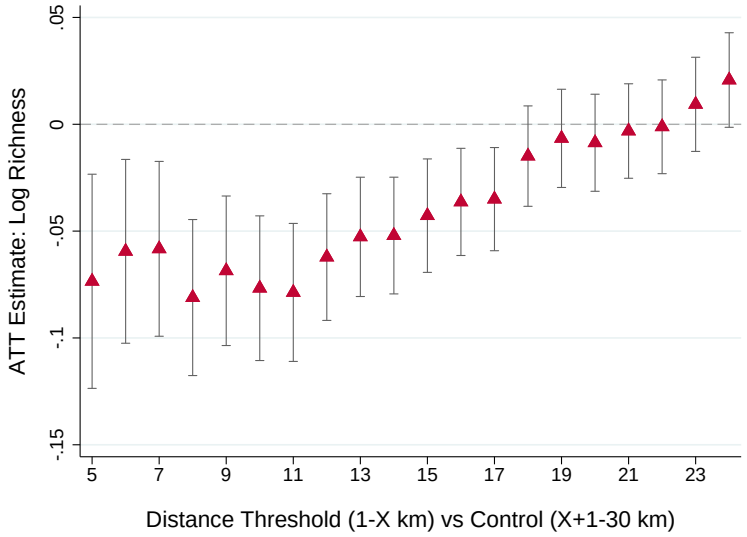
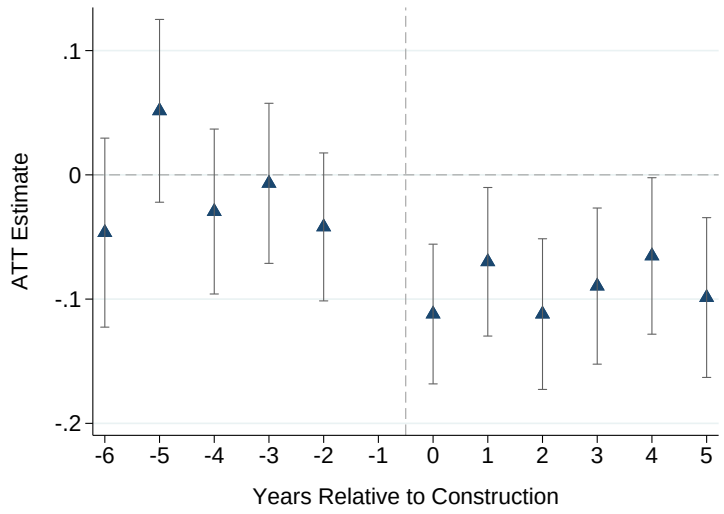


Figure 2: Ecological impact of industrial facility construction

Event study estimates of the effect of TRI facility construction on nearby bird populations. Each panel plots Callaway–Sant'Anna doubly robust ATT estimates by event time, with 95% confidence intervals clustered at the facility level. Treatment sites lie 1–14 km from the facility; control sites lie 15–30 km away. Controls include seasonal coverage, NLCD impervious surface growth, total observer-hours, and total observers. The dashed vertical line marks $t = -1$. Panel (a): log bird abundance declines 6% to 9% after construction. Panel (b): log species richness declines 3% to 5%.

(a) Abundance



(b) Richness

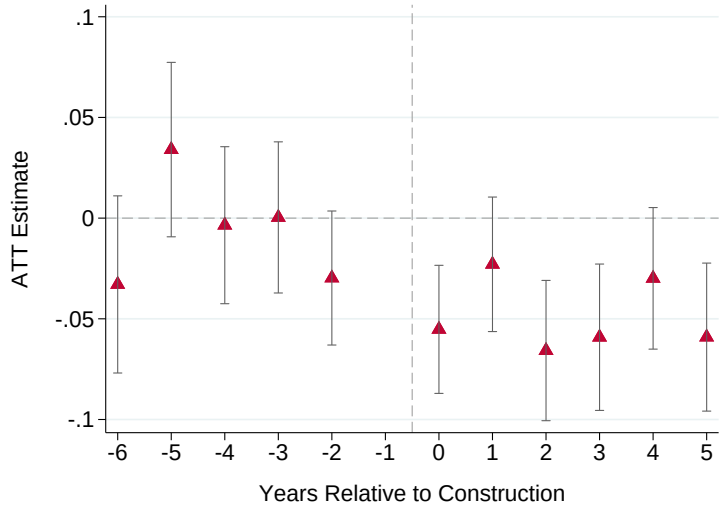
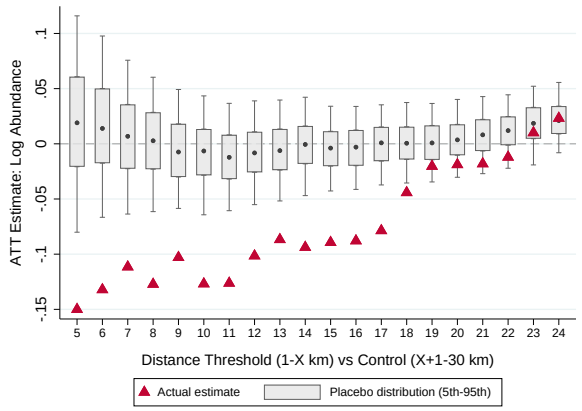


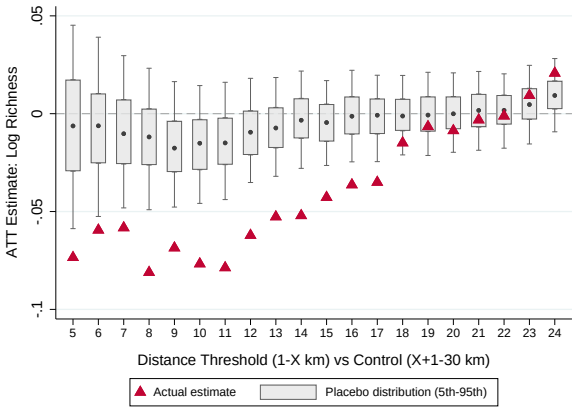
Figure 3: Placebo distributions with actual estimates

Placebo distributions from 1,000 permuted-timing iterations of the staggered difference-in-differences estimator, varying the treatment radius from 5 to 24 km. At each threshold, the treated ring spans $(1, X]$ km and the control ring spans $(X+1, 30]$ km. Gray boxes show the interquartile range (25th–75th percentile) of placebo estimates; whiskers extend to the 5th and 95th percentiles; dots mark the placebo median. Red triangles show the actual estimates from Section 3. Panels (a) and (b) plot ATT point estimates; panels (c) and (d) plot z-statistics, with dashed horizontal lines at ± 1.96.

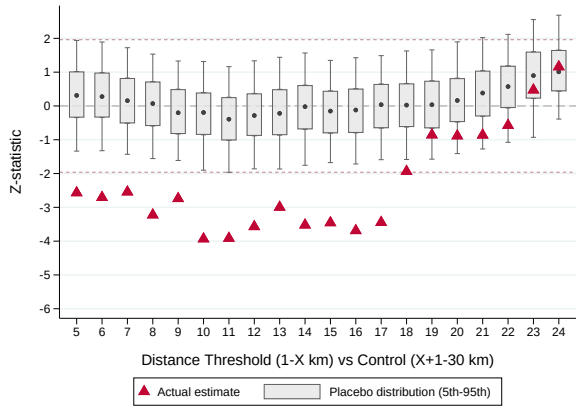
(a) Abundance ATT estimates



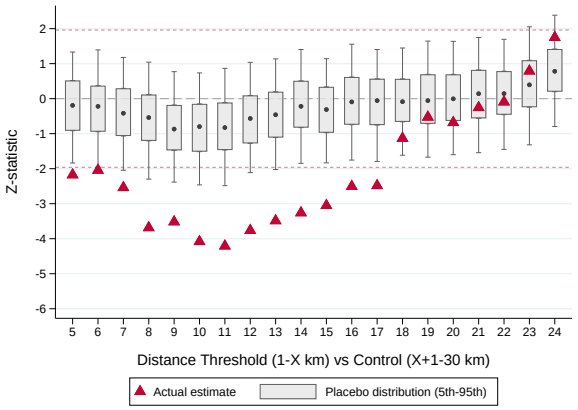
(b) Richness ATT estimates



(c) Abundance z-statistics



(d) Richness z-statistics



Tables

Table 1: Variable definitions

Continuous variables in Fama-MacBeth regressions are standardized. Portfolio sorts use unstandardized values.

Variable	Definition
Panel A: Key independent variables	
BPI	ln(Bird Abundance/Total Assets). Bird Abundance is the sum across a firm's TRI facilities of facility-level exponentiated means of effort-adjusted log bird counts from eBird checklists in the 15–30 km ring. Higher values = larger ecological footprint per dollar of assets. Source: eBird, TRI, Compustat.
Releases	ln(Total TRI Releases/Total Assets). Total on-site air, water, and ground releases in pounds. Source: TRI, Compustat.
Abatement indicator	Equal to one if any facility reports source reduction to the TRI P2 program. Source: TRI.
ln(bird abundance)	Natural log of effort-adjusted bird abundance in the 15–30 km ring, before asset scaling. The ecological numerator of BPI. Source: eBird, TRI.
Panel B: Return and operating performance variables	
Monthly return	Monthly holding-period return, inclusive of dividends. Source: CRSP.
COGS/Sales	Cost of goods sold divided by net sales, quarterly. Source: Compustat.
Gross margin	$(1 - \text{COGS/Sales}) \times 100$.
SGA/Sales	Selling, general, and administrative expenses divided by net sales, quarterly. Source: Compustat.
Sales growth	Year-over-year change in quarterly sales: $(Sale_q - Sale_{q-4})/\lvert Sale_{q-4}\rvert$. Source: Compustat.
SUE (seasonal random walk)	$(EPS_q - EPS_{q-4})/\sigma(EPS_q - EPS_{q-4})$, estimated over the prior eight quarters. Source: Compustat.
SUE (analyst-based)	Actual EPS minus mean consensus forecast, scaled by price. Source: I/B/E/S.
Analyst-implied return	ln(Median PT/|prc|), where Median PT is the median analyst 12-month price target and prc is the closing stock price, following Engelberg et al. (2020). Source: I/B/E/S Summary, CRSP.
Panel C: Firm characteristics and controls	
Log market cap	ln(price × shares outstanding), in thousands. Source: CRSP.
Log book-to-market	ln(book equity/market equity). Source: Compustat, CRSP.
Momentum	Cumulative return, months $t-12$ to $t-2$. Source: CRSP.
ROA	Income before extraordinary items / lagged total assets, annualized. Source: Compustat.
Leverage	Total debt / total assets. Source: Compustat.
Investment rate (I/K)	Capital expenditures / lagged net PP&E. Source: Compustat.
Log tangibility	ln(net PP&E/total assets). Source: Compustat.
Idiosyncratic volatility	SD of daily FF3 residuals, prior 60 trading days. Source: CRSP.
Gross profitability (GP/A)	(Sales − COGS)/Total Assets, quarterly. Source: Compustat.
Asset turnover	Sales/Total Assets, quarterly. Source: Compustat.
Altman Z-score	Modified four-component: $1.2(\text{WC/TA}) + 3.3(\text{EBIT/TA}) + 0.6(\text{MVE/TL}) + 1.0(\text{Sales/TA})$, following Altman (1968). Source: Compustat, CRSP.
Panel D: Information environment variables	
Institutional ownership	Fraction of shares held by 13F institutions, quarterly. Source: Thomson Reuters.
Analyst coverage	Number of distinct estimates in I/B/E/S consensus. Source: I/B/E/S.
PT dispersion	Standard deviation of analyst 12-month price targets, scaled by their median, monthly. Source: I/B/E/S Summary.
Amihud illiquidity	Monthly average of daily $\lvert r_d\rvert/\text{DolVol}_d$ ($\times 10^6$), following Amihud (2002). Source: CRSP.
Turnover	Monthly volume / shares outstanding (%). Source: CRSP.
MSCI environmental score	Industry-adjusted environmental pillar score. Source: MSCI ESG Ratings.
10-K biodiversity disclosure	Biodiversity keywords in 10-K filings, following Giglio et al. (2026). Source: SEC EDGAR.
CDP disclosure	Voluntary reporting to Carbon Disclosure Project. Source: CDP.
Panel E: Additional variables	
Number of TRI facilities	Active TRI facilities operated by the firm per year. Source: TRI.
Log total assets	ln(ATQ), in millions. Source: Compustat.
Short interest ratio (SIR)	Shares short / shares outstanding, monthly. Source: Compustat.
Board size	Number of directors. Source: BoardEx.
Board independence	Fraction of non-executive directors. Source: BoardEx.
Director network size	Average director NetworkSize (raw levels in Table 2; log-transformed in regressions). Source: BoardEx.
Long-term growth forecast	Median analyst long-term EPS growth forecast (%). Source: I/B/E/S.

Table 2: Summary statistics

The sample comprises 1,873 TRI-linked firms from October 1988 through December 2024. BPI is ln(bird abundance/total assets); Releases is ln(TRI releases/total assets). Panel A return-predicting characteristics are lagged one month relative to the return they predict; BPI and Releases follow the vintage-year timing rule in Section 2.1. All continuous variables are winsorized at the 1st and 99th percentiles. Panel D reports BoardEx variables (available from 1999). Observation counts vary because the table is not restricted to a common sample.

	Mean	SD	p5	p25	Median	p75	p95	N
Panel A: Monthly variables								
Monthly return (%)	1.202	12.981	-17.003	-5.054	0.808	6.743	20.000	235,586
BPI	-3.529	1.808	-6.708	-4.722	-3.447	-2.275	-0.638	232,581
Releases	3.076	3.237	-3.245	1.371	3.609	5.374	7.388	216,653
ln(bird abundance)	3.882	1.357	1.797	2.881	3.842	4.845	6.172	235,586
Log market cap	14.099	2.120	10.500	12.655	14.140	15.538	17.592	232,581
Log book-to-market	-0.701	0.753	-2.016	-1.125	-0.650	-0.215	0.427	225,861
Momentum $(-12, -2)$	0.132	0.376	-0.460	-0.063	0.127	0.314	0.741	215,717
ROA	10.820	24.565	-23.262	4.012	12.284	21.543	41.181	232,548
Leverage	0.280	0.171	0.004	0.159	0.272	0.384	0.583	232,581
Investment rate (I/K)	0.145	0.200	0.014	0.041	0.083	0.163	0.497	229,982
Idiosyncratic volatility	0.022	0.016	0.008	0.012	0.018	0.027	0.053	232,333
Number of TRI facilities	5.8	9.0	1.0	1.0	3.0	6.0	21.0	235,586
Log total assets	7.421	1.920	4.300	6.075	7.405	8.740	10.625	232,581
Log tangibility (PPE/A)	-1.351	0.671	-2.533	-1.759	-1.281	-0.870	-0.365	232,465
Amihud illiquidity ($\times 10^6$)	0.256	1.161	0.000	0.000	0.002	0.021	1.015	231,810
Turnover	6.625	6.250	0.803	2.503	4.841	8.429	19.147	231,810
Institutional ownership	0.735	0.263	0.134	0.617	0.813	0.941	1.000	52,437
Analyst coverage	7.3	5.2	2.0	3.0	6.0	10.0	18.0	100,566
Short interest ratio (SIR)	0.041	0.062	0.001	0.010	0.022	0.047	0.145	185,257
Panel B: Quarterly operating variables								
COGS/Sales (%)	69.2	21.4	37.0	60.2	70.9	79.8	91.0	78,318
Gross margin (%)	30.8	21.4	9.0	20.2	29.1	39.8	63.0	78,318
SGA/Sales (%)	17.2	12.4	0.0	8.2	14.9	23.8	41.8	78,281
Sales growth (%, YoY)	8.9	27.2	-26.1	-3.7	5.4	15.9	55.0	71,049
Gross profitability (GP/A)	0.079	0.045	0.018	0.047	0.072	0.102	0.167	78,316
Asset turnover	0.278	0.144	0.082	0.179	0.256	0.350	0.545	78,697
Altman Z-score	9.791	34.761	0.636	1.385	2.514	4.947	29.443	72,045
Abatement indicator	0.448	0.497	0.000	0.000	0.000	1.000	1.000	26,561
Panel C: Earnings announcement variables								
SUE (seasonal random walk)	-0.018	1.646	-2.576	-0.504	0.065	0.624	2.473	74,164
SUE (analyst-based)	0.977	3.411	-3.768	-0.403	0.668	2.178	6.708	59,902
Long-term growth forecast	12.437	10.680	-0.300	7.200	11.000	15.000	30.000	34,053
Analyst-implied return	0.085	0.512	-0.748	-0.030	0.100	0.248	0.730	142,828
Panel D: Board governance variables (BoardEx subsample)								
Board size	11.51	3.49	6.00	9.00	11.00	14.00	17.00	138,516
Board independence	0.710	0.150	0.470	0.580	0.670	0.880	0.910	138,516
Director network size	1737	1115	358	938	1530	2292	3784	138,516

Table 3: Ecological impact: progressive controls

This table reports staggered difference-in-differences estimates of the effect of TRI facility construction on local bird populations. The estimator is Callaway and Sant'Anna (2021), using doubly-robust inverse probability tilting with not-yet-treated units as controls. Treated units are facility-ring-years for the inner ring (1–14 km); control units are the corresponding outer ring (15–30 km). The outcome is log adjusted abundance (Panel A) or log adjusted species richness (Panel B). Each column progressively adds controls: (1) none; (2) seasonal coverage (month indicators); (3) plus NLCD urbanization growth; (4) plus observer effort (log observers, log hours). Post ATT is the average treatment effect across post-construction event times ($t = 0$ to $+5$); Pre ATT averages pre-construction periods ($t = -5$ to -1). Standard errors are clustered at the facility level. z-statistics in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1)	(2)	(3)	(4)
Panel A: Log Abundance				
Post ATT	-0.064^{***}	-0.074^{***}	-0.074^{***}	-0.092^{***}
	(-2.66)	(-2.94)	(-2.92)	(-3.68)
Pre ATT	-0.034	-0.027	-0.024	-0.010
	(-1.49)	(-1.13)	(-1.01)	(-0.42)
Panel B: Log Species Richness				
Post ATT	-0.032^{**}	-0.039^{***}	-0.040^{***}	-0.050^{***}
	(-2.38)	(-2.77)	(-2.84)	(-3.59)
Pre ATT	-0.022^{*}	-0.016	-0.016	-0.006
	(-1.69)	(-1.20)	(-1.20)	(-0.42)
Coverage (month indicators)	×	✓	✓	✓
Urbanization (NLCD growth)	×	×	✓	✓
Effort (observers, hours)	×	×	×	✓
Observations	58,489	58,489	57,911	57,906
Facilities (clusters)	3,583	3,583	3,583	3,583

Table 4: Size × BPI double sort: multi-factor alphas

This table reports alphas of high-minus-low BPI spread portfolios from a conditional double sort on market equity terciles and BPI terciles. Each vintage year, firms are first sorted into size terciles, then into BPI terciles within each size group. Value-weighted portfolios are formed within each of the nine size–BPI cells. Each cell reports the alpha of the high-BPI minus low-BPI spread portfolio, in percent per month. Column (1) reports CAPM alphas; column (2) controls for the three Fama and French (1993) factors; column (3) uses the five Fama and French (2015) factors; column (4) adds Betting Against Beta (Frazzini and Pedersen, 2014); column (5) adds the momentum factor (Carhart, 1997); and column (6) adds traded liquidity (Pástor and Stambaugh, 2003); and column (7) adds the Pástor et al. (2021) green factor to FF5 over the 2009–2022 subsample where green factor data are available. Pooled averages the high and low portfolios across size groups before computing the spread. Newey-West(12) standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1) CAPM	(2) FF3	(3) FF5	(4) FF5+BAB	(5) FF5+Mom	(6) FF5+Liq	(7) FF5+Green
H−L BPI α within size tercile (%/month):							
Small	+0.760***	+0.786***	+0.693***	+0.690***	+0.622***	+0.714***	+0.708***
	(0.112)	(0.105)	(0.115)	(0.120)	(0.121)	(0.118)	(0.225)
Mid	+0.559***	+0.555***	+0.539***	+0.557***	+0.462***	+0.546***	+0.320*
	(0.096)	(0.097)	(0.106)	(0.111)	(0.106)	(0.109)	(0.186)
Big	+0.103	+0.052	−0.123	−0.120	−0.112	−0.149	+0.173
	(0.160)	(0.122)	(0.123)	(0.125)	(0.118)	(0.124)	(0.149)
Pooled	+0.474***	+0.464***	+0.370***	+0.376***	+0.324***	+0.370***	+0.400***
	(0.073)	(0.064)	(0.071)	(0.073)	(0.071)	(0.073)	(0.126)
N months	434	434	434	432	434	434	159

Table 5: Fama-MacBeth regressions: BPI and stock returns

This table reports monthly Fama-MacBeth cross-sectional regressions of excess stock returns on the Biodiversity Pressure Index (BPI) and progressively expanded control sets. The dependent variable is the monthly excess return in percent. BPI is ln(bird abundance/total assets) and Releases is ln(TRI releases/total assets), both cross-sectionally standardized to zero mean and unit variance each month. All other controls enter as winsorized levels: log book-to-market equity, cumulative return from $t-12$ to $t-2$ (momentum), log market equity, return on assets, leverage, idiosyncratic volatility (standard deviation of daily FF3 residuals), and the investment rate (capital expenditures over lagged PP&E). Column (8) adds asset turnover (Sales/Assets), gross profitability ((Sales−COGS)/Assets), and tangibility (log PP&E/Assets) as a denominator defense: these ratios share BPI's 1/Assets scaling. Column (9) adds the modified Altman (1968) Z-score ($Z = 1.2 \times \text{WC/TA} + 3.3 \times \text{EBIT/TA} + 0.6 \times \text{MVE/TL} + 1.0 \times \text{Sales/TA}$) as a distress control. Specifications (1)–(7) use a common sample of 194,036 firm-months spanning 428 months; specifications (8) and (9) use slightly fewer observations due to additional data requirements. All specifications include Fama-French 12-industry fixed effects. The negative signs on log book-to-market and momentum reflect the restricted TRI reporting sample, which excludes financial firms and overweights capital-intensive industrials; these signs are stable across specifications and do not affect the BPI coefficient. Newey-West(6) t-statistics in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
BPI (standardized)	+0.130***	+0.250***	+0.252***	+0.180***	+0.154***	+0.185***	+0.167***	+0.168***	+0.174***
	(2.65)	(5.29)	(4.95)	(4.06)	(3.26)	(4.16)	(3.83)	(3.63)	(3.59)
Releases (standardized)					+0.055*	+0.039	+0.055*	+0.027	+0.029
					(1.72)	(1.33)	(1.84)	(0.94)	(1.04)
Log book-to-market		−0.864***	−0.862***	−0.666***	−0.679***	−0.680***	−0.675***	−0.691***	−0.827***
		(−7.49)	(−8.67)	(−6.59)	(−6.81)	(−7.03)	(−7.00)	(−6.97)	(−8.15)
Momentum ($t-12, t-2$)		−0.221	−0.247	−0.684***	−0.684***	−0.649***	−0.664***	−0.706***	−0.746***
		(−0.85)	(−0.95)	(−2.66)	(−2.67)	(−2.69)	(−2.76)	(−2.87)	(−3.01)
Log ME			+0.002	−0.083**	−0.089**	+0.008	+0.024	+0.029	−0.002
			(0.04)	(−2.05)	(−2.16)	(0.22)	(0.70)	(0.79)	(−0.06)
ROA				+0.033***	+0.033***	+0.035***	+0.034***	+0.033***	+0.034***
				(11.57)	(11.47)	(12.91)	(12.72)	(11.69)	(11.33)
Leverage				+0.018	−0.012	−0.224	−0.259	−0.224	−0.760**
				(0.05)	(−0.03)	(−0.65)	(−0.77)	(−0.66)	(−2.10)
Idiosyncratic volatility						+24.328***	+25.731***	+26.353***	+26.379***
						(3.84)	(4.08)	(4.12)	(4.09)
Investment rate (I/K)							−1.588***	−1.670***	−1.750***
							(−5.55)	(−5.45)	(−5.65)
Asset turnover (S/A)								+0.494	+0.372
								(1.52)	(1.09)
Gross profitability (GP/A)								−0.774	−0.301
								(−0.61)	(−0.24)
Tangibility (ln PPE/A)								+0.118*	+0.174**
								(1.76)	(2.48)
Altman Z-score									−0.000***
									(−2.56)
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Avg. firms/month	453	453	453	453	453	453	453	446	435
Avg. R^2	0.07	0.11	0.12	0.13	0.14	0.15	0.15	0.16	0.17

Table 6: Biodiversity exposure and operating performance

This table examines the relationship between biodiversity exposure and operating performance. The sample is a quarterly panel of TRI-reporting firms from 1989 to 2024. The dependent variables are Income/Assets (quarterly income before extraordinary items divided by lagged total assets), Sales Growth (year-over-year quarterly sales growth), Gross Margin ((Sales − COGS)/Sales), and Operating Margin ((Sales − XOPR)/Sales), each expressed in percentage points. The key independent variable is BPI, standardized within TRI vintage year, where higher values indicate greater ecological liability. All regressions include the one-quarter lagged dependent variable to isolate within-firm changes in performance beyond mean reversion. A negative BPI coefficient indicates that firms with larger biodiversity footprints have lower profitability and margins relative to their own recent history. Controls (log assets, log book-to-market, ROA, leverage, log PPE/Assets) are lagged one quarter. All specifications include firm and year fixed effects. t-statistics based on standard errors clustered by firm appear in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1) Inc/A	(2) Sales Gr.	(3) Gross Margin	(4) Oper. Margin
BPI (standardized)	-0.073^{***}	-1.324^{***}	-0.399^{***}	-0.474^{***}
	(-3.52)	(-7.06)	(-4.47)	(-5.00)
Lagged DV	0.271^{***}	0.644^{***}	0.578^{***}	0.324^{***}
	(5.92)	(94.78)	(31.73)	(19.22)
Log assets	-0.075^{**}	-1.937^{***}	0.114	1.149^{***}
	(-2.12)	(-8.41)	(0.71)	(6.88)
Log B/M	-0.998^{***}	-2.682^{***}	-1.103^{***}	-2.113^{***}
	(-27.31)	(-12.10)	(-10.52)	(-14.68)
ROA	0.002	-0.009^{*}	-0.002	0.043^{***}
	(0.35)	(-1.71)	(-0.57)	(11.09)
Leverage	-2.038^{***}	3.347^{***}	-0.652	-0.362
	(-15.54)	(3.58)	(-1.35)	(-0.64)
Log PPE/Assets	-0.032	-1.787^{***}	-0.188	-0.038
	(-0.71)	(-5.01)	(-0.76)	(-0.13)
Firm FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Adj. R^2	0.443	0.551	0.911	0.742
Observations	70,725	63,336	72,807	72,744

Table 7: Biodiversity exposure and cost structure

This table decomposes the operating performance effect of biodiversity exposure into cost components. The sample is a quarterly panel of TRI-reporting firms from 1989 to 2024. Columns (1)–(2) use levels as dependent variables: COGS/Sales (cost of goods sold divided by sales) and SGA/Sales (selling, general, and administrative expenses divided by sales), each in percentage points. Columns (3)–(4) use the year-over-year change in each cost ratio to test whether the effects are levels penalties or escalating drags; coefficients indistinguishable from zero confirm that BPI raises production costs in levels rather than as a trend, while administrative costs are unaffected in both levels and changes. The key independent variable is BPI, standardized within TRI vintage year. Columns (1)–(2) include the one-quarter lagged dependent variable to isolate within-firm changes beyond mean reversion. Controls (log assets, log book-to-market, ROA, leverage, log PPE/Assets) are lagged one quarter. All specifications include firm and year fixed effects. t-statistics based on standard errors clustered by firm appear in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1) COGS/Sales	(2) SGA/Sales	(3) ΔSGA/Sales	(4) ΔCOGS/Sales
BPI (standardized)	0.399***	−0.026	−0.000	−0.025
	(4.47)	(−0.54)	(−1.06)	(−0.30)
Lagged DV	0.578***	0.666***		
	(31.73)	(28.19)		
Log assets	−0.114	−0.320***	0.002***	0.390***
	(−0.71)	(−3.18)	(4.43)	(3.64)
Log B/M	1.103***	0.553***	0.001***	0.525***
	(10.52)	(9.89)	(3.66)	(5.57)
ROA	0.002	0.009***	−0.000***	−0.031***
	(0.57)	(4.76)	(−11.56)	(−9.71)
Leverage	0.652	0.153	−0.008***	−1.530***
	(1.35)	(0.51)	(−4.13)	(−3.26)
Log PPE/Assets	0.188	−0.109	0.001	0.179
	(0.76)	(−0.68)	(0.93)	(1.01)
Firm FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Adj. R^2	0.911	0.929	0.044	0.035
Observations	72,807	73,510	65,433	64,808

Table 8: Analyst-implied returns and biodiversity pressure

This table reports panel regressions of analyst-implied 12-month returns on standardized BPI in a monthly firm-month panel from 1999 to 2024. The dependent variable is ln(MEDPTG/|prc|), the natural log of the median analyst price target divided by the closing stock price, following Engelberg et al. (2020). The price target consensus is from the I/B/E/S Summary file; the closing price is the month-end CRSP value. BPI and Releases are standardized within TRI vintage year. Lagged controls ln(ME) and ln(B/M) are measured one month before the price-target consensus date. The regressions also include ln(# Analysts), the natural log of the number of analysts issuing targets, and PT Dispersion, the standard deviation of price targets across analysts scaled by the median. Column (1) reports OLS without fixed effects; column (2) adds year fixed effects; column (3) adds Fama-French 12-industry fixed effects; column (4) adds standardized Releases as an additional control. The sample is restricted to firms with non-missing BPI and absolute closing price exceeding one dollar. The implied return is winsorized at the 1st and 99th percentiles within TRI vintage year. Standard errors are double-clustered by firm and year-month and reported in parentheses. *, **, *** denote statistical significance at the 10%, 5%, and 1% levels.

	(1)	(2)	(3)	(4)
BPI (std.)	-0.042***	-0.060***	-0.063***	-0.064***
	(0.016)	(0.017)	(0.019)	(0.020)
Releases (std.)				0.013
				(0.011)
ln(ME)	-0.043***	-0.052***	-0.044***	-0.044***
	(0.011)	(0.012)	(0.013)	(0.013)
ln(B/M)	0.083***	0.079***	0.103***	0.088***
	(0.017)	(0.018)	(0.019)	(0.019)
ln(# Analysts)	0.044**	0.027	0.017	0.008
	(0.022)	(0.022)	(0.022)	(0.023)
PT Dispersion	0.414***	0.518***	0.451***	0.483***
	(0.103)	(0.110)	(0.108)	(0.113)
Industry FE	×	×	✓	✓
Year FE	×	✓	✓	✓
Adj. R^2	0.048	0.089	0.101	0.099
N	127,019	127,019	127,019	119,105
Firms	1,103	1,103	1,103	1,045

Table 9: Earnings surprises and announcement returns

Panel A regresses earnings surprise measures on standardized BPI and controls using the quarterly earnings announcement sample. SUE1 is the seasonal random-walk surprise ($EPS_q - EPS_{q-4}$, scaled by rolling standard deviation), following Chan et al. (1996). SUE2 is the analyst-based surprise (actual − consensus mean, scaled by forecast dispersion) from I/B/E/S. Odd-numbered columns in Panel A include year fixed effects only; even-numbered columns add Fama-French 12-industry fixed effects. The attenuation of BPI in the SUE2 specification when industry fixed effects are added (columns 3–4) indicates that analysts partially capture biodiversity-related costs through industry peer benchmarking. Panel B regresses three-day $[-1, +1]$ cumulative abnormal returns (CARs) on standardized BPI and controls. Abnormal returns are market-adjusted and winsorized at the 1st and 99th percentiles. Column (8) restricts to the I/B/E/S universe where all surprise measures are available. BPI and Releases are standardized within TRI vintage year. Standard errors in parentheses, clustered by firm in Panel A and double-clustered by firm and year-month in Panel B. *, **, *** denote significance at 10%, 5%, 1%.

	Panel A: BPI predicts surprises				*Panel B: Announcement returns*			
	SUE1		SUE2		CAR $[-1, +1]$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
BPI (std.)	0.056***	0.065***	0.102**	0.076*	0.124***	0.078**	0.053	0.074
	(0.011)	(0.011)	(0.040)	(0.042)	(0.038)	(0.039)	(0.040)	(0.059)
Releases (std.)	-0.001	-0.001	-0.107***	-0.070**				-0.012
	(0.008)	(0.009)	(0.028)	(0.030)				(0.043)
SUE1						0.537***		0.120***
						(0.026)		(0.023)
SUE2							0.653***	0.606***
							(0.016)	(0.021)
LTG								0.020***
								(0.004)
ln(1+Analysts)	-0.069***	-0.068***	0.043	-0.007				0.057
	(0.010)	(0.010)	(0.046)	(0.047)				(0.093)
ln(ME)	0.052***	0.052***	0.130***	0.163***	-0.023	-0.039*	-0.127***	-0.074*
	(0.008)	(0.008)	(0.027)	(0.028)	(0.021)	(0.022)	(0.026)	(0.043)
ln(B/M)	-0.188***	-0.200***	-0.324***	-0.279***	-0.061	0.044	0.043	0.092
	(0.014)	(0.014)	(0.040)	(0.041)	(0.049)	(0.051)	(0.051)	(0.058)
Industry FE	×	✓	×	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Adj. R^2	0.039	0.039	0.048	0.052	0.002	0.017	0.103	0.095
N	66,633	66,624	54,550	54,549	74,161	70,959	57,596	28,815
Firms	1,632	1,631	1,448	1,448	1,806	1,723	1,531	1,187

Table 10: Earnings announcement realization and post-announcement drift

This table examines how much of the annual BPI premium is realized during earnings announcements and how the High−Low BPI spread evolves at extended post-announcement windows. Panel A reports the EA realization fraction by size tercile. For each size group, the annualized EA return is four times the quarterly $[-1,+1]$ High−Low CAR spread; the realization fraction is this annualized EA return divided by the annual portfolio alpha from the FF5+Mom model in Table 4 (small: $0.622\% \times 12 = 7.46\%$; mid: $0.462\% \times 12 = 5.54\%$), which uses a more conservative factor specification than the baseline FF3 alpha. Panel B reports the High−Low BPI tercile spread in cumulative abnormal returns at windows from $[-1,+1]$ through $[-1,+60]$ trading days, using the full sample and by size tercile. CARs are market-adjusted. Portfolio assignments use lagged BPI terciles within TRI vintage year. Size terciles are formed within calendar year using lagged market capitalization. Newey-West (4 lags) t-statistics in parentheses. Panel A tercile means are time-series averages of annual cross-sectional means. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Panel A: EA Realization Fraction by Size Tercile

	T1 (Low)	T3 (High)	H−L	Annual α	Fraction
Small	−0.067	0.678	0.745***	7.46%	39.9%
			(2.95)		
Mid	0.143	0.425	0.283	5.54%	20.4%
			(0.93)		
Big	0.156	0.291	0.135	−2.10%	—
			(1.17)		

Panel B: Post-Announcement Drift — H−L BPI Spread (%)

	$[-1,+1]$	$[-1,+5]$	$[-1,+20]$	$[-1,+40]$	$[-1,+60]$
Full sample	0.415***	0.500***	0.697***	1.478***	2.183***
	(4.36)	(4.69)	(3.60)	(4.77)	(5.12)
Small	0.745***	0.932***	1.036***	2.122***	2.474***
	(3.29)	(3.55)	(2.63)	(3.41)	(3.33)
Mid	0.283	0.305	−0.011	0.071	0.266
	(0.49)	(0.47)	(−0.19)	(0.06)	(0.66)
Big	0.135	0.373	0.621	0.320	0.583
	(0.81)	(1.35)	(1.48)	(−0.72)	(0.72)
N (events)	77,320	77,297	77,151	76,734	76,250

Table 11: Pollution prevention and the BPI cost channel

This table reports regressions of COGS/Sales (%) on BPI, an abatement indicator, and their interaction. Abatement equals one if at least one of the firm's facilities filed a Pollution Prevention (P2) source reduction report with the EPA's Toxic Release Inventory in the concurrent TRI reporting year. BPI is standardized within TRI vintage year. Columns (1) and (3) report pooled OLS; columns (2) and (4) add firm and year fixed effects. t-statistics based on standard errors clustered by firm appear in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	COGS/Sales$_t$		COGS/Sales$_{t+4}$	
	(1)	(2)	(3)	(4)
BPI (standardized)	0.514***	0.484***	0.548***	0.562***
	(9.61)	(5.07)	(6.02)	(3.45)
Abatement	−0.015	−0.044	−0.049	−0.167
	(−0.29)	(−0.53)	(−0.51)	(−1.20)
BPI × Abatement	−0.139**	−0.206**	−0.194*	−0.240
	(−2.43)	(−2.14)	(−1.91)	(−1.52)
Lagged DV	0.912***	0.578***	0.882***	0.370***
	(164.58)	(31.67)	(99.58)	(17.15)
Log assets	0.081***	−0.114	−0.038	0.231
	(3.15)	(−0.71)	(−0.83)	(0.89)
Log B/M	0.948***	1.106***	1.155***	0.863***
	(13.89)	(10.56)	(10.75)	(5.60)
ROA	0.027***	0.002	0.059***	0.020***
	(10.30)	(0.58)	(13.16)	(5.49)
Leverage	1.413***	0.662	2.647***	0.690
	(5.93)	(1.36)	(5.92)	(0.85)
Log PPE/Assets	0.255***	0.190	0.449***	0.378
	(4.15)	(0.77)	(4.22)	(0.95)
Firm FE	×	✓	×	✓
Year FE	×	✓	×	✓
Adj. R^2	0.886	0.911	0.795	0.856
Observations	72,831	72,807	62,897	62,867

Table 12: CDP biodiversity disclosure and continuous BPI

This table tests whether firms with greater biodiversity exposure are less likely to voluntarily disclose biodiversity information through the CDP survey. The dependent variables are five binary indicators constructed from CDP biodiversity questionnaire responses, coded as one if the firm answered "Yes" and zero otherwise. *Near Bio Areas* indicates operations near biodiversity-important areas (column 1). *Indicators* indicates the firm monitors biodiversity indicators (column 2). *Actions* indicates the firm has taken biodiversity actions or commitments (column 3). *Commitments* indicates biodiversity-specific public commitments (column 4). *Vol. Disc.* indicates voluntary publication of biodiversity information (column 5). *BPI* (std) is ln(bird abundance/total assets), standardized cross-sectionally. *Releases* (std) is ln(TRI releases/total assets), standardized analogously. *Enforcement* ($t-1$) is an indicator equal to one if the firm faced an EPA enforcement action with a monetary penalty in the prior year. *Abatement* is an indicator equal to one if any of the firm's TRI facilities reported pollution prevention activities. *BPI* × *Abatement* is the product of standardized BPI and the abatement indicator. Controls include ln(Assets), *Leverage*, *ROA*, ln(PPE/Assets), and ln(B/M). All regressions include Fama–French 12-industry and year fixed effects. t-statistics are in parentheses, computed from standard errors clustered by firm. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1) Near Bio	(2) Indicators	(3) Actions	(4) Commitments	(5) Vol. Disc.
Enforcement ($t-1$)	-0.344^{***}	$+0.702^{***}$	$+0.535^{***}$	-0.024	-0.425^{***}
	(-3.96)	(9.58)	(6.41)	(-0.09)	(-4.51)
Abatement	-0.031	-0.003	$+0.028$	$+0.078$	$+0.031$
	(-0.73)	(-0.08)	(0.62)	(1.62)	(0.72)
BPI (std)	-0.050	-0.034	-0.096^{**}	-0.056	-0.076^{*}
	(-1.09)	(-0.79)	(-2.08)	(-1.28)	(-1.84)
BPI × Abatement	$+0.044$	$+0.084^{**}$	$+0.082^{**}$	$+0.041$	$+0.055$
	(1.12)	(2.44)	(2.09)	(0.92)	(1.39)
Releases (std)	$+0.021$	$+0.037$	$+0.027$	$+0.028$	$+0.023$
	(0.82)	(1.50)	(0.99)	(1.07)	(1.03)
ln(Assets)	$+0.027$	$+0.051^{**}$	$+0.033$	$+0.011$	$+0.016$
	(1.02)	(2.06)	(1.23)	(0.41)	(0.68)
ln(B/M)	-0.024	-0.043	-0.061^{*}	-0.020	-0.048^{*}
	(-0.70)	(-1.34)	(-1.76)	(-0.62)	(-1.75)
ROA	-0.000	$+0.000$	-0.002^{*}	$+0.000$	-0.000
	(-0.35)	(0.27)	(-1.70)	(0.28)	(-0.41)
ln(PPE/Assets)	$+0.104^{***}$	$+0.107^{***}$	$+0.157^{***}$	$+0.123^{***}$	$+0.145^{***}$
	(3.06)	(3.15)	(4.35)	(3.46)	(4.70)
Leverage	-0.031	-0.032	-0.258	-0.186	-0.150
	(-0.16)	(-0.22)	(-1.33)	(-1.18)	(-0.95)
Industry FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓
R^2	0.189	0.209	0.219	0.169	0.143
N	629	853	848	918	997
Firms	276	299	301	310	320

Table 13: MSCI rating components and ecological exposure

This table tests whether MSCI rating components encode ecological exposure. The sample is an annual panel (2007–2024) of TRI firms matched to MSCI ratings. *E Score* is the MSCI environmental pillar score. *Biodiv Score* is the biodiversity and land use key issue score. *Biodiv Exp.* is the biodiversity exposure sub-score. *Biodiv Mgmt.* is the biodiversity management sub-score. *Biodiv Wt.* is the weight assigned to the biodiversity key issue in the overall rating. BPI and Releases are standardized within TRI vintage year. CDP Disclosure equals one if the firm responds to the Carbon Disclosure Project biodiversity questionnaire. Abatement equals one if the firm reports source reduction activities to TRI. All specifications include firm and industry×year fixed effects. Controls: ln(ME), ln(B/M), ROA, leverage, ln(assets), ln(PPE/assets). Standard errors clustered by firm in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1) E Score	(2) Biodiv Score	(3) Biodiv Exp.	(4) Biodiv Mgmt.	(5) Biodiv Wt.
BPI (std)	0.012	0.155	-0.227**	-0.355	-0.883**
	(0.088)	(0.249)	(0.108)	(0.241)	(0.386)
Releases (std)	-0.023	0.612*	-0.097	0.375	0.089
	(0.111)	(0.346)	(0.069)	(0.332)	(0.141)
Abatement (P2)	0.028	-0.003	0.057	-0.005	-0.179
	(0.069)	(0.283)	(0.073)	(0.207)	(0.153)
BPI (std) × Abatement	-0.012	0.015	-0.060	0.097	0.382**
	(0.070)	(0.223)	(0.064)	(0.147)	(0.173)
Disclosure (CDP)	0.129	1.067***	0.160	-0.087	0.065
	(0.149)	(0.378)	(0.129)	(0.431)	(0.191)
BPI (std) × CDP	0.083	0.635**	0.061	0.380	-0.128
	(0.168)	(0.274)	(0.102)	(0.565)	(0.161)
ln(ME)	-0.036	0.049	0.030	0.339	-0.048
	(0.153)	(0.431)	(0.154)	(0.328)	(0.392)
ln(B/M)	0.032	0.074	0.121	0.205	-0.413
	(0.131)	(0.390)	(0.131)	(0.297)	(0.289)
ROA	-0.001	-0.002	-0.000	-0.004	0.000
	(0.001)	(0.003)	(0.001)	(0.004)	(0.002)
Leverage	0.184	1.132	0.263	2.050	2.516
	(0.560)	(2.074)	(0.648)	(1.772)	(1.662)
ln(Assets)	0.283	0.028	0.424*	-0.557	-0.541
	(0.219)	(0.729)	(0.239)	(0.453)	(0.664)
ln(PP&E/Assets)	-0.017	-1.193	0.268	-0.307	-0.231
	(0.191)	(0.730)	(0.174)	(0.479)	(0.600)
Firm FE	✓	✓	✓	✓	✓
Ind.-Year FE	✓	✓	✓	✓	✓
Adj. R^2	0.524	0.709	0.659	0.535	0.967
Observations	6,283	495	3,453	495	1,399

Table 14: Selection into MSCI biodiversity coverage

Logit selection model for MSCI biodiversity score coverage. The dependent variable equals one if MSCI reports a non-missing biodiversity and land use score for a given firm-year (2.6% of observations). Average marginal effects reported. BPI is cross-sectionally standardized within TRI vintage year. CDP Disclosure equals one if the firm responds to the Carbon Disclosure Project biodiversity questionnaire. Abatement equals one if the firm reports source reduction activities to TRI. Column (3) adds only ln(ME) to the fixed effects specification. Columns (4)–(5) add the full control set: lagged ln(ME), ln(B/M), ROA, leverage, ln(assets), ln(PPE/assets). Annual panel, firm-clustered standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1)	(2)	(3)	(4)	(5)
BPI (std)	-0.011***	-0.027***	-0.023*	-0.009	-0.016
	(0.003)	(0.009)	(0.012)	(0.013)	(0.012)
CDP Disclosure	0.031***	0.023	0.020	-0.007	0.005
	(0.008)	(0.025)	(0.024)	(0.025)	(0.024)
Abatement	-0.019***	-0.002	-0.006	-0.011	-0.000
	(0.005)	(0.017)	(0.017)	(0.016)	(0.017)
ln(ME)			0.004	0.044*	0.051**
			(0.006)	(0.024)	(0.025)
ln(B/M)				0.061**	0.065**
				(0.025)	(0.027)
ROA				-0.000***	-0.000*
				(0.000)	(0.000)
Leverage				-0.056	-0.033
				(0.080)	(0.083)
ln(Assets)				-0.024	-0.031
				(0.023)	(0.025)
ln(PPE/Assets)				0.103***	0.084***
				(0.018)	(0.019)
Year FE	×	✓	✓	✓	✓
Industry FE	×	✓	✓	×	✓
Pseudo R^2	0.042	0.225	0.226	0.221	0.295
N	22,394	5,951	5,951	6,234	5,425

Table 15: Enforcement CARs by penalty status

This table reports cumulative abnormal returns (CARs) around EPA enforcement actions against TRI-reporting facilities, split by whether the action carried a monetary penalty. Each event is an EPA enforcement conclusion linked to a publicly traded firm via the TRI facility registry. Daily abnormal returns are Fama–French three-factor residuals from a market model estimated over trading days $[-250, -31]$ relative to the event, using log stock returns. CARs are summed log abnormal returns over each window, expressed in percentage points. Firms are sorted into within-year BPI terciles based on their lagged (year $-$ 1) biodiversity pressure index. Panels A and B restrict to post-2000 events; Panels C and D use the full sample (1988–2024). Panels A and C include events with a recorded combined penalty amount greater than zero; Panels B and D include events with no monetary penalty. H–L is the difference between T3 and T1 means. t-statistics based on standard errors clustered by firm appear in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	T1 (Low)	T2	T3 (High)	H–L	t-stat
Panel A: Post-2000, with monetary penalty (366 events, 193 firms)					
CAR $[-1, +1]$	-0.40	-0.45	-0.08	0.32	(0.42)
CAR $[-5, +5]$	-0.68	1.21	$-2.01*$	-1.32	(-0.98)
CAR $[-10, +30]$	-0.49	0.52	$-4.92**$	-4.43	(-1.64)
CAR $[-20, +45]$	-0.04	0.24	$-5.45*$	-5.40	(-1.56)
CAR $[-20, +60]$	0.11	1.83	$-7.67**$	$-7.78**$	(-2.08)
N events	134	115	106		
Panel B: Post-2000, without monetary penalty (3177 events, 549 firms)					
CAR $[-1, +1]$	$0.20*$	0.02	0.13	-0.07	(-0.31)
CAR $[-5, +5]$	0.09	-0.09	-0.14	-0.22	(-0.68)
CAR $[-10, +30]$	-0.00	$-0.84*$	-0.68	-0.67	(-1.06)
CAR $[-20, +45]$	0.08	-0.96	-1.10	-1.18	(-1.25)
CAR $[-20, +60]$	-0.11	-0.74	-1.25	-1.14	(-1.07)
N events	952	1176	967		
Panel C: Full sample, with monetary penalty (1056 events, 347 firms)					
CAR $[-1, +1]$	-0.09	-0.08	-0.12	-0.03	(-0.09)
CAR $[-5, +5]$	-0.04	0.27	-0.67	-0.63	(-1.08)
CAR $[-10, +30]$	0.48	0.48	-1.05	-1.53	(-1.29)
CAR $[-20, +45]$	0.47	-0.06	-1.53	-2.00	(-1.30)
CAR $[-20, +60]$	0.56	0.20	-2.32	$-2.88*$	(-1.71)
N events	400	352	284		
Panel D: Full sample, without monetary penalty (3272 events, 566 firms)					
CAR $[-1, +1]$	0.17	0.01	0.17	0.00	(0.01)
CAR $[-5, +5]$	0.14	-0.09	-0.12	-0.27	(-0.82)
CAR $[-10, +30]$	0.03	$-0.89**$	-0.59	-0.62	(-0.98)
CAR $[-20, +45]$	-0.01	$-1.01*$	-1.03	-1.01	(-1.09)
CAR $[-20, +60]$	-0.22	-0.82	-1.24	-1.02	(-0.97)
N events	998	1204	987		

Appendix

This appendix collects results that support the main analysis. Section A.1 validates the satellite-inferred construction dates with visual confirmation and tests the spatial identification strategy. Section A.2 decomposes BPI to rule out a mechanical size artifact. Section A.3 presents the joint horse race between BPI and toxic releases. Section A.4 tests whether the regulatory pipeline responds to ecological context, documenting that EPA enforcement targets emissions rather than ecology.

A.1 Ecological validation checks

The staggered difference-in-differences design in Section 3 anchors the causal estimates to satellite-inferred construction dates rather than to the delayed administrative event of a facility's first TRI filing. Two identification concerns require direct validation. First, the construction dates must be correctly identified. Second, the bird abundance decline must reflect genuine ecological damage rather than spatial displacement.

Construction date detection. A facility's true construction year often precedes its first regulatory filing by several years, so I extract construction timing directly from satellite imagery. The National Land Cover Database (NLCD) reports the fraction of each 30-meter pixel covered by impervious surfaces such as pavement, concrete, and rooftops. For each facility, I sum impervious area within a 500-meter buffer of the TRI-reported coordinates, compute annual log growth rates, and flag the year of maximum growth as the inferred construction date, restricting to pre-reporting years so the algorithm captures the true buildup. I retain facilities where the inferred date falls within three years before the first TRI report, which focuses the analysis on new facilities rather than legacy sites.

Figure A.1 validates this procedure across all facilities. Average impervious surface growth spikes at the detected construction year (year 0) and is near zero in surrounding years, confirming that the algorithm isolates a discrete construction event rather than noisy

year-to-year variation.

Figure A.2 provides visual corroboration. I obtain high-resolution (∼1 m) aerial imagery from the USDA National Agriculture Imagery Program (NAIP) before and after the detected construction dates. Each image pair is warped to an identical geographic grid so that landscape features align across panels, covering approximately 1.8 km × 1.8 km centered on the facility's TRI-reported coordinates. The images confirm that the statistical jumps in impervious surface correspond to physical land clearing and infrastructure development. Panel (a) shows Oklahoma cropland converted to an industrial site around 2017; Panel (b) shows a South Carolina field replaced by a tire manufacturing plant around 2013.

Displacement test. The bird abundance decline documented in Section 3 could reflect spatial displacement rather than genuine ecological damage. If birds flee the impact zone, abundance should rise in rings just outside it. I test for this pattern using a staggered difference-in-differences design in which the treated unit is a series of six-kilometer sliding rings at increasing distances from the facility rather than the inner circle. For each start distance $X \in \{5, \ldots, 20\}$, I define the treated ring as $[X, X{+}5]$ km and use the far edge $[25, 30]$ km as a stable reference group, keeping the same construction-year cohort assignment and event window $[-6, +5]$ as in the main specification. The Callaway and Sant'Anna (2021) estimator is applied with doubly-robust imputation and facility-level clustering.

Figure A.3 reports the results. Bird abundance declines significantly in rings out to approximately 15 km, with peak damage of −10.1% at 6–11 km. No ring shows a significant positive effect at any distance. The losses in the near-field are not offset by gains farther out, ruling out spatial reshuffling as an alternative explanation. The smooth monotonic attenuation with distance is consistent with genuine ecological degradation from a site-specific source.

Observer selection test. Citizen science data raises the concern that facility construction alters the local landscape in ways that shift where birders choose to survey, independently

of actual ecological change. I test this directly by replacing the bird abundance outcome with log(checklist count+1) and re-estimating the primary staggered difference-in-differences specification. If facilities repel observers from the inner ring, checklist density should decline differentially after construction. The ATT on log checklist density is −0.024, economically negligible relative to the 6–9% bird abundance decline and statistically indistinguishable from zero. The pre-versus-post equality test also fails to reject. Observer spatial sorting cannot account for the ecological damage documented in Section 3.

A.2 Alternative BPI constructions and the ecological numerator

BPI equals log(bird abundance/total assets), which decomposes additively into log(birds) − log(assets). One concern is that the return-predictive power of BPI derives from the asset-scaling denominator rather than the ecological numerator. Table A.1 addresses this by replacing BPI with its unscaled numerator, log(birds), in Fama-MacBeth regressions with full controls.

With standard controls (columns 1–2), log(birds) is statistically significant but economically weaker than BPI. The key test is columns 3–4, which add log(assets) as a separate control, allowing both components of BPI to carry distinct coefficients. log(birds) strengthens in column 4, while log(assets) enters with a negative coefficient, confirming that both the ecological numerator and the intensity denominator contribute independently to return prediction.

Three additional pieces of evidence rule out a mechanical size explanation. First, toxic releases intensity shares the same 1/assets denominator yet BPI retains its predictive power after controlling for releases in every specification (Table 5, columns 5–7). Second, the denominator captures exposure *intensity* rather than a disguised size effect. log(assets) is significant alongside log(birds) because it measures how much damage matters per dollar of assets, not because smaller firms mechanically earn higher returns. Third, columns (6)–(7) of Table A.1 replace total assets with alternative denominators. $\text{BPI}_{sales} \equiv \ln(\text{birds/sales})$ produces an

FM coefficient of 0.168, comparable to the asset-scaled BPI. $\text{BPI}_{PP\&E} \equiv \ln(\text{birds/PP\&E})$ yields 0.103. All three denominators produce statistically significant premiums. Because assets, sales, and PP&E are correlated proxies for firm size, this test confirms that the effect is not uniquely sensitive to the choice of scaling variable, but does not by itself isolate the ecological numerator's contribution. The randomization test below provides that isolation. Conditional double sorts on size $\times$ BPI_{sales} terciles yield pooled FF3 alphas of 33 bps/month, concentrated in small and mid-cap firms (FF3 $\alpha = 51$ bps for small; 44 bps for mid).

A randomization inference test directly isolates the ecological numerator. For each of 1,000 iterations, I shuffle bird abundance across firms within the same Fama-French 12 industry and vintage year, preserving the cross-sectional distribution but breaking the link between a firm and the ecology surrounding its facilities. I recompute placebo BPI from the shuffled birds and estimate the baseline Fama-MacBeth specification with log(assets) as an additional control. The placebo t-statistics center on zero, and only 39 of the 1,000 placebos exceed the actual, yielding a one-sided empirical p-value of 0.039 (two-sided: 0.078). I use a one-sided test because the directional prediction (that firms located near richer ecologies earn higher returns) is established ex ante by the theoretical framework. The ecology of a firm's facility locations predicts returns beyond what asset scaling, industry membership, and standard controls explain.

A.3 BPI versus toxic releases in Fama-MacBeth regressions

Table A.2 reports Fama-MacBeth regressions of BPI and release intensity, extending the specification in Table 5 with size interactions. When both variables enter jointly (column 3), BPI retains a significant coefficient, indicating that the ecological dimension adds incremental pricing information. Column 4 adds interactions with log market equity; the slope-difference t-statistic tests whether the BPI and Releases size-interaction slopes are equal, and they are indistinguishable. Column 5 adds idiosyncratic volatility interactions, and BPI carries incremental return-predictive power beyond releases in all specifications.

A.4 Regulatory burden and ecological exposure

Table A.3 tests whether facilities near ecologically rich areas face greater regulatory attention from the EPA. The sample is a facility-year panel of 14,681 TRI-reporting facilities from 1987 to 2024, matched to bird abundance in the 15–30 km control ring. The table spans the full regulatory pipeline: EPA inspections (columns 1–2), informal enforcement actions such as warning letters and notices of violation (columns 3–4), the probability of any formal penalty (columns 5–6), and penalty amounts (columns 7–8). Odd-numbered columns include EPA region, industry, and year fixed effects. Even-numbered columns add log toxic releases, isolating whether ecological context attracts regulatory attention beyond what emissions alone predict.

Bird abundance has no predictive power for any outcome once emissions are controlled. Across all four measures of regulatory attention, the coefficient on ecological richness is economically small and statistically indistinguishable from zero. Log releases, by contrast, is highly significant for every outcome. The regulatory apparatus responds to what firms emit, not to the ecological context in which they emit. This result reinforces the main finding: ecological exposure raises production costs through a channel that the entire intermediary chain (financial analysts, ESG raters, corporate disclosures, and regulators) treats as indistinguishable from ordinary operational expense.

Appendix Figures

Figure A.1: NLCD validation of satellite-inferred construction dates

Average growth in impervious surface area within 500 meters of TRI facility coordinates, measured annually from the National Land Cover Database (NLCD). The x-axis plots years relative to the satellite-inferred construction date (year 0), identified by the largest discrete jump in built-up area. The spike at year 0 confirms that the jump detection algorithm successfully isolates the construction event. Error bars show 95% confidence intervals. Growth is near zero in pre-construction years and decays gradually post-construction as surrounding development fills in. $N = 3{,}583$ facilities.

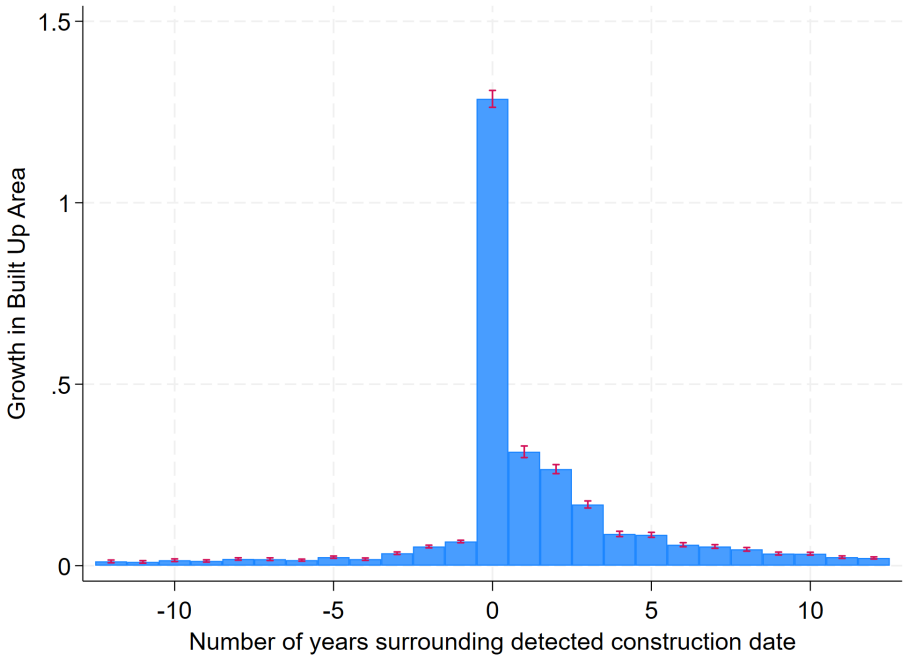
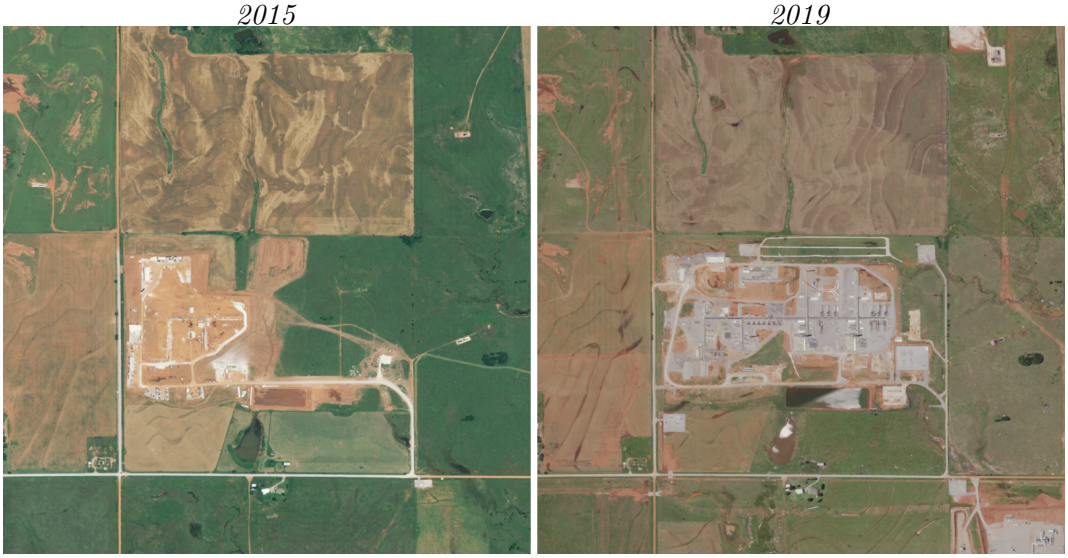


Figure A.2: Visual validation of satellite-inferred construction dates

High-resolution (∼1 m) aerial imagery from the USDA National Agriculture Imagery Program (NAIP) before and after the construction years detected by the NLCD impervious surface algorithm (Section 3; Figure A.1). Each pair is warped to an identical geographic grid so that landscape features align across panels. Each image covers approximately 1.8 km × 1.8 km centered on the facility's TRI-reported coordinates.

(a) Oklahoma, cropland to industrial (built ∼2017)



(b) South Carolina, field to tire manufacturing plant (built ∼2013)

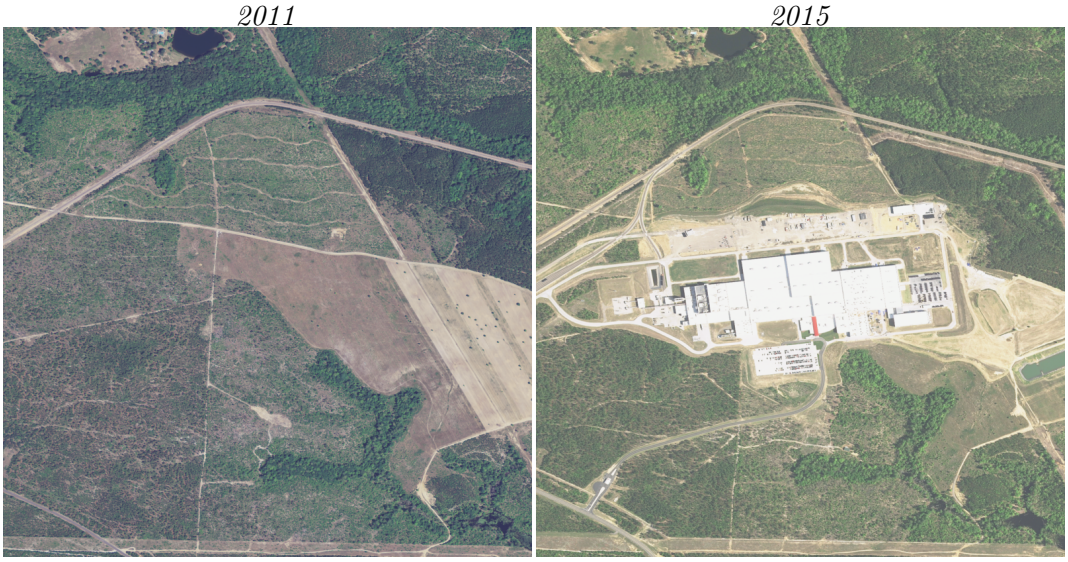
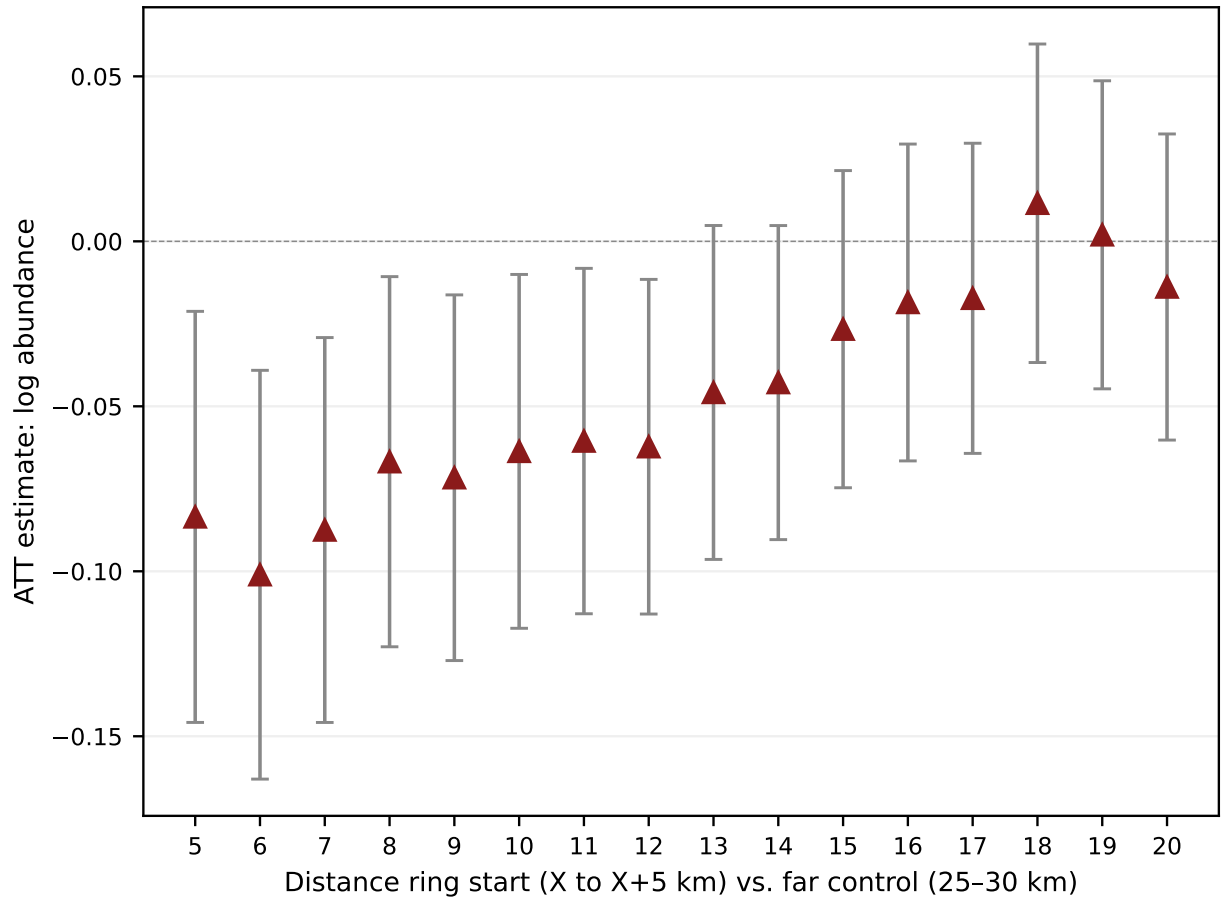


Figure A.3: Spatial attenuation of ecological decline

Callaway and Sant'Anna (2021) staggered difference-in-differences estimates of the effect of TRI facility construction on log bird abundance in 5 km-wide sliding rings at increasing distances from the facility. Each ring is compared against a far control ring (25–30 km). The specification includes month-of-year indicators, observer counts, birding effort, and urbanization growth as controls. Abundance declines significantly out to approximately 15 km. Peak damage is -10.1% at 6–11 km ($z = -3.20$). No ring shows a significant positive effect, ruling out displacement of birds to adjacent areas. $N = 3{,}623$ facilities with construction year ≥ 1990; event window $[-6, +5]$ years; clustering by facility.



Appendix Tables

Table A.1: Alternative BPI constructions and the ecological numerator

This table tests whether the return-predictive power of BPI depends on the choice of scaling denominator and whether it survives a comprehensive set of size-related controls. Columns (1)–(4) replace BPI with its unscaled numerator, ln(birds), cross-sectionally standardized each month; columns (3)–(4) add ln(assets) as a separate control. Column (5) reports the baseline BPI $\equiv$ ln(birds/assets) from Table 5. Column (6) replaces total assets with trailing twelve-month sales: $\text{BPI}_{sales} \equiv$ ln(birds/sales). Column (7) uses net PP&E: $\text{BPI}_{PP\&E} \equiv$ ln(birds/PP&E). Column (8) adds the full battery of size-related controls to the baseline specification: ln(ME)^2, Amihud illiquidity, share turnover, idiosyncratic volatility, and ln(1 + Analysts). All alternative measures are cross-sectionally standardized. All columns include log book-to-market, momentum, log market equity, ROA, leverage, and Fama-French 12-industry fixed effects. Newey-West(6) t-statistics in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ln(birds) (std.)	+0.058**	+0.058**	+0.090***	+0.097***				
	(2.17)	(2.23)	(3.35)	(3.65)				
ln(assets)			−0.725***	−0.821***				
			(−4.41)	(−5.29)				
BPI_{assets} (std.)					+0.180***			+0.198***
					(4.06)			(4.43)
BPI_{sales} (std.)						+0.168***		
						(4.31)		
$\text{BPI}_{PP\&E}$ (std.)							+0.103**	
							(2.12)	
Log book-to-market	✓	✓	✓	✓	✓	✓	✓	✓
Momentum	✓	✓	✓	✓	✓	✓	✓	✓
Log ME	✓	✓	✓	✓	✓	✓	✓	✓
(Log ME)^2	×	×	×	×	×	×	×	✓
ROA, Leverage	✓	✓	✓	✓	✓	✓	✓	✓
Idiosyncratic volatility	×	✓	×	✓	×	×	×	✓
Amihud, Turnover	×	×	×	×	×	×	×	✓
ln(1 + Analysts)	×	×	×	×	×	×	×	✓
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓
N months	428	428	428	428	428	428	324	428
Avg. firms/month	489	489	489	489	489	489	386	453

Table A.2: BPI versus toxic releases

Monthly Fama-MacBeth regressions of excess returns on lagged firm characteristics. All variables of interest are cross-sectionally standardized each month. Interaction terms are products of standardized constituents. Controls include log book-to-market, cumulative returns ($t-12$ to $t-2$, labeled "Momentum"), and Fama-French 12-industry fixed effects. "Size" in interactions denotes standardized log market equity. Coefficients are in percent per month. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1)	(2)	(3)	(4)	(5)
BPI	+0.269***		+0.241***	+0.242***	+0.265***
	(6.06)		(5.00)	(4.50)	(5.51)
Releases		+0.107***	+0.038	+0.038	+0.036
		(3.34)	(1.15)	(1.14)	(1.10)
Log ME				+0.013	+0.157**
				(0.15)	(2.15)
Idiovol					+0.241**
					(2.38)
Momentum	−0.243	−0.187	−0.222	−0.252	−0.175
	(−0.95)	(−0.73)	(−0.86)	(−0.98)	(−0.74)
BPI × Size				+0.002	+0.010
				(0.05)	(0.27)
Rel × Size				+0.004	+0.056
				(0.12)	(1.56)
BPI × Idiovol					−0.018
					(−0.36)
Rel × Idiovol					+0.091
					(1.39)
Industry FE	✓	✓	✓	✓	✓
N months	428	428	428	428	428
Avg. firms/month	490	458	458	458	458
Avg. R^2	0.103	0.104	0.110	0.128	0.150

Table A.3: Regulatory burden and ecological exposure

This table tests whether facilities near ecologically rich areas face greater regulatory attention, conditional on toxic emissions. The sample is a facility-year panel of 14,681 TRI-reporting facilities from 1987 to 2024, matched to bird abundance in the 15–30 km control ring from eBird citizen science data. The dependent variables span the full regulatory pipeline: annual EPA inspection counts (columns 1–2), informal enforcement action counts (columns 3–4), an indicator for any formal penalty (columns 5–6), and the log of one plus total penalty amount (columns 7–8). Ecological richness is the log effort-adjusted bird abundance in the 15–30 km ring surrounding each facility, the same measure used to construct BPI. Log releases is the natural log of one plus total toxic releases reported to the TRI. All specifications include EPA region, two-digit NAICS industry, and year fixed effects. All coefficients and standard errors are multiplied by 100 for readability; this corresponds to percentage points for Prob(penalty) and to percent change for Log(1+penalty). Standard errors clustered by facility appear in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

	Inspections		Informal actions		Prob(penalty)		Log(1+penalty)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ecological richness	0.15	0.03	0.01	-0.01	0.03	-0.01	0.36	0.14
	(0.09)	(0.10)	(0.01)	(0.01)	(0.05)	(0.05)	(0.27)	(0.26)
Log releases		0.26***		0.02***		0.09***		0.24***
		(0.03)		(0.01)		(0.01)		(0.04)
EPA Region FE	✓	✓	✓	✓	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
N	270,889	270,889	270,889	270,889	270,889	270,889	270,889	270,889
Facilities	14,681	14,681	14,681	14,681	14,681	14,681	14,681	14,681

Internet Appendix

eBird effort adjustment methodology

Section 2.1 summarizes the effort-adjustment procedure used to construct bird abundance measures from eBird checklists. I provide the full specification here.

The unit of observation is the eBird checklist. I keep complete checklists (all species reported) under structured protocols (Traveling, Stationary, Area) and deduplicate shared group outings. For each checklist i, I sum species counts to obtain total birds and take the log,

$$y_i = \log(\text{TotalBirds}_i + 1). \tag{5}$$

I then partial out observer effort and predictable detectability by estimating

$$y_i = \beta_1 \,\text{Hours}_i + \beta_2 \,\text{Observers}_i + \delta_h + \gamma_{s,m} + \varepsilon_i, \tag{6}$$

where δ_h is an hour-of-day fixed effect and $\gamma_{s,m}$ is a state-by-month fixed effect. The residual $\hat{\varepsilon}_i$ is the effort-adjusted log abundance for checklist i. Because the model includes no local fixed effects below the state level, $\hat{\varepsilon}_i$ retains all within-state spatial variation in bird populations. In a biologically rich area the average residual is positive, and in a degraded area it is negative.

For each facility f in year t, I average the residuals across all checklists within the ring and exponentiate to obtain the facility-level abundance measure.

$$\text{BirdAbundance}_{ft} = \exp\left(\frac{1}{N_{ft}} \sum_{i=1}^{N_{ft}} \hat{\varepsilon}_i\right), \tag{7}$$

where N_{ft} is the number of checklists in the 15–30 km control ring around facility f in year t. Facility-years with no checklists are treated as missing. Coverage grows from approximately 7,600 checklists per year in the mid-1980s to over 6.5 million by 2024, so the fraction of

missing facility-years declines over the sample.

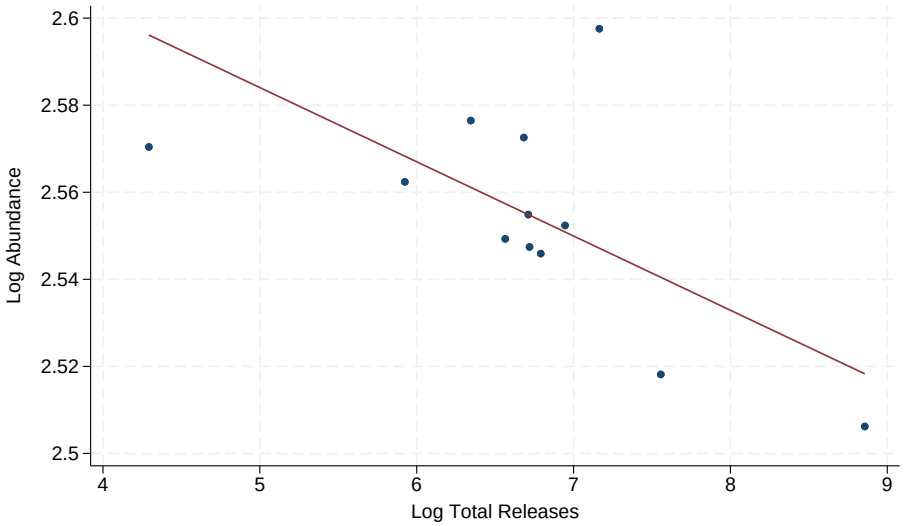
Because effort-adjusted abundance is a residual used as a dependent variable in subsequent regressions, the first-step regressors must appear in those specifications to avoid the biased coefficients documented by Chen et al. (2018). The first stage is estimated at the checklist level, while the second stage is aggregated to the facility-ring-year level. I include aggregated versions of the first-stage regressors (observer counts, birding hours, and month-of-year indicators) as controls in the difference-in-differences specifications of Section 3. Table 3 shows that the core results are robust regardless of whether these controls are included (comparing columns 1 and 4).

Section A: Ecological Validation

The printed Appendix validates the satellite-inferred construction dates, provides visual confirmation, and rules out spatial displacement. This section provides three additional pieces of supporting evidence. First, a within-facility dose-response test confirms that years with higher toxic releases produce lower bird abundance at the same site (IA Figure IA.1). Second, ambient water-quality monitoring shows elevated lead concentrations near TRI facilities (IA Figure IA.2). Third, randomization inference confirms that the lead effects are correctly timed rather than artifacts of generic trends (IA Figure IA.3).

Figure IA.1: Within-facility dose-response in bird abundance

Binscatter of log effort-adjusted bird abundance on log total toxic releases within the 1–14 km impact zone around TRI facilities, absorbing facility fixed effects. Each point represents a dodecile (one of twelve equal-sized bins) of the within-facility distribution of log releases, averaging approximately 1,975 facility-year observations per bin. The negative slope indicates that, conditional on all time-invariant facility characteristics, years with higher reported emissions are associated with lower proximate bird abundance ($\beta = -0.008$, $t = -2.14$; $N = 23{,}693$ facility-years across 2,456 facilities). Ecological degradation scales with the intensity of industrial activity within a given facility over time.



Lead contamination pathway. I use ambient water-quality monitoring data to test whether TRI facility construction raises local lead concentrations. I match each TRI facility to nearby monitoring stations, restrict to routine water samples, harmonize units to μg/L, and summarize lead within each station-year using the median to reduce sensitivity to extreme readings. I then implement the same ring-based staggered DiD that I use for bird abundance, replacing the outcome with median lead concentration. For each distance threshold $X \in \{5, \ldots, 24\}$ km, the inner ring spans $(1, X]$ km and the outer ring spans $(X+1, 30]$ km. As in equation (3), the estimand is

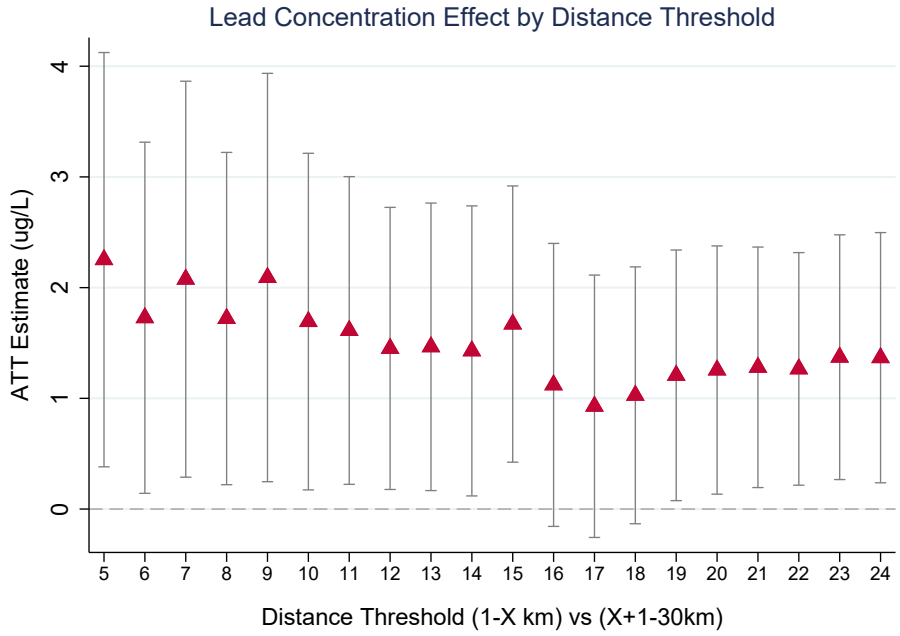
$$\text{ATT}(e) \;=\; \mathbb{E}\Big[\, \Delta Y_{f,t}^{\text{Inner}} - \Delta Y_{f,t}^{\text{Outer}} \;\Big|\; t - \hat{T}_f = e \,\Big], \tag{8}$$

where Y is now median lead concentration (μg/L) rather than log bird abundance, $\hat{T}_f$ is the satellite-inferred construction year, and ΔY denotes the change relative to the pre-construction period. I aggregate across event times $e \geq 0$ and cluster standard errors at the facility level.

IA Figure IA.2 reports the results. Lead concentrations increase by 1.7 to 2.3 μg/L within 5 to 10 km of a new facility. The effect attenuates with distance, dropping to roughly 1 μg/L at 16 to 18 km and stabilizing near 1.2 to 1.4 μg/L for thresholds above 19 km. This spatial decay is the expected signature of a localized pollution source. Expanding the treated radius dilutes exposure by including less affected areas, while the control ring thins as it retreats toward 30 km. Confidence intervals are wide at tight thresholds, reflecting sparse monitoring coverage, but the point estimates are positive across all 20 distance cutoffs.

Figure IA.2: Lead concentration by distance from facility

Staggered difference-in-differences estimates of the effect of TRI facility construction on ambient water lead concentrations (μg/L), varying the treatment radius from 5 to 24 km. At each threshold X, the treated ring spans $(1, X]$ km and the control ring spans $(X+1, 30]$ km. The estimator is Callaway and Sant'Anna (2021) with doubly-robust inverse probability tilting and not-yet-treated controls. Lead concentrations increase by approximately 1.7 to 2.3 μg/L within 5 to 10 km of a new facility and attenuate with distance. Error bars show 95% confidence intervals clustered at the facility level.



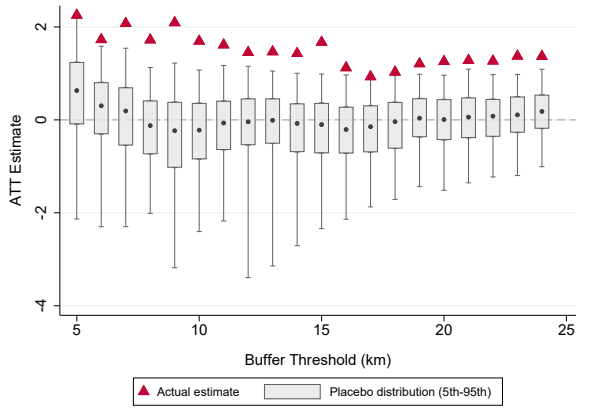
Lead placebo tests. I apply the same randomization inference used for the biodiversity placebos, keeping the spatial design fixed and permuting only facility construction years. Each of the 1,000 iterations randomly reassigns construction timing across facilities and re-estimates the staggered DiD under the permuted cohort structure. This generates an empirical null distribution for ATT_X that preserves cross-sectional heterogeneity in monitoring, distance, and facility characteristics but eliminates any systematic alignment between true construction and subsequent lead dynamics.

IA Figure IA.3 shows that the actual lead estimates sit above the placebo envelope across most thresholds, often exceeding the 95th percentile of the shuffled-timing distribution. In the z-statistic panel, the placebo distribution centers near zero and rarely approaches conventional significance, while the true estimates cluster around $z > 2$ for most distance cutoffs. If the baseline results were driven by generic time trends or estimator artifacts, the permuted cohorts would frequently produce similarly large positive effects. They do not. The lead results therefore provide a chemical pathway linking TRI facility operations to the ecological damage that the biodiversity pressure index captures.

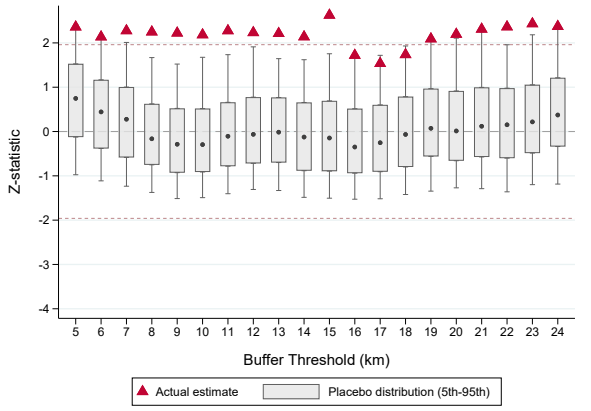
Figure IA.3: Randomization inference for lead concentration effects

Placebo distributions from 1,000 permuted-timing iterations of the lead concentration difference-in-differences estimator, varying the treatment radius from 5 to 24 km. Gray boxes show the interquartile range (25th–75th percentile) of placebo estimates; whiskers extend to the 5th and 95th percentiles; dots mark the placebo median. Red triangles show the actual estimates. Panel (a) plots ATT point estimates in μg/L. Panel (b) plots z-statistics, with dashed horizontal lines at ± 1.96. Actual lead concentration effects lie above the placebo distribution at close range and remain elevated across most thresholds.

(a) ATT estimates (μg/L)



(b) z-statistics



Section B: Robustness of the Biodiversity Premium

IA.2 Releases × BPI cross-sort

Table IA.1 complements the Fama-MacBeth horse race in Table A.2 with a non-parametric cross-sort. Firms are first sorted into terciles of toxic releases intensity, then into BPI terciles within each releases group. The BPI spread is positive within all three releases groups under FF3 controls. Under three-factor benchmarks, the spread is largest among low-releases firms (0.32%), but this ranking varies across augmented factor models: under FF5+Mom and FF5+Liq, the low-releases spread is the smallest of the three groups. The two measures capture complementary facets of environmental risk, and BPI adds pricing information that emissions intensity alone does not convey regardless of the releases group or factor model used.

Table IA.1: Releases × BPI double sort: multi-factor alphas

This table reports alphas of high-minus-low BPI spread portfolios from a conditional double sort on toxic releases terciles and BPI terciles. Each vintage year, firms are first sorted into terciles of ln(TRI releases/total assets), then into BPI terciles within each releases group. Value-weighted portfolios are formed within each of the nine releases–BPI cells. Each row reports the alpha of the high-BPI minus low-BPI spread portfolio within the indicated releases group, in percent per month. Column (1) reports CAPM alphas; column (2) controls for the three Fama-French (1993) factors; column (3) uses the five Fama-French (2015) factors; column (4) adds Betting Against Beta (Frazzini and Pedersen, 2014); column (5) adds the momentum factor (Carhart, 1997); and column (6) adds traded liquidity (Pástor and Stambaugh, 2003). Pooled averages the high and low portfolios across releases groups before computing the spread. Newey-West(12) standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1) CAPM	(2) FF3	(3) FF5	(4) FF5+BAB	(5) FF5+Mom	(6) FF5+Liq
H−L BPI α within releases tercile (%/month):						
Low Releases	+0.343	+0.322**	+0.125	+0.157	+0.095	+0.098
	(0.209)	(0.160)	(0.156)	(0.155)	(0.154)	(0.159)
Mid Releases	+0.168	+0.151	+0.117	+0.239	+0.206	+0.124
	(0.166)	(0.151)	(0.162)	(0.175)	(0.155)	(0.165)
High Releases	+0.142	+0.152	+0.142	+0.173	+0.182	+0.159
	(0.208)	(0.187)	(0.183)	(0.191)	(0.173)	(0.183)
Pooled	+0.218	+0.208*	+0.128	+0.190	+0.161	+0.127
	(0.155)	(0.118)	(0.116)	(0.122)	(0.109)	(0.118)
N months	434	434	434	432	434	434

IA.3 Granular size × BPI double sort (10 deciles × 5 quintiles)

Tables IA.2 and IA.3 report high-minus-low BPI spread alphas from a 10×5 conditional double sort. Each vintage year, firms are assigned to ten size deciles, then to five BPI quintiles within each decile. Table IA.2 uses unconditional size breakpoints computed across all sample firms. Under FF3 controls, the spread is positive and significant in eight of the ten deciles (D1–D8). Table IA.3 uses NYSE market-equity breakpoints following the standard Fama and French (1993) methodology, which concentrates more firms in the smallest deciles. Results are similar or stronger. The premium is robust through D7 across all six factor models in both tables. At D6, the spread falls below conventional significance under FF5+Mom but holds under the other five models. D8 is significant under the unconditional breakpoints except under FF5+Mom, and loses significance under NYSE breakpoints in four of the six models (FF5, FF5+BAB, FF5+Mom, and FF5+Liq). The spread is indistinguishable from zero in the two largest deciles across all specifications. The premium is therefore not confined to micro-caps and is economically large and statistically reliable across the broad cross-section of small and mid-cap firms.

Table IA.2: Size × BPI double sort: unconditional breakpoints (10 × 5)

Each vintage year, firms are sorted into ten size deciles by market equity, then into five BPI quintiles within each decile. Size breakpoints are unconditional, computed across all sample firms. The table reports the high-minus-low BPI spread (Q5−Q1) alpha within each size decile, in percent per month. D1 is the smallest-cap decile, D10 the largest. Column (1) reports CAPM alphas; column (2) controls for the three Fama-French (1993) factors; column (3) uses the five Fama-French (2015) factors; column (4) adds Betting Against Beta (Frazzini and Pedersen, 2014); column (5) adds the momentum factor (Carhart, 1997); and column (6) adds traded liquidity (Pástor and Stambaugh, 2003). Value-weighted portfolios. Newey-West(12) t-statistics in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1) CAPM	(2) FF3	(3) FF5	(4) FF5+BAB	(5) FF5+Mom	(6) FF5+Liq
H−L BPI spread α (%/month)						
D1 (Small)	+1.850***	+1.927***	+1.820***	+1.632***	+1.578***	+1.801***
	(4.76)	(5.59)	(5.09)	(4.25)	(4.20)	(4.99)
D2	+1.312***	+1.310***	+0.980***	+0.865***	+0.828***	+1.020***
	(4.52)	(4.47)	(3.70)	(2.89)	(3.00)	(3.80)
D3	+1.060***	+1.113***	+1.080***	+1.149***	+0.936***	+1.094***
	(3.99)	(4.36)	(3.99)	(3.93)	(3.29)	(3.99)
D4	+0.761***	+0.827***	+0.956***	+1.002***	+0.956***	+0.976***
	(3.77)	(4.15)	(4.47)	(4.64)	(4.39)	(4.46)
D5	+0.732***	+0.734***	+0.756***	+0.767***	+0.599***	+0.775***
	(3.54)	(3.48)	(3.17)	(2.96)	(2.83)	(3.17)
D6	+0.457**	+0.451**	+0.394*	+0.427**	+0.324	+0.408**
	(2.53)	(2.43)	(1.90)	(1.98)	(1.53)	(2.02)
D7	+0.579***	+0.591***	+0.504**	+0.507**	+0.411**	+0.499**
	(2.61)	(2.74)	(2.45)	(2.40)	(2.13)	(2.39)
D8	+0.603***	+0.592***	+0.388*	+0.441*	+0.242	+0.410**
	(3.15)	(3.16)	(1.96)	(1.93)	(1.11)	(2.09)
D9	+0.256	+0.253	+0.140	+0.212	+0.084	+0.112
	(1.47)	(1.47)	(0.93)	(1.44)	(0.53)	(0.71)
D10 (Big)	+0.227	+0.178	−0.075	−0.052	−0.106	−0.103
	(1.10)	(1.03)	(−0.42)	(−0.29)	(−0.60)	(−0.58)
N months	434	434	434	434	434	434

Table IA.3: Size × BPI double sort: NYSE breakpoints (10 × 5)

Each vintage year, firms are sorted into ten size deciles by market equity, then into five BPI quintiles within each decile. Size breakpoints use NYSE market-equity cutoffs following the Fama and French (1993) methodology, which concentrates more firms in the smallest deciles. The table reports the high-minus-low BPI spread (Q5−Q1) alpha within each size decile, in percent per month. D1 is the smallest-cap decile, D10 the largest. Column (1) reports CAPM alphas; column (2) controls for the three Fama-French (1993) factors; column (3) uses the five Fama-French (2015) factors; column (4) adds Betting Against Beta (Frazzini and Pedersen, 2014); column (5) adds the momentum factor (Carhart, 1997); and column (6) adds traded liquidity (Pástor and Stambaugh, 2003). Value-weighted portfolios. Newey-West(12) t-statistics in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1) CAPM	(2) FF3	(3) FF5	(4) FF5+BAB	(5) FF5+Mom	(6) FF5+Liq
H−L BPI spread α (%/month)						
D1 (Small)	+2.212***	+2.333***	+2.356***	+2.179***	+2.223***	+2.359***
	(5.85)	(6.26)	(5.70)	(5.03)	(5.10)	(5.83)
D2	+1.038***	+1.065***	+0.897***	+0.894***	+0.782***	+0.930***
	(3.96)	(4.02)	(3.04)	(2.96)	(2.78)	(3.14)
D3	+0.668***	+0.684***	+0.626**	+0.697**	+0.569**	+0.616**
	(2.61)	(2.64)	(2.39)	(2.57)	(2.01)	(2.39)
D4	+0.908***	+0.961***	+1.015***	+1.003***	+0.884***	+1.032***
	(4.22)	(4.68)	(4.43)	(4.14)	(4.08)	(4.48)
D5	+0.666***	+0.644***	+0.721***	+0.789***	+0.657***	+0.749***
	(3.15)	(2.96)	(2.98)	(3.25)	(2.72)	(3.05)
D6	+0.538**	+0.544**	+0.444*	+0.443*	+0.309	+0.470**
	(2.53)	(2.52)	(1.95)	(1.91)	(1.40)	(2.12)
D7	+0.592***	+0.596***	+0.556**	+0.555**	+0.485**	+0.574**
	(2.64)	(2.61)	(2.43)	(2.32)	(2.08)	(2.52)
D8	+0.415**	+0.424**	+0.267	+0.268	+0.037	+0.275
	(2.36)	(2.49)	(1.55)	(1.34)	(0.19)	(1.58)
D9	+0.154	+0.147	+0.022	+0.100	−0.009	−0.016
	(0.88)	(0.86)	(0.15)	(0.64)	(−0.06)	(−0.11)
D10 (Big)	+0.174	+0.113	−0.155	−0.145	−0.176	−0.177
	(0.88)	(0.70)	(−0.92)	(−0.87)	(−1.06)	(−1.05)
N months	434	434	434	434	434	434

IA.4 Industry × size × BPI triple sort

Table IA.4 controls for industry composition within size bins using a triple sort on Fama-French 12 industries, market equity terciles within industry, and BPI terciles within each industry-size cell. The small-firm BPI spread remains economically and statistically significant under all factor models, while the big-firm spread is statistically indistinguishable from zero. Industry sorting does not eliminate the size gradient documented in Table 4.

Table IA.4: Industry × size × BPI triple sort: multi-factor alphas

This table reports Fama-French three-factor (column 1) and five-factor (columns 2–5) alphas of high-minus-low BPI spread portfolios from a triple sort on Fama-French 12 industry, market equity terciles within industry, and BPI terciles within each industry–size cell. Each vintage year, firms are first allocated to FF12 industries and size terciles within industry, then sorted into BPI terciles within each of the 36 industry–size bins. Value-weighted portfolios are formed by pooling across industries within each size–BPI cell. Each cell reports the alpha of the high-BPI minus low-BPI spread portfolio, in percent per month. Column (1) controls for the three Fama-French (1993) factors; column (2) uses the five Fama-French (2015) factors; column (3) adds Betting Against Beta (Frazzini and Pedersen, 2014); column (4) adds the momentum factor (Carhart, 1997); and column (5) adds traded liquidity (Pástor and Stambaugh, 2003). Pooled averages the high and low portfolios across size groups before computing the spread. Newey-West(12) t-statistics in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1)	(2)	(3)	(4)	(5)
	FF3	FF5	FF5+BAB	FF5+Mom	FF5+Liq
H−L BPI α within industry–size cell (%/month):					
Small	+0.814***	+0.863***	+0.905***	+0.825***	+0.885***
	(5.28)	(5.61)	(5.72)	(4.92)	(5.67)
Mid	+0.479***	+0.430***	+0.454***	+0.348***	+0.404***
	(4.27)	(3.52)	(3.93)	(3.05)	(3.59)
Big	+0.096	−0.077	−0.096	−0.081	−0.107
	(0.82)	(−0.69)	(−0.89)	(−0.75)	(−0.95)
Pooled	+0.463***	+0.405***	+0.421***	+0.364***	+0.394***
	(6.29)	(5.84)	(6.04)	(4.98)	(5.93)
N months	434	434	432	434	434

IA.5 Asset characteristics and governance controls

Table IA.5 tests whether BPI proxies for the age, tangibility, or governance quality of a firm's asset base. Asset Age, defined as one minus the ratio of net to gross PP&E, captures the depreciation state of the capital stock. If BPI merely identified old, heavily depreciated facilities, Asset Age should absorb the coefficient. It does not. In the univariate specification (column 2), Asset Age is marginally significant (-0.070, $t = -1.84$) but BPI strengthens. Tangibility and institutional ownership concentration are individually insignificant (columns 3–4) and leave BPI intact. In the saturated specification with all three controls and Releases (column 6), BPI retains a coefficient of 0.182% per month, significant at the 1% level. The premium is not an artifact of old, tangible, or poorly governed asset bases.

Table IA.5: Fama-MacBeth regressions: asset characteristics and governance controls

This table reports monthly Fama-MacBeth regressions of excess stock returns on standardized BPI and progressively richer controls. Asset Age is one minus the ratio of net to gross PP&E (higher = more depreciated capital stock). Tangibility is PP&E divided by total assets. IO Concentration is the Herfindahl index of institutional ownership from 13-F filings. All independent variables except factor controls are standardized to zero mean and unit variance within each cross-section. Base controls include log book-to-market, momentum, log market equity, ROA, leverage, idiosyncratic volatility, and Fama-French 12-industry fixed effects. t-statistics use Newey-West standard errors with six lags. The sample covers October 1988 to December 2024. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)
BPI (standardized)	0.217***	0.248***	0.236***	0.207***	0.204***	0.182***
	(5.15)	(5.15)	(4.81)	(4.07)	(3.82)	(3.25)
Asset age		−0.070*			−0.058	−0.061
		(−1.84)			(−1.43)	(−1.49)
Tangibility			0.045		0.004	−0.007
			(1.01)		(0.10)	(−0.15)
IO concentration				0.025	0.035	0.033
				(0.50)	(0.63)	(0.58)
Releases (standardized)						0.047
						(1.13)
Base controls	✓	✓	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓	✓	✓
N months	428	324	324	324	324	324
Avg. firms/month	454	293	357	338	277	277
Avg. R^2	0.14	0.18	0.16	0.17	0.20	0.20

IA.6 Earnings surprises and abatement status

Table IA.6 tests whether the positive relationship between BPI and earnings surprises documented in Table 9 differs by abatement status. If non-remediating firms inflated short-term earnings by deferring environmental costs, the positive BPI coefficient should concentrate among non-abaters and disappear for firms that actively invest in source reduction.

Columns 1–3 regress SUE1 (seasonal random-walk surprise) on standardized BPI, an abatement indicator, and their interaction. BPI is positive and significant at the 1% level across all three specifications. The BPI $\times$ Abatement interaction is statistically insignificant once industry and year fixed effects are included (column 2, coefficient = 0.017), and adding releases leaves the result unchanged (column 3). Both abaters and non-abaters show positive BPI-SUE relationships of similar magnitude, ruling out the possibility that non-remediating firms inflate short-term earnings by deferring environmental costs. The positive SUE reflects the persistent cost structure documented in Section 5 rather than strategic cost deferral by a subset of firms.

Table IA.6: Earnings surprises, BPI, and abatement status

This table regresses the seasonal random-walk standardized unexpected earnings (SUE1) on BPI, an abatement indicator, and their interaction. BPI and Releases are standardized within TRI vintage year. Abatement equals one if any of a firm's TRI facilities report source reduction activities in the P2 module. Controls include log market equity and log book-to-market. Standard errors clustered by firm in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	SUE1 (Random Walk)		
	(1)	(2)	(3)
BPI	0.0624***	0.0600***	0.0601***
	(0.0109)	(0.0120)	(0.0130)
Abatement	−0.0303**	−0.0094	−0.0103
	(0.0153)	(0.0176)	(0.0184)
BPI × Abatement	0.0249*	0.0165	0.0206
	(0.0146)	(0.0146)	(0.0152)
Releases			−0.0016
			(0.0088)
ln(ME)	0.0327***	0.0236***	0.0242***
	(0.0049)	(0.0060)	(0.0063)
ln(B/M)	−0.217***	−0.203***	−0.203***
	(0.0131)	(0.0129)	(0.0140)
Industry FE	×	✓	✓
Year FE	×	✓	✓
N	71,451	71,442	66,624
Firms	1,729	1,728	1,631

IA.7 Return premium by abatement status

Table IA.7 reports Fama-French three-factor alphas for the high-minus-low BPI spread estimated separately for non-abating and abating firms. The premium concentrates in non-abaters, where the FF3 alpha is 0.52% per month and significant. For firms that actively invest in source reduction, the spread is economically and statistically indistinguishable from zero. This pattern is consistent with the operating-cost evidence in Table 11, where abatement offsets roughly 43% of the BPI cost penalty, and firms that eliminate the persistent cost drag also eliminate the associated return premium.

Table IA.7: BPI Return Premium by Abatement Status

This table reports factor model alphas for the high-minus-low BPI spread, estimated separately for non-abating and abating firms. Abatement equals one if at least one of the firm's facilities filed a Pollution Prevention (P2) source reduction report with the EPA's Toxic Release Inventory in the concurrent TRI reporting year. Portfolios are formed by sorting into BPI terciles within each abatement group and TRI vintage year. Returns are value-weighted. t-statistics use Newey-West(12) standard errors. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	Non-Abaters	Abaters
FF3 α (%)	0.524^{***}	-0.079
	(3.03)	(-0.49)
FF5 α (%)	0.453^{**}	-0.213
	(2.47)	(-1.32)
FF5+BAB α (%)	0.465^{**}	-0.173
	(2.34)	(-0.98)
FF5+Mom α (%)	0.426^{**}	-0.111
	(2.34)	(-0.80)
FF5+Liq α (%)	0.449^{**}	-0.214
	(2.44)	(-1.28)
Months	434	372
Unique firms	1,776	1,141
Avg. firms/month	342	289

IA.8 Post-2002 eBird subsample

eBird launched in 2002; pre-2002 observations come from retroactively digitized historical checklists with sparser spatial coverage. To confirm that results do not depend on noisy early bird data, I re-estimate the core tests on the subsample where BPI is constructed from post-2002 eBird vintages only (October 2002 through December 2024, 267 months, 458 firms per month on average).

Panel A of Table IA.8 reports Fama-MacBeth regressions. BPI enters with a coefficient of 0.175% per month, statistically significant, in the univariate specification and remains significant in the horse-race and saturated specifications that control for toxic releases. Releases is insignificant in all specifications. All coefficients are comparable to or larger than those in the full-sample Table 5.

Panel B reports double-sort alphas. The small-firm FF3 alpha is 81 basis points per month, declining to 46 basis points in mid-caps and insignificant in large firms. The size gradient documented in Section 4 is, if anything, sharper in the post-2002 period, and the COGS/Sales regression yields a BPI coefficient stronger than the full-sample estimate.

Table IA.8: Post-2002 eBird subsample robustness

The sample is restricted to observations with tri_vintage $\geq$ 2002 (October 2002 through December 2024; 267 months, 458 firms per month on average). *Panel A*: Monthly Fama-MacBeth regressions of excess returns on lagged firm characteristics. BPI, idiosyncratic volatility, and releases are cross-sectionally standardized each month. All specifications include Fama-French 12-industry fixed effects. Newey-West(6) t-statistics in parentheses. Coefficients are in percent per month. *Panel B*: High-minus-low BPI spreads within size terciles (independent sorts on size and BPI within each vintage year; value-weighted). Alphas from Newey-West(12) regressions of the H-L spread on the Fama-French three factors. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Panel A: Fama-MacBeth regressions

	(1)	(2)	(3)	(4)
BPI (standardized)	$+0.175^{***}$	$+0.254^{***}$	$+0.170^{***}$	$+0.197^{***}$
	(2.86)	(3.78)	(2.71)	(3.47)
Releases (standardized)			+0.039	+0.031
			(0.85)	(0.75)
Log book-to-market		-0.816^{***}	-0.628^{***}	-0.634^{***}
		(−6.92)	(−5.58)	(−5.73)
Momentum ($t-12, t-2$)		−0.437	-0.804^{**}	-0.772^{**}
		(−1.20)	(−2.24)	(−2.29)
Log market equity		−0.030	-0.120^{**}	−0.003
		(−0.49)	(−2.26)	(−0.08)
ROA			$+0.031^{***}$	$+0.032^{***}$
			(7.83)	(8.94)
Leverage			+0.093	−0.174
			(0.21)	(−0.44)
Idiovol (standardized)				$+30.65^{***}$
				(3.41)
Investment rate				-0.943^{***}
				(−3.45)
Industry FE	✓	✓	✓	✓
N months	267	267	267	267
N observations	122,242	122,242	122,242	122,242

Panel B: Size × BPI double-sort alphas (FF3, Newey-West 12)

Size tercile	Alpha (bps/mo)	t-stat	N
Small	+81.0	5.73	267
Mid	+45.6	3.65	267
Large	+12.6	0.97	267
Pooled	+46.4	5.88	267

Section C: Mechanisms of Information Failure

IA.9 Media coverage and biodiversity exposure

Table IA.9 examines whether aggregate biodiversity attention affects the BPI premium. Panel A shows that the interaction of high-BPI quartile membership with the NYT biodiversity attention index significantly predicts RepRisk news counts, indicating that media scrutiny of biodiversity-exposed firms rises with public attention. Panel C tests whether the BPI return premium varies with the NYT attention index in Fama-MacBeth regressions; the interaction is insignificant, indicating that the premium does not depend on aggregate attention cycles. Panels D–E confirm this finding using time-series and subsample approaches. The premium exists in both high- and low-attention months.

Table IA.9: NYT biodiversity attention tests

Panel A: PPML regression of RepRisk news counts on HighBioQuart × NYT biodiversity index with firm and year fixed effects (annual). Panel B: Campbell-Shiller VAR(1) R-squared for the BPI spread with 6 state variables including NYT_{bio}. Panel C: Fama-MacBeth cross-sectional regressions of excess returns on BPI × NYT interactions (monthly). Panel D: Time-series regression of BPI long-short return on NYT_{bio} controlling for FF3 factors. Panel E: BPI Fama-MacBeth coefficient estimated separately in high versus low NYT attention months.

Test	Estimate	t-stat	p-value	N
Panel A: Annual (PPML, Firm+Year FE)				
HighBioQuart × NYT_{bio}	+0.209		0.003***	10,110
Panel B: Campbell-Shiller R^2 (with vs without NYT_{bio})				
Main portfolio	3.95%			267
Small firms	3.58%			267
Big firms	0.97%			267
Panel C: Monthly Fama-MacBeth				
BPI × NYT_{bio}	−0.016%	−0.91	0.366	288 mo
BPI × $\text{NYT}_{\text{climate}}$	−0.024%	−1.59	0.113	288 mo
Panel D: Time-Series Portfolio				
NYT_{bio} β (TS reg)	−0.056%	−0.31	0.755	267 mo
Panel E: Subsample Split				
Low attention months	+0.138%	2.16**	0.032	207 mo
High attention months	+0.213%	3.06***	0.003	81 mo

IA.10 Institutional ownership and analyst coverage

Tables IA.10 and IA.11 test whether institutional ownership and analyst coverage moderate the BPI premium. Both tests face the same limitation, as their data requirements (13-F filings for IO, IBES estimates for coverage) act as size filters that disproportionately exclude the micro-cap firms where the premium is strongest.

Table IA.10 sorts on institutional ownership. The BPI spread concentrates in the low-IO tercile (FF3 alpha 0.69%). Panel B sorts on residual IO orthogonalized to log(ME) following Nagel (2005). After removing the mechanical size correlation, the BPI spread is statistically indistinguishable from zero across all residual IO groups, confirming that IO provides no moderation beyond size.

Table IA.11 sorts on the number of IBES analyst estimates. The BPI spread is statistically indistinguishable from zero in all three coverage groups under FF3, and the pooled H-L spread is zero under all six factor models. Coverage and size are correlated ($\rho = 0.741$), and the IBES data requirement filters out the small firms where the premium is strongest, so the null result reflects the absence of the relevant size segment rather than a moderating effect of coverage.

Panel B orthogonalizes coverage to log(ME) following Nagel (2005). After removing the mechanical size component, the BPI spread shows no economically or statistically significant moderation across residual coverage groups (pooled FF3 $\alpha = 0.08\%$), confirming that residual coverage provides no moderation beyond size.

Table IA.12 complements the portfolio sorts with a regression-based test. I split the sample at the median number of IBES analysts (17) and estimate the BPI coefficient on earnings-announcement CARs separately for below-median and above-median coverage firms. BPI is statistically significant on both sides of the split: 0.173% for below-median and 0.106% for above-median coverage. In column (4), the interaction BPI $\times$ High Coverage is economically small and statistically insignificant (0.037), confirming that analyst coverage does not attenuate the BPI earnings-announcement premium.

Table IA.10: Institutional ownership × BPI double sort: multi-factor alphas

This table reports alphas of high-minus-low BPI spread portfolios from a conditional double sort on institutional ownership (IO) terciles and BPI terciles. Each vintage year, firms are first sorted into IO terciles, then into BPI terciles within each IO group. Value-weighted portfolios are formed within each of the nine IO–BPI cells. Each row reports the alpha of the high-BPI minus low-BPI spread portfolio within the indicated IO group, in percent per month. Panel A sorts on raw IO (total institutional shares / shares outstanding). Panel B sorts on residual IO orthogonalized to ln(ME) following Nagel (2005). IO is lagged by one quarter to avoid look-ahead bias from the 45-day 13-F filing window. IO is linked to CRSP via the WRDS FactSet-CRSP linking table using historical CUSIP matching with date-range validation. Column (1) reports CAPM alphas; column (2) controls for the three Fama-French (1993) factors; column (3) uses the five Fama-French (2015) factors; column (4) adds Betting Against Beta (Frazzini and Pedersen, 2014); column (5) adds the momentum factor (Carhart, 1997); and column (6) adds traded liquidity (Pástor and Stambaugh, 2003). Pooled averages the high and low portfolios across IO groups before computing the spread. Newey-West(12) standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1) CAPM	(2) FF3	(3) FF5	(4) FF5+BAB	(5) FF5+Mom	(6) FF5+Liq
Panel A: Raw Institutional Ownership						
H−L BPI α within IO tercile (%/month):						
Low IO	+0.783**	+0.694***	+0.518**	+0.503**	+0.541**	+0.563**
	(0.314)	(0.244)	(0.243)	(0.218)	(0.240)	(0.248)
Mid IO	+0.236	+0.135	+0.006	+0.047	+0.066	+0.014
	(0.226)	(0.161)	(0.151)	(0.152)	(0.134)	(0.153)
High IO	−0.108	−0.201	−0.225	−0.164	−0.174	−0.252
	(0.226)	(0.181)	(0.168)	(0.182)	(0.157)	(0.170)
Pooled	+0.289	+0.207	+0.101	+0.128	+0.148	+0.110
	(0.209)	(0.132)	(0.128)	(0.137)	(0.115)	(0.133)
N months	321	321	321	319	321	309
Panel B: Residual IO (Nagel 2005)						
H−L BPI α within IO tercile (%/month):						
Low Resid IO	+0.297	+0.203	−0.040	−0.020	−0.009	+0.003
	(0.270)	(0.205)	(0.192)	(0.185)	(0.190)	(0.189)
Mid Resid IO	+0.068	−0.005	−0.059	+0.007	+0.027	−0.061
	(0.216)	(0.181)	(0.178)	(0.202)	(0.162)	(0.186)
High Resid IO	+0.119	+0.082	+0.047	+0.069	+0.042	+0.017
	(0.141)	(0.143)	(0.158)	(0.160)	(0.156)	(0.157)
Pooled	+0.158	+0.097	−0.004	+0.031	+0.033	−0.002
	(0.168)	(0.127)	(0.129)	(0.141)	(0.124)	(0.135)
N months	315	315	315	313	315	303

Table IA.11: Analyst coverage × BPI double sort: multi-factor alphas

This table reports alphas of high-minus-low BPI spread portfolios from a conditional double sort on analyst coverage terciles and BPI terciles. Each vintage year, firms are first sorted into terciles of IBES analyst coverage (number of annual estimates), then into BPI terciles within each coverage group. Value-weighted portfolios are formed within each of the nine coverage–BPI cells. Each row reports the alpha of the high-BPI minus low-BPI spread portfolio, in percent per month. Panel A sorts on raw IBES analyst coverage. Panel B sorts on residual coverage orthogonalized to ln(ME) following Nagel (2005). Column (1) reports CAPM alphas; column (2) controls for the three Fama-French (1993) factors; column (3) uses the five Fama-French (2015) factors; column (4) adds Betting Against Beta (Frazzini and Pedersen, 2014); column (5) adds the momentum factor (Carhart, 1997); and column (6) adds traded liquidity (Pástor and Stambaugh, 2003). Pooled averages the high and low portfolios across coverage groups before computing the spread. Newey-West(6) standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1) CAPM	(2) FF3	(3) FF5	(4) FF5+BAB	(5) FF5+Mom	(6) FF5+Liq
Panel A: Raw Analyst Coverage						
H−L BPI α within coverage tercile (%/month):						
Low Coverage	0.185	0.156	0.162	0.231	0.153	0.181
	(0.165)	(0.159)	(0.162)	(0.172)	(0.176)	(0.164)
Mid Coverage	−0.049	−0.100	−0.074	0.016	−0.007	−0.082
	(0.147)	(0.141)	(0.153)	(0.168)	(0.155)	(0.158)
High Coverage	0.016	−0.095	−0.313**	−0.333**	−0.272*	−0.339**
	(0.182)	(0.140)	(0.140)	(0.149)	(0.139)	(0.140)
Pooled	0.055	−0.008	−0.070	−0.024	−0.037	−0.076
	(0.116)	(0.101)	(0.106)	(0.121)	(0.113)	(0.109)
N months	396	396	396	396	396	396
Panel B: Residual Coverage (Nagel 2005)						
H−L BPI α within coverage tercile (%/month):						
Low Resid Cov	0.315*	0.227	0.164	0.222	0.257*	0.136
	(0.173)	(0.142)	(0.155)	(0.171)	(0.152)	(0.158)
Mid Resid Cov	0.145	0.084	0.029	0.075	0.057	0.051
	(0.200)	(0.177)	(0.175)	(0.187)	(0.180)	(0.176)
High Resid Cov	0.061	−0.069	−0.358*	−0.344*	−0.311	−0.367*
	(0.273)	(0.209)	(0.206)	(0.208)	(0.195)	(0.205)
Pooled	0.175	0.082	−0.054	−0.015	0.003	−0.060
	(0.178)	(0.129)	(0.125)	(0.131)	(0.119)	(0.127)
N months	396	396	396	396	396	396

Table IA.12: Earnings announcement CARs by analyst coverage

	Full Sample	Below Median	Above Median	Interaction
BPI (std.)	0.124***	0.173***	0.106**	0.106*
	(3.26)	(2.77)	(2.32)	(1.90)
High Coverage				0.175**
				(2.26)
BPI × High Coverage				0.037
				(0.61)
ln(ME)	-0.023	0.024	-0.108***	-0.051**
	(-1.08)	(0.61)	(-3.46)	(-2.05)
ln(B/M)	-0.061	0.037	-0.141**	-0.066
	(-1.26)	(0.47)	(-2.57)	(-1.35)
N	74,161	35,493	38,668	74,161
Firms	1,806	1,653	1,028	1,806
Adj. R^2	0.002	0.002	0.002	0.002

Notes: Dependent variable is CAR[-1,+1] (market-adjusted, %). BPI is standardized within TRI vintage year. High Coverage equals one for firms with above-median analyst coverage (median = 17). Firms without IBES coverage are coded as zero analysts. All specifications include FF12 industry dummies (coefficients suppressed) and year fixed effects. Standard errors clustered by firm and year-month; t-statistics in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

IA.11 Board governance and network connectedness

Tables IA.10 and IA.11 show that institutional ownership and analyst coverage provide no moderation of the BPI premium beyond firm size. I next test whether governance quality at the board level helps firms recognize and price their ecological exposures. Directors with extensive professional networks might diffuse information to more market participants, correcting the forecast errors that generate the premium. More independent boards might demand better environmental disclosure, and larger boards might bring more diverse expertise.

Table IA.13 tests these hypotheses using BoardEx data matched to 98% of BPI firms from 2005 onward. I construct three board-level variables: average director network size (log-transformed), the fraction of non-executive directors, and total board size. Each variable is orthogonalized to log(ME) cross-sectionally each year to remove the mechanical size correlation ($\rho = 0.59$ for network size, -0.29 for independence, 0.59 for board size), then standardized monthly. Monthly Fama-MacBeth regressions include book-to-market and Fama-French 12-industry fixed effects.

The board variables themselves predict returns individually. More connected and larger boards earn lower returns, while more independent boards earn higher returns. Board characteristics carry real return effects, but none of the three BPI interactions is statistically significant. In a saturated specification with all three board variables and all three interactions, BPI remains significant and every interaction remains insignificant. Well-governed boards confer no advantage in pricing ecological exposures.

Table IA.13: Board governance and the BPI premium

This table reports monthly Fama-MacBeth regressions of excess returns on BPI, board governance variables, and their interactions. Board variables are from BoardEx, matched to 98% of BPI firms from 2005 onward via the WRDS linking table (preferred matches, score ≤ 6). Network is the log of average director network size. Independence is the fraction of non-executive directors. Board Size is the number of directors. All board variables are orthogonalized to log(ME) cross-sectionally each year and standardized monthly (^{orth}). Controls include log(ME) and book-to-market. All specifications include Fama-French 12-industry fixed effects. Newey-West(6) standard errors in parentheses. Coefficients in percentage points. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	(1) Baseline	(2) Network	(3) Independence	(4) Board Size	(5) All
BPI (std)	0.364***	0.333***	0.361***	0.356***	0.322***
	(0.066)	(0.062)	(0.064)	(0.066)	(0.060)
Network^{orth} (std)		−0.205***			−0.175***
		(0.065)			(0.066)
BPI × Network^{orth}		0.021			0.012
		(0.037)			(0.041)
Independence^{orth} (std)			0.108***		0.163***
			(0.034)		(0.040)
BPI × Independence^{orth}			0.013		−0.005
			(0.050)		(0.050)
Board Size^{orth} (std)				−0.233***	−0.212***
				(0.047)	(0.046)
BPI × Board Size^{orth}				−0.039	−0.035
				(0.038)	(0.038)
Controls	✓	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓	✓
Months	288	288	288	288	288
Avg. firms/month	468	468	468	468	468

IA.12 Cost adjustment after EPA enforcement

The return-based enforcement drift in Section 6 shows that regulatory penalties trigger gradual price declines for high-BPI firms. To test whether the price response reflects real cost escalation, I construct a binary indicator $\text{Post_Penalty}_{j,q}$ equal to one if firm j received an EPA enforcement action with a monetary penalty (combined_cost > 0) in the preceding four quarters. I interact this indicator with standardized BPI in the COGS/Sales panel regression from Equation 4, retaining firm and year fixed effects with firm-clustered standard errors.

The BPI main effect is stable at 0.42 percentage points, significant at the 1% level, confirming that high-BPI firms face structurally higher costs regardless of enforcement status. The Post_Penalty main effect is economically and statistically zero (-0.01), indicating that enforcement events do not raise costs for the average firm. The interaction BPI $\times$ Post_Penalty is -0.25, marginally significant at the 10% level, indicating that following enforcement the cost gap between high- and low-BPI firms narrows. High-BPI firms subject to a penalty move toward the cost level of their low-BPI peers, suggesting that penalized firms begin compliance-driven operational adjustments of the kind documented in Section 5. The adjustment is firm-specific: the null Post_Penalty main effect confirms that enforcement does not raise costs uniformly.

Table IA.14: COGS/Sales and EPA enforcement

This table reports quarterly panel regressions of COGS/Sales on standardized BPI, a post-penalty indicator, and their interaction. Post_Penalty equals one if the firm received an EPA enforcement action with a monetary penalty in the preceding four quarters. Controls include lagged COGS/Sales, log assets, log book-to-market, ROA, leverage, and tangibility. All specifications include firm and year fixed effects with standard errors clustered by firm. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	(1)	(2)	(3)
BPI	0.399***	0.399***	0.420***
	(0.089)	(0.089)	(0.091)
Post-Penalty		−0.012	−0.015
		(0.150)	(0.150)
BPI × Post-Penalty			−0.248*
			(0.140)
Lagged COGS/Sales	0.578***	0.578***	0.578***
	(0.018)	(0.018)	(0.018)
Log assets	−0.114	−0.114	−0.115
	(0.161)	(0.161)	(0.161)
Log BM	1.103***	1.103***	1.107***
	(0.105)	(0.105)	(0.105)
ROA	0.002	0.002	0.002
	(0.004)	(0.004)	(0.004)
Leverage	0.652	0.652	0.647
	(0.484)	(0.485)	(0.485)
Log PPE/Assets	0.188	0.189	0.193
	(0.247)	(0.247)	(0.247)
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
N	72,807	72,807	72,807

IA.13 10-K biodiversity disclosure

Table IA.15 tests whether BPI predicts corporate biodiversity disclosure in annual reports. The dependent variables are three 10-K biodiversity disclosure scores from Giglio et al. (2026): an aggregate indicator for any biodiversity sentence, a count of negative-sentiment sentences, and a regulation indicator. BPI is statistically indistinguishable from zero in all three specifications, confirming that ecological exposure does not predict voluntary disclosure.

Table IA.15: EPA enforcement and corporate disclosure avoidance

This table tests whether EPA enforcement suppresses corporate biodiversity disclosure. The dependent variables are three 10-K biodiversity disclosure scores from Giglio et al. (2026): *Aggregate* is a binary indicator for any biodiversity-related sentence (column 1). *Negative Sentiment* is a discrete count of sentences expressing negative sentiment or rising biodiversity risk (column 2). *Regulation* is a binary indicator for sentences discussing biodiversity-related regulatory risk (column 3). *Enforcement* $(t-1)$ is an indicator equal to one if the firm faced a material EPA enforcement action (combined penalty > \$10K) in the prior fiscal year. *BPI* (std) is ln(bird abundance/total assets), standardized cross-sectionally. *Releases* (std) is ln(TRI releases/total assets), standardized cross-sectionally. Controls are lagged one year: ln(Assets) is log total assets; ln(B/M) is log book-to-market equity; *ROA* is return on assets; ln(PPE/Assets) is log property, plant, and equipment scaled by assets; and *Leverage* is total debt over total assets. All regressions include firm and year fixed effects. t-statistics are in parentheses, computed from standard errors clustered by firm. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1) Aggregate	(2) Negative Sentiment	(3) Regulation
Enforcement $(t-1)$	−0.017	+0.012	−0.002
	(−1.50)	(1.38)	(−0.23)
BPI (std)	−0.008	−0.010	−0.007
	(−1.29)	(−1.31)	(−1.42)
Releases (std)	+0.007	+0.007	+0.005
	(1.51)	(1.42)	(1.19)
ln(Assets)	+0.008	−0.016	+0.001
	(0.78)	(−1.05)	(0.13)
ln(B/M)	−0.002	+0.003	−0.000
	(−0.40)	(0.54)	(−0.08)
ROA	+0.000	+0.000	+0.000
	(0.40)	(0.20)	(0.17)
ln(PPE/Assets)	+0.014	+0.025	+0.005
	(1.17)	(1.41)	(0.46)
Leverage	−0.046	−0.025	−0.040*
	(−1.63)	(−0.90)	(−1.67)
Firm FE	✓	✓	✓
Year FE	✓	✓	✓
R^2	0.530	0.444	0.577
N	7,770	7,770	7,770
Firms	648	648	648

Section D: Ruling Out Limits to Arbitrage

IA.14 Size, liquidity, and idiosyncratic volatility interactions

Table IA.16 asks whether any standard limits-to-arbitrage proxy absorbs the BPI premium in a single Fama-MacBeth specification. The baseline BPI coefficient is economically and statistically significant. Adding interactions with log market equity and idiosyncratic volatility leaves both insignificant. The BPI level coefficient rises when interactions are included, consistent with the mechanical recentering that occurs when interaction terms are added to a specification with standardized variables. The positive but insignificant BPI $\times$ Log ME interaction (+0.042) may appear to conflict with the monotonic size gradient in the portfolio sorts (Tables 4 and IA.2), but two features of the regression setting explain the difference: the equal-weighted Fama-MacBeth estimator gives different weight to size bins than value-weighted portfolio returns, and the strong negative correlation between log ME and idiosyncratic volatility ($\rho \approx -0.60$) makes the individual interaction coefficients unstable when both enter simultaneously. BPI retains its predictive power in the presence of releases and all interaction terms, confirming that the ecological dimension contains independent pricing information. The table includes the full set of standard controls (book-to-market, momentum, ROA, leverage) used in Table 5, and all specifications include Fama-French 12-industry fixed effects.

Table IA.16: Limits to arbitrage: size and idiosyncratic volatility interactions

Monthly Fama-MacBeth regression of excess returns on BPI, toxic releases, and their interactions with size (log market equity) and idiosyncratic volatility, all included simultaneously. All variables of interest are cross-sectionally standardized each month. Interaction terms are products of standardized constituents. Controls (log book-to-market, cumulative returns $t-12$ to $t-2$, ROA, leverage) and Fama-French 12-industry fixed effects are included but not reported. Newey-West(6) t-statistics in parentheses. Coefficients in percent per month. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	BPI (std)	Releases (std)
Level	$+0.285^{***}$	-0.013
	(4.75)	(-0.36)
$\times$ Log ME	$+0.042$	$+0.013$
	(1.13)	(0.34)
$\times$ Idiosyncratic vol.	$+0.062$	-0.010
	(0.99)	(-0.14)
Log ME	$+0.238^{**}$ (2.48)	
Idiosyncratic vol.	$+0.143$ (1.25)	
Controls	✓	
Industry FE	✓	
N months	428	

IA.15 Short-selling constraints

The short interest ratio (SIR) measures the fraction of shares outstanding that are currently sold short. A high SIR indicates that many investors have borrowed and sold shares, often reflecting bets against the stock or difficulties in further shorting (limited available inventory). If the BPI premium reflected limits-to-arbitrage from short-selling frictions rather than an information gap, it should be larger among high-SIR stocks where short positions are already stretched.

Tables IA.17 and IA.18 find no evidence for this. The Fama-MacBeth BPI × SIR interaction is statistically insignificant. SIR is nearly orthogonal to size ($\rho = -0.033$), so this null is not mechanically produced by the size correlation. The triple sort confirms that the BPI spread persists at similar magnitudes across all SIR terciles, ruling out shorting frictions as an explanation.

Table IA.17: Fama-MacBeth: BPI × short interest interaction

Monthly Fama-MacBeth regressions of excess returns on lagged firm characteristics. SIR = short-intadj / (shrout × 1000). All variables of interest are cross-sectionally standardized each month. Interaction terms are products of standardized constituents. Controls include log book-to-market, cumulative returns ($t-12$ to $t-2$), and Fama-French 12-industry fixed effects. Newey-West(6) t-statistics in parentheses. Coefficients are in percent per month. corr(SIR, Size) = −0.033. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1) Base	(2) +Interact	(3) +Size
BPI	+0.161***	+0.159***	+0.180***
	(3.44)	(3.53)	(3.54)
SIR	+0.027	+0.026	+0.037
	(0.57)	(0.59)	(0.87)
BPI × SIR		−0.039	−0.040
		(−1.03)	(−1.24)
Log ME			+0.028
			(0.25)
BPI × Size			−0.035
			(−0.99)
Controls	✓	✓	✓
Industry FE	✓	✓	✓
N months	427	427	427

Table IA.18: Triple sort: size × short interest × BPI

Conditional triple sort ($3 \times 3 \times 3$). Within each vintage year, firms are sorted into size terciles, then SIR terciles within each size group, then BPI terciles within each size×SIR cell. SIR = shortintadj / (shrout × 1000), lagged one month. Value-weighted portfolios. The table reports the Fama-French three-factor alpha of the high-minus-low BPI spread within each size×SIR cell, in percent per month. Pooled averages the H−L series across size groups. DiD is the difference in pooled H−L between High and Low SIR terciles. Newey-West(12) t-statistics in parentheses. Sample: 175,841 firm-months (1993–2022), 1,402 firms. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	Low SIR	Mid SIR	High SIR
FF3 α of H−L BPI spread within size × SIR cell (%/month):			
Small	+0.559*	+0.574**	+0.717***
	(1.82)	(2.35)	(3.01)
Mid	+0.031	+0.350*	+0.241
	(0.18)	(1.88)	(1.20)
Big	+0.185	−0.159	+0.041
	(1.13)	(−0.79)	(0.20)
Pooled	+0.259*	+0.255*	+0.333***
	(1.88)	(1.72)	(2.65)
N months	360	360	360