A bird's-eye view: firm biodiversity footprints and earnings expectations

Atreya Dey

University of Cambridge, Judge Business School

Abstract

Using bird-watching data matched to U.S. industrial facilities, I show that firms causally reduce bird populations and species. I use these estimates to construct a firm-level biodiversity footprint index that identifies firms with many ecologically consequential facilities relative to their size. Sorting on the footprint generates significant abnormal returns among small- and mid-cap firms from 1988 to 2024. The return premium reflects forecast errors: investors misinterpret the higher production costs of high-footprint firms as operational decline, and their pessimism corrects when earnings are announced. Although firms' facility locations are public, firms with larger biodiversity footprints provide less voluntary biodiversity information.

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 Igor Lončarski (discussant), Lasse Heje Pedersen, David Hirshleifer, Elroy Dimson, Alexandre Garel, Sarah Kroechert, Raghu Rau, Diego Guisande (discussant), Luca Taschini, Glen Gostlow, and Tim Newbold. I also thank seminar participants at the University of Cambridge, Judge Business School Brown Bag Seminar; the 2026 International Symposium on Climate, Finance, and Sustainability; the 3rd Workshop on Recent Trends and New Developments in Sustainable, Green & International Finance at EDHEC; the 2nd Financial Fraud, Misconduct and Market Manipulation Conference at Lancaster University; the Global Research Alliance for Sustainable Finance and Investment 2025; and the Sustainable and Impact Investments International Conference 2025. I am grateful for funding from the Cambridge Endowment for Research in Finance.

1 Introduction

Firm operations are a major driver of biodiversity loss, with an estimated $5 trillion in annual private finance supporting activities that directly harm nature.[1] In response, biodiversity loss has been framed as the next environmental externality, after carbon emissions, for green finance to address.[2] Investors, governments, and standard setters increasingly seek firm-level information on nature impacts that is comparable across firms and relevant for financial decision-making (Dimson et al., 2026). However, green finance can only discipline these externalities to the extent that measures of firms' externality-to-value ratio are accurate and comparable (Pedersen, 2026).

A central obstacle is that biodiversity is difficult to summarize as a firm-level characteristic, creating two measurement problems. First, firms' effects on biodiversity are context-specific— they arise from land-use change, pollution, and resource extraction whose consequences depend on both location and the spatial extent of ecological impact (Giglio et al., 2024). Greenhouse gas emissions, in contrast, can be standardized into CO_2 equivalents, aggregated across operations, scaled by firm value, and linked directly to climate damages. Second, biodiversity is difficult to define as it depends on how many organisms are present and how many species they represent, making cross-firm comparisons difficult. In the absence of a direct measure of ecological impact, investors rely on disclosures and environmental, social, and governance (ESG) scores to compare firms. Yet only 2% of firms discuss biodiversity in their 10-K filings (Giglio et al., 2026).

I propose a standardized measure of firms' ecological impacts by constructing a biodiversity footprint index using bird observations, industrial facility locations from the Toxics Release Inventory (TRI) program, and U.S. equity-market data.[3] I begin by providing causal

[1] United Nations Environment Programme, *State of Finance for Nature*. The report traces private financial flows that fund economic activity driving biodiversity loss, including habitat conversion, water and soil pollution, and unsustainable extraction of natural resources.

[2] I use the term green finance broadly to describe investment policies and regulatory frameworks designed specifically to address environmental challenges (Pedersen, 2026). See, "Environmental Accounting Must Go Beyond Climate to Include Nature," *The Wall Street Journal*, June 30, 2021.

[3] Birds are well suited to ecological monitoring because they are widespread, readily detected, extensively

evidence that industrial facilities reduce nearby bird populations: satellite-implied facility construction dates, used in a staggered difference-in-differences design, shows that bird abundance falls 6–9% and species richness 3–5% around new industrial facilities within a 14 km radius. The effect attenuates with distance from the facility and is stronger in years with higher toxic releases—two patterns that confirm the relevance of bird abundance as a measure of real ecological responses to industrial activity. Combined with market data, the resulting measure proxies for the biodiversity surrounding each firm's operations per dollar of assets, and varies continuously across firms and over time. I then show that larger ecologically consequential footprints predict value-relevant firm outcomes that are difficult to incorporate into expectations (Edmans, 2011; Hirshleifer et al., 2013). At the same time, these footprints are not well captured by existing information channels such as firm disclosures.

Specifically, I document four sets of results. First, I find a biodiversity footprint premium concentrated in small and mid-cap firms. Conditional double sorts on size and the footprint index generate a factor-adjusted high-minus-low return spread of 89 basis points per month in small firms and 52 basis points in mid-caps, whereas large firms show no significant spread. The premium is concentrated in the footprint's extensive margin, measured as the number of ecologically consequential facilities relative to assets, and is distinct from pollution intensity after controlling for toxic releases in a joint regression.[4] In Fama-MacBeth regressions, the premium survives controls for analyst coverage, liquidity, idiosyncratic volatility, and short interest. Controls for facility-neighborhood demographics leave the premium unchanged, ruling out demographic and urbanization differences around firm facilities as explanations. Nor does the premium covary with the biodiversity attention index of Giglio et al. (2026), meaning the return premium is inconsistent with an attention-driven channel.

Second, high-footprint firms have higher production costs because they operate more

observed, and sensitive to environmental changes (Fraixedas et al., 2020).

[4]Although both are built around the same EPA TRI facilities, the footprint uses TRI to identify firms' facility networks and measures surrounding biodiversity from eBird, so the two capture different dimensions of environmental risk. Releases measures a standardized quantity that investors can compare (Hsu et al., 2023); the footprint measures the ecological consequences of the firm's operating network.

ecologically consequential sites per dollar of assets. A one-standard-deviation increase in the biodiversity footprint index raises the cost of goods sold (COGS) by 0.35 percentage points of sales, with no effect on selling and administrative expenses. The cost penalty is stable year over year, consistent with a persistent feature of the firm's operations rather than a transitory shock. The cost is related to firms' production inputs, ruling out a demand-side channel in which consumers cut spending in response to environmental news (Derrien et al., 2025). The penalty partly reflects labor-intensive production technologies, where firms in more biodiverse locations employ more workers per dollar of fixed capital.[5] This interpretation differs from the regulatory-enforcement channels documented by Chen (2025) and Akbari et al. (2026), as EPA inspections, informal actions, and penalties are unrelated to a facility's ecological richness once its emissions are held fixed.

Third, prices adjust when earnings reveal information about these operational costs. Analysts are systematically pessimistic about high-footprint firms: a one-standard-deviation increase in the index is associated with a 4.6% lower 12-month implied return from analyst price targets. Yet these firms repeatedly beat pessimistic forecasts, generating positive earnings surprises against both seasonal random-walk and analyst-consensus benchmarks. Returns concentrate around the same information events, with four quarterly announcement windows accounting for 32% of the small-firm alpha, in line with the concentration documented for other anomalies (Engelberg et al., 2018). The correction begins at earnings announcements and continues in small firms over the following weeks (Hong and Stein, 1999).

Fourth, the premium persists because standard information channels do not translate firms' publicly disclosed facility networks into biodiversity operating footprints. Firms report their facility locations through the TRI, and I construct the footprint by pairing those facilities with ecological data. Equity analysts, however, do not incorporate this information, as greater coverage does not attenuate the premium. Firms also do not reveal the information

[5] Standard environmental metrics (state-level mini-NEPA exposure, industry-level water intensity, firm-level waste generation, and the share of releases subject to Comprehensive Environmental Response, Compensation, and Liability Act (CERCLA)) do not absorb the index's cost effect (Internet Appendix Table IA.13).

through their own ESG disclosures. Responses to the Carbon Disclosure Project's (CDP) voluntary biodiversity questionnaire show that larger footprints are associated with less reporting of biodiversity actions and less voluntary publication of biodiversity information.

The results support an errors-in-expectations hypothesis. Investors observe the persistent cost structure of high-footprint firms but cannot identify the underlying cause. They interpret the lower margins as operational decline and set earnings forecasts too low. The firms subsequently beat these forecasts, consistent with the market underprocessing a public but nonstandard source of firm value, as in Edmans (2011).

I rule out three rational alternatives. First, ambiguity aversion could generate a premium for hard-to-assess biodiversity-related uncertainty; however, it would not predict positive earnings surprises. Second, information-production models (Veldkamp, 2006) predict that the premium should shrink with more intermediary resources; instead, the effect persists across varying levels of analyst coverage. Third, a Merton (1987) investor-recognition premium compensates investors for holding firms with a limited investor base and does not predict biased forecasts. The positive earnings surprises of high-footprint firms therefore point to forecast errors rather than rational compensation for limited recognition.

I contribute to the green finance literature by showing that a firm's biodiversity footprint is financially material and by directly measuring—rather than imputing—firms' effects on biodiversity. Existing work studies biodiversity as a financial risk that affects expected returns and firms' cost of capital, typically using measures that infer a firm's ecological impact from its sector-and-country revenue mix (Garel et al., 2024), its industry's dependence on ecosystem services (Garel et al., 2026), or its news coverage and disclosures (Giglio et al., 2026). A subset of studies geolocate facilities but assume the firm's ecological effect: they fix a radius around each facility to count nearby protected areas (Akbari et al., 2026) or measure biodiversity at a set distance (Chen et al., 2024). My measure instead records bird abundance at each facility, confirms that this abundance is causally linked to industrial activity, and estimates rather than assumes the extent of a firm's ecological impact. The biodiversity footprint is an externality-to-value ratio in the sense of Pedersen (2026). The externality aggregates across facilities, and the cost penalty is concentrated in its extensive margin, measured by the number of ecologically consequential facilities per dollar of assets. This identifies a cash-flow channel rather than the discount-rate channel that much of the biodiversity-pricing literature emphasizes. The premium exists throughout the 1988–2024 sample rather than emerging only after the 2021 Kunming Declaration (Garel et al., 2024), and is distinct from the emissions-intensity premium of Hsu et al. (2023) and Bolton and Kacperczyk (2021).

I also contribute to the literature on how financial markets value public but nonstandard firm attributes. Previous studies document similar valuation failures in other settings. Edmans (2011) shows that markets undervalue employee satisfaction, even when it is independently verified and publicly visible, while Hirshleifer et al. (2013) show that innovative efficiency predicts returns when investors underprocess public patent information. In this paper, the underlying information is similarly public and independently sourced. Yet investors observe the higher production costs of high-footprint firms but fail to connect them to the firms' operating footprints. The premium therefore reflects forecast errors rather than compensation for recognized risk. Consistent with this interpretation, high-footprint firms generate positive earnings surprises, and returns concentrate around earnings announcements. Unlike the innovative-efficiency premium in Hirshleifer et al. (2013), greater analyst coverage does not attenuate the premium.

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. Chen et al. (2024) study critical habitat designations, per the Endangered Species Act, using TRI and eBird data. They find that affected establishments relocate pollution offsite and that affected firms experience lower valuations. I identify a cost channel that is independent of regulatory designation—firms have elevated production costs due to their distributed operations where their activity affects local biodiversity, whether or not that area receives formal protection. Without a discrete policy event to observe, the cost appears as elevated COGS, 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.

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 return premium and its relationship to emissions intensity. Section 5 traces the cash-flow expectations channel from operating costs through earnings surprises. Section 6 documents the information gap that sustains the premium. Section 7 concludes.

2 Data

The dataset combines eBird observations from industrial facility sites identified in the EPA Toxics Release Inventory (TRI) with stock returns and balance-sheet data from CRSP and Compustat. The merged panel contains 210,560 firm-month observations for 1,762 unique U.S. firms from October 1988 through December 2024. I describe the data in four parts: the construction of the Biodiversity Footprint Index (BFI), a parallel emissions-intensity measure (Releases), the financial panel and controls, and additional datasets used to test the mechanism, covering regulatory enforcement, voluntary disclosure, and news coverage. Table A.1 defines all variables and Table 1 reports summary statistics.

2.1 Measuring biodiversity footprints

Green finance disciplines a firm's environmental externality through its externality-to-value ratio, the total externality the firm imposes scaled by its value (Pedersen, 2026). For carbon, this ratio is emissions over value; for biodiversity, the difficulty is measuring the externality itself. A credible measure must track a real ecological quantity that the firm genuinely affects. The Biodiversity Footprint Index (BFI) measures the biodiversity surrounding a firm's facilities per dollar of assets. The surrounding area, undisturbed by the facility, gives a clean measure of the biodiversity a firm's locations potentially hold, while Section 3 estimates the biodiversity loss caused by facility operations.

Its construction begins with the precise geographic location of industrial facilities, which I obtain from the TRI. Established by the Emergency Planning and Community Right-to-Know Act of 1986, the TRI provides a comprehensive public registry of facility-level industrial activity in the U.S., containing annual reports on over 800 listed chemicals from approximately 21,000 facilities (U.S. Environmental Protection Agency, 2024).

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 (eBird, 2024). I use structured checklists from the eBird Basic Dataset, in which volunteer observers record which bird species have been detected, their counts, survey duration, distance traveled, and location, covering 1985 through 2024. Coverage grows from 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 BFI numerator is time-varying throughout the sample period, constructed from checklist data available in each year rather than backfilled from a single, recent abundance map. I treat facility-years with no nearby checklists as missing and exclude them from the firm-level sum.

Raw checklist totals reflect observer effort and timing as much as underlying bird populations, so I implement an 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 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. The state-by-month fixed effects absorb cross-state differences in biodiversity, so the adjusted abundance measure captures local deviations from each state's seasonal average rather than absolute ecological levels. Differencing against the state-season average also removes cross-state differences in eBird coverage and any state-level economic or regulatory conditions. For each facility, ring, and year, I average these effort-adjusted log counts across all checklists falling within the ring, and exponentiate to obtain a relative effort-adjusted abundance index. Appendix IA.1 provides the full specification and fixed-effect structure.

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 declines in bird populations and species within the treatment ring, relative to the control ring, after a facility is constructed.[6]

The outer ring plays two roles. In the difference-in-differences design in Section 3, it provides a counterfactual for the bird population that would have prevailed without the facility. The 15–30 km ring lies beyond the facility's primary impact area. In constructing BFI, outer-ring abundance measures the biodiversity exposed to the average causal effect of the average facility, whereas abundance near the facility measures the biodiversity remaining after that impact. With the average effect held constant across facilities, aggregating outer-ring abundance produces a firm-level measure proportional to the firm's total ecological externality. I use this aggregate as the index numerator and scale it by total assets to express the ecological externality per dollar of assets (Pedersen, 2026). I define BFI as

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

where Bird Abundance is the sum across all facilities operated by firm i in year t of the effort-adjusted abundance in each facility's 15–30 km 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

[6] Abundance counts the individual birds; species richness counts distinct species—two standard, complementary biodiversity measures.

facility-level measures additively. The index is thus higher for firms with more consequential facilities, biodiverse surroundings, or a smaller asset base.

Like emissions intensity, which sums a firm's total pollution, BFI sums the surrounding bird abundance across all of a firm's facilities, an aggregate rather than a per-site quantity.[7] BFI correlates -0.70 with log market capitalization because the index scales surrounding bird abundance by total assets. All return tests condition on size, and a denominator decomposition confirms that log(birds) and log(assets) each independently predict returns; the result also holds with unscaled bird counts and alternative denominators such as plant, property, and equipment (Appendix A.1).

I match this annual BFI measure to monthly returns using a single timing rule applied uniformly across all empirical tests in the paper. TRI data for reporting year Y are published in October of year $Y+1$. For example, 2020 TRI data become available in October 2021, so I assign that vintage to observations dated from November 2021 through October 2022. Zhang (2025) shows that the carbon premium stems from a look-ahead bias, since emissions are released with a long lag and, being estimated from sales, reflect contemporaneous performance. BFI avoids both problems. TRI's release date is official and known, so every test uses only already-public data, and the numerator is not constructed from firms' accounting outcomes. The eBird Basic Dataset is released within a year of collection, so a vintage's checklists are public before its November formation date. Pre-2002 records are retroactively digitized, and the post-2002 subsample (Internet Appendix Table IA.6) confirms the premium over the period when the bird data was collected in real time.

2.2 Emissions intensity

To distinguish ecological context from pollution output, and separate BFI from the pollution premium documented by Hsu et al. (2023), I construct a parallel measure of emissions

[7] When facilities are geographically clustered, their control rings overlap and nearby checklists contribute to multiple facility measures. A firm with more facilities in the same ecosystem is correspondingly more exposed to it, so the additive structure is appropriate.

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 (lbs) for firm i in year t. I deliberately mirror BFI's construction using the same TRI facilities, asset denominator, and vintage timing, which makes their moderate pairwise correlation of $r = 0.3$ notable. It shows that BFI captures the ecological context of a firm's location rather than its pollution output. While some heavily polluting facilities are located in rich ecosystems, many high-BFI firms emit little, and many high-Releases firms are located in ecologically sparse areas.

2.3 Financial data and controls

I merge the BFI and Releases measures onto 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 210,560 firm-month observations for 1,762 unique firms from October 1988 through December 2024. Firm coverage peaks at 642 firms around 2000 and declines to 468 firms by 2024, as TRI-reporting industries are consolidated. The 210,560 observations are the full BFI/CRSP/Compustat merge; requiring non-missing lagged controls yields the Fama-MacBeth estimation sample of 172,938 firm-months, spanning 428 calendar months with an average of 404 firms per cross-section. I classify firms using SIC codes according to the Fama-French 12 industries.

The baseline set of control variables includes log market equity, log book-to-market equity, the prior 12-month return (momentum), return on assets, leverage, and idiosyncratic volatility, the standard deviation of daily returns in excess of the market within each month.

2.4 Enforcement, disclosure, and intermediary data

I draw on additional datasets to probe the mechanisms through which firms realize these ecological costs, and by which 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.[8] These records detail case timing, linked facilities, and assessed penalties, allowing me to identify 4,161 enforcement events affecting 637 publicly traded firms in my sample.

Second, to examine the role of biodiversity news, I source 12,264 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. I also measure shifts in aggregate attention using the New York Times news index constructed by Giglio et al. (2026), which counts articles mentioning biodiversity-related terms.

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

2.5 Summary statistics

Table 1 reports summary statistics for the merged panel. BFI averages -3.59 with a panel standard deviation of 1.77, and Releases averages 3.11 with a standard deviation of 3.15. The standard control variables (log market equity, log book-to-market, momentum, return on assets, leverage, and idiosyncratic volatility) show typical cross-sectional dispersion. Two pairwise correlations shape the empirical strategy: BFI's strong negative correlation with log market equity ($\rho = -0.613$) means every return test conditions on size, and BFI's moderate

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

correlation with Releases ($\rho = 0.3$) leaves substantial independent variation that allows a distinction between the biodiversity premium and emissions intensity.

3 Ecological Validation

For the BFI to capture a value-relevant attribute, the ecology it measures must be real and it must genuinely be affected by industrial facilities within their locale. The spatial extent of that effect is itself an empirical object, yet prior work selects a radius by convention (Chen et al., 2024) rather than estimating it from the data. I use local bird populations as a proxy for ecological condition because they respond to habitat disturbance and pollution, and citizen-science surveys provide dense spatial coverage (Fraixedas et al., 2020). Simply comparing bird populations across space is not sufficient, because facilities locate in places that differ systematically from non-facility areas. I therefore use a staggered difference-in-differences design that varies the treatment radius from 5 to 24 kilometers and measure the effect of new TRI facility construction on local bird populations.

A central challenge is that a facility's construction date may precede its first regulatory filing by several years. I therefore extract 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—average impervious surface growth spikes at the detected construction year and is near zero in surrounding years (Internet Appendix Figure IA.1). Visual inspection of high-resolution aerial imagery confirms that these statistical jumps correspond to physical land clearing and infrastructure development (Figure 2).

To isolate a facility's causal effect 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, is 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 tests for spatial attenuation of the effect. Any spillover contamination of the control ring at the boundary becomes observable as X widens. For each facility f, ring $u \in \{\text{inner}, \text{outer}\}$, and year t, the outcome $Y_{f,u,t}$ is log effort-adjusted bird abundance or the log number of species. This equation describes the Difference-in-Differences (DiD) logic:

$$Y_{f,u,t} \;=\; \text{Facility-Ring FE}_{f,u} + \text{Year FE}_t + \beta \cdot \text{Inner}_u \times \text{Post}_{f,t} + X'_{f,t}\gamma + \varepsilon_{f,u,t}, \tag{3}$$

where Inner_u indicates the inner ring and $\text{Post}_{f,t}$ indicates years at or after the facility's satellite-inferred construction year $\hat{T}_f$. The vector $X_{f,t}$ contains time-varying controls for observer effort and local urbanization. The dynamic event-study version replaces $\text{Post}_{f,t}$ with a full set of event-time indicators—one for each year relative to construction—with the pre-construction year omitted as the reference.

I estimate the specification with the staggered DiD estimator of Callaway and Sant'Anna (2021), which is robust to the bias that affects two-way fixed effects under heterogeneous treatment timing. The estimator computes cohort-time-specific comparisons and aggregates them to cohort-size-weighted event-time coefficients β_e. The control set comprises all outer-ring observations (never treated) and inner-ring observations from facilities not yet constructed at t (not yet treated). I cluster standard errors at the facility level. The urbanization control in all specifications uses impervious surface growth in a fixed 15–30 km band around the facility, which proxies for regional development trends without absorbing the facility's own physical footprint. The identification assumption is not that facility locations are randomly assigned; it is that bird populations in the inner ring would have trended similarly to those in the outer ring in the absence of facility construction.

Figure 3 presents the estimated difference-in-differences at every integer threshold from 5 to 24 kilometers and plots the average treatment effect on the treated (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 statistically significant up to 18 kilometers and attenuates beyond. 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 use 14 kilometers as the validation threshold, a conservative cutoff well inside the significant range that captures the strongly-affected zone without reaching the marginal boundary near 18 kilometers. Table 2 reports the DiD estimates at the 14-kilometer 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). Panel C reports the same Post ATT under a more conservative specification that restricts the control set to never-treated outer-ring observations only. The abundance ATT moves from −9.2% to −8.9% and the richness ATT from −5.0% to −4.8% in the full-controls column (4). The result is robust to the choice of control group. Figure 4 traces the event-time path and shows that the decline begins at the time of construction and persists through the post-treatment window.[9]

For 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.3), indicating that ecological degradation persists as an ongoing byproduct of industrial activity rather than a one-time consequence of initial land use.

I test for four identification threats that could plausibly explain this pattern in the absence of a real facility effect. First, the decline could reflect pre-existing differences in ring trends

[9] 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, economically negligible relative to the 6–9% abundance decline. Observer spatial sorting does not drive the result.

or anticipatory behavior, rather than the facility itself. The dynamic event-study pre-period coefficients are economically and statistically zero, ruling both out: inner- and outer-ring abundance trend together before construction, and birds do not respond in anticipation of it. These pre-period tests also mitigate the convergence concern raised by Sager and Singer (2025), who warn that high-abundance areas may regress toward a lower mean regardless of treatment. Second, to discount the possibility of control-ring contamination, a displacement test finds no offsetting increase in bird populations in the rings immediately surrounding the impact zone (Internet Appendix Figure IA.2). The losses in the near-field are not offset by gains farther out, so spatial reshuffling does not contaminate the outer control ring. Third, for treatment-timing exogeneity, a randomization inference exercise permutes facility construction dates 1,000 times at each distance threshold (Figure 5). The placebo ATTs cluster around zero at every threshold, while the actual abundance estimates remain significant through 18 km, outside the 5th percentile of the placebo distribution. Richness follows the same pattern. Facility construction timing rather than ring geometry therefore drives the result. Fourth, the decline could reflect changes in observer effort rather than ecology. Adding effort controls (log observers, log hours) to the doubly-robust estimator (Column 4 of Table 2) shifts the post-construction ATT, to -9.2%. Differential observer effort therefore does not drive the result.

Furthermore, a direct chemical pathway corroborates the ecological finding. 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.4). A similar randomization inference exercise confirms that this lead contamination effect lies outside the placebo distribution (Figure IA.5).

Together, these results establish the ecological basis for BFI. The validation exercise shows that industrial facilities cause measurable bird declines within the 1–14 kilometer impact ring. The decline attenuates with distance, persists through the post-treatment window, survives placebo tests on facility timing and ring geometry, and coincides with elevated chemical contamination of the surrounding environment.

The index numerator instead uses abundance in the surrounding 15–30 km ring, separating causal validation from firm-level measurement. The outer ring approximates the biodiversity each site would support absent the facility's direct impact. Near-facility abundance would instead measure biodiversity remaining after that impact and could assign smaller footprints to facilities that caused greater losses. With the average causal decline held constant across facilities, outer-ring abundance is proportional to each facility's ecological externality. Aggregating this abundance across a firm's facilities and scaling by assets produces an externality intensity analogous to the externality-to-value ratio in Pedersen (2026).

4 The Biodiversity Footprint Return Premium

Having shown that industrial facilities affect nearby bird populations, I next test whether this location-based biodiversity characteristic is priced in equity markets. I do this by first showing that long-short portfolios sorted on BFI generate economically and statistically significant return spreads among small and mid-cap firms. Second, I distinguish the premium from the pollution premium, documented in prior work, showing that the biodiversity footprint carries pricing information beyond emissions intensity. Third, I document that the premium declines monotonically with firm size, which motivates the mechanism tested in Section 5.

4.1 Cross-sectional asset pricing tests

I test whether firms with a larger biodiversity footprint—those running more operations in biodiverse areas relative to their size—earn higher subsequent returns. Table 3 reports returns for portfolios formed on conditional double-sorts of market capitalization and BFI, using the vintage-year timing from Section 2 so that all portfolio assignments rely on data that was publicly available before the return month. In the universe of small firms, conditional sorts generate a Fama-French three-factor-adjusted high-minus-low return spread of 0.89% per month (roughly 11% per year). The premium remains economically large in mid-cap firms at 0.52% per month and attenuates to zero for large firms. The premium survives factor models augmented with a liquidity factor and with the Pástor et al. (2021) green factor, which distinguishes it from the return spread associated with ESG tilts.

I show that the premium is distinct from the pollution premium documented by Hsu et al. (2023). BFI and Releases are built from the same TRI facilities, but BFI uses each facility's location to measure the surrounding biodiversity, while Releases uses the emissions those facilities report. Columns 5–9 in Table 4 shows that BFI carries a premium of around 0.16% per month (2% per year), and adding standardized Releases leaves this estimate unchanged while Releases itself contributes little independent explanatory power. Internet Appendix IA.4 extends this analysis with size and idiosyncratic-volatility interactions and shows that the premium on BFI is robust, while Releases is not (Internet Appendix Table IA.3).

Table 4 also shows that the BFI index remains distinct from known return predictors. The BFI coefficient is 0.14% per month with industry fixed effects alone and rises to 0.27% once book-to-market and momentum enter, since both correlate with BFI and offset the spread in the univariate specification. The full-control estimate survives a modified Altman (1968) Z-score distress control, following Akey and Appel (2021).

Another concern is that the premium is driven by scaling BFI by total assets rather than by the ecological footprint itself. I address this concern by using aggregate bird abundance without asset scaling and find the same return pattern (Appendix A.1). Replacing total assets with sales or PP&E also yields comparable premiums (Table A.2), showing that the result is not specific to the choice of denominator. Appendix A.2 then separates BFI into abundance per facility, its intensive margin, and facility count relative to assets, its extensive margin. The return premium is concentrated in the extensive margin, with facility intensity predicting returns of 0.26% per month while abundance per facility is insignificant. This is consistent with the externality-to-value structure of Pedersen (2026), in which a firm's total externality aggregates across its facilities before being scaled by value, so that only the aggregate footprint is priced.

The small-firm concentration may instead reflect illiquid microcap stocks, yet I show that the BFI premium declines broadly with firm size. Internet Appendix Table IA.4 reports a granular sort into ten size deciles crossed with BFI quintiles. The return spread falls from nearly 1.7% per month in the smallest stocks to an economically negligible spread in the largest, significant across most deciles rather than forming a discrete small-firm effect. An industry-conditioned triple sort produces nearly identical results (Appendix IA.6, Table IA.5).

Finally, the premium may reflect the sparsity of early eBird data. Restricting the sample to post-2002 eBird vintages, when citizen science coverage is densest, strengthens rather than weakens the results. The small-firm alpha rises to 85 to 94 basis points per month across five factor models, and the Fama-MacBeth BFI coefficient remains significant across all specifications (Internet Appendix Table IA.6). All of this points to each firm's own operating footprint in biodiverse areas—rather than sparse early data or firm size—as the source of the premium.

5 The Cash-Flow Expectations Channel

The premium documented in the cross-sectional return tests could be interpreted in two different ways: it may be compensation for a cost that the market prices as risk, or it could reflect expectations that misjudge that cost. A risk premium for BFI should load broadly across the size distribution, since large firms retain wide cross-sectional dispersion in bird abundance (σ = 1.35 versus 1.10 for small firms); however, since the premium attenuates for large firms, I test the cash-flow expectations channel first. I start by showing that the biodiversity footprint raises operating costs through COGS. Second, I document how investors misinterpret this cost structure, which generates positive earnings surprises and concentrated announcement-date returns.

5.1 Operating costs

To tie the return premium to firm fundamentals, I estimate a quarterly panel specification linking the biodiversity footprint to operating costs,

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

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 $\mathrm{BFI}_{j,y}$ denotes the TRI vintage year, not the calendar year of the quarter.[10] The control vector $\mathbf{X}$ includes log assets, log book-to-market, return-on-assets, leverage, and tangibility, with standard errors clustered by firm. I estimate the same specification for Selling, General, and Administrative Expenses (SG&A) scaled by sales, and for year-over-year changes in both cost ratios, which separate a stable cost level from a trending one.

Firms with larger biodiversity footprints have lower operating performance across multiple measures. Table 5 shows that a one-standard-deviation increase in BFI lowers quarterly income-to-assets by 0.06%, reduces year-over-year sales growth by 1.2%, and decreases operating margins by 0.43%. Decomposing operating costs into components, a one-standard-deviation increase in BFI raises cost of goods sold by 0.35% of sales, with no corresponding effect on selling, general, or administrative expenses (Table 5). The COGS effect appears in the level but not in the year-over-year change, indicating a stable cost structure associated with a larger biodiversity footprint rather than an escalating cost shock.

The cost effect follows the same extensive-margin pattern as the return premium. Ap-

[10]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.

pendix Table A.4 decomposes BFI into abundance per facility and facility count. Per-facility abundance is insignificant, whereas facility count predicts higher production costs, indicating that the penalty increases with the breadth of a firm's operating footprint rather than the biodiversity of its average site. The difference-in-differences evidence in Section 3 shows that industrial facilities reduce local biodiversity on average, giving this operating breadth a direct ecological interpretation.

Although Derrien et al. (2025) document consumer-demand effects for high-profile ESG events, the biodiversity footprint does not attract greater media coverage (Internet Appendix Table IA.8). Its low public salience distinguishes it from a reputational demand shock, whose signature is defensive spending on marketing and customer retention. Consistent with a production-side channel, COGS rises while SGA stays flat, both in terms of levels and in year-over-year changes (Table 5). High-footprint firms do grow sales more slowly, but with no accompanying SGA response or media attention, pointing to an elevated cost structure rather than a consumer-driven loss of demand.

Evidence from regulatory data indicates that the cost penalty is independent of the regulatory pipeline. In a facility-year panel of 236,325 observations across 13,731 TRI facilities 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.5).[11] Emissions, by contrast, is strongly correlated with all four outcomes. Regulators appear to target emissions levels rather than the ecology of the surrounding area. The 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.

Other observable cost channels do not absorb the BFI cost effect. Including proxies for regulatory stringency, the share of facilities in mini-NEPA (State Environmental Policy Act) states, an indicator for water-intensive industries, off-site waste intensity, and the share of releases subject to CERCLA cleanup liability as controls in the COGS regression leaves

[11] The count and amount specifications use Poisson pseudo-maximum likelihood, and the penalty-probability specification uses a linear probability model.

the BFI coefficient unchanged, with each candidate channel individually insignificant (Internet Appendix Table IA.13). The null results are consistent with an extensive-margin cost channel. The penalty is not tied to any single regulatory or industry exposure but reflects operating breadth, measured by the number of sites a firm operates relative to its asset base. The insignificant CERCLA share separates it from the cleanup-liability channel of Bellon (2025).

While external operating-environment proxies do not absorb the cost effect, the production technology of high-BFI firms accounts for part of it. A one-standard-deviation increase in BFI raises the log labor-to-capital ratio by 5.5%, the labor share of value added (Donangelo, 2014) by 1.4%, and the variable share of operating expense (Novy-Marx, 2011) by 0.6% (Internet Appendix Table IA.14). Adding the labor-to-capital ratio to the cost regression of Equation (4) attenuates the BFI coefficient by roughly one-fourth on the labor-channel subsample. The labor channel reflects employee headcount rather than wages. Firm-level compensation per employee and BLS state by 6-digit NAICS wage exposure are unrelated to BFI, while headcount per dollar of fixed capital rises with BFI. Operating breadth raises costs partly because high-BFI firms staff more facilities per dollar of capital.

5.2 Earnings expectations

If the return premium reflects errors in expectations rather than risk compensation, three patterns should follow. First, analysts should hold ex-ante pessimistic views of high-BFI firms and embed that pessimism in their published forecasts. Second, BFI should predict positive earnings surprises, because the market underestimates future earnings for firms with an extensive footprint. Third, returns should concentrate around the dates when hard information corrects those expectations (La Porta, 1997; Engelberg et al., 2018). In other words, investors observe elevated costs but interpret them too pessimistically, meaning 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 6 reports panel regressions of this implied return on standardized BFI, 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 BFI corresponds to a 4.6% lower implied return (Column 3), and the coefficient is unchanged when toxic releases enter as an additional control (Column 4). Column 5 interacts BFI with within-vintage size terciles in a saturated specification where every control, industry effect, and year effect is also interacted with the size tercile. The pessimism declines monotonically across size groups, with a marginal implied-return reduction of 10.4% per standard deviation of BFI in small firms, 3.3% in mid-cap firms, and 1.0% in large firms (not statistically significant). A joint test of whether the BFI effect is the same across the three size terciles is rejected at the 10% level. The pattern is consistent with a mechanism in which analysts observe the depressed margins and expect further deterioration, rather than recognizing those margins as a stable cost of the firm's broad operating footprint.

Next, I test whether firms beat these pessimistic forecasts using two earnings-surprise benchmarks. The analyst consensus forecast (SUE2) tests whether sell-side analysts are too pessimistic about high-BFI firms. The seasonal random-walk benchmark (SUE1) compares current earnings to the same quarter one year earlier, and therefore does not depend on analyst forecasts (Chan et al., 1996). BFI predicts positive surprises against both benchmarks (Table 7, Panel A). For SUE2 the coefficient attenuates when industry fixed effects are added, indicating that analysts partially capture the operating-cost component of BFI through industry peer benchmarking. I/B/E/S coverage also skews the SUE2 sample toward larger firms, where the BFI premium is weakest. For SUE1, the coefficient remains stable with industry fixed effects, showing that high-BFI firms also beat a mechanical seasonal benchmark. Releases are negatively related to analyst-based surprises but does not predict seasonal surprises, suggesting that the market partially incorporates observable pollutant information into analyst forecasts.

Panel B of Table 7 shows that the earnings surprises predicted by BFI account for much of the announcement-return premium. In the baseline specification, a one-standard deviation increase in BFI predicts a 0.10% higher announcement returns. Adding the analyst consensus surprise absorbs most of this effect, although the SUE2 specification 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 is measurable for the entire sample, absorbs roughly one-third of the baseline BFI coefficient. Including both surprises together leaves BFI with no residual effect on the announcement return (Column 8), consistent with the surprises mediating the announcement-date return correction. Analyst expectations absorb more of the announcement return than the seasonal benchmark in the covered subsample because the market reacts to analyst forecasts rather than last year's earnings. BFI therefore predicts which firms will beat analyst expectations, and the market reprices when those expectations are corrected at the earnings date.

The announcement window accounts for a large but incomplete fraction of the return premium. In small firms, four quarterly announcement windows capture 32% of the annualized five-factor alpha, falling to 11% in mid-caps and negligible in large firms (Table 8, Panel A). The correction continues after the three-day announcement window, consistent with gradual information diffusion among heterogeneous investors (Hong and Stein, 1999). In the full sample, the high−low BFI spread grows from 0.40% over $[-1, +1]$ to 2.12% over $[-1, +60]$ trading days, with the strongest drift among small firms (Panel B). Earnings announcements therefore begin, but do not complete, the repricing of high-footprint firms. Atilgan et al. (2023) document a parallel pattern for carbon emissions, with 30–50% of the emission premium concentrating at earnings announcements.

The two surprise benchmarks distinguish analyst pessimism from a broader seasonal expectation error. SUE2 shows that human forecasters are too pessimistic about high-BFI firms: analysts observe depressed margins but do not identify the broad operating footprint behind them, so they extrapolate further deterioration. The cost penalty documented in Section 5.1 is stable rather than escalating, such that maintaining the existing cost structure is enough to beat pessimistic expectations. The attenuation of the SUE2 coefficient with industry fixed effects indicates that analysts partially learn these costs through peer benchmarking. SUE1, on the other hand, shows that the effect is not confined to analyst forecasts or the firms they cover. High-BFI firms also beat a mechanical seasonal benchmark, and the coefficient is stable with industry fixed effects.[12] Thus, the earnings surprise evidence reflects both analyst pessimism in the analyst covered sample and a broader seasonal expectation error.

Greater analyst coverage does not resolve these expectational errors. Adding analyst coverage to the Panel A consensus surprise specification has no effect on the BFI coefficient, and the interaction of analyst coverage with BFI is similarly null for announcement returns (Internet Appendix Table IA.11). Analysts observe the persistent costs but appear to interpret them as a decline rather than as a feature of the firm's operating footprint, consistent with the omitted-attribute mechanism of Hong et al. (2007). Greater coverage does not correct this, and the premium persists in part because high-BFI firms are disproportionately outside the analyst-covered universe. In fact, 29% of high-BFI firm-quarters have zero analyst coverage, compared with 12% for low-BFI firms, so the seasonal forecast error goes uncorrected for the firms that generate the largest return spread.

Together, these results establish the cash-flow expectations channel. High-BFI firms bear a stable cost penalty that investors misinterpret as deteriorating fundamentals. They beat both analyst forecasts and seasonal benchmarks, and the associated returns concentrate around the earnings announcements that correct those expectations. The evidence therefore links the breadth of firms' biodiversity operating footprints to returns through predictable cash-flow news rather than compensation for risk.

[12] A stable cost ratio does not imply flat earnings. The elevated cost of goods sold holds as a share of sales, so margins persist and earnings grow with sales, and the seasonal surprise is the gap between the margin decline investors extrapolate and the firms' actual growth.

6 The Information Environment

Having established persistent forecasting errors as the channel of mispricing, I next examine why the errors persist. The operating footprint needed to construct BFI is publicly visible in TRI facility disclosures, yet standard information channels do not translate it into a measure of biodiversity exposure. Firms rarely disclose their biodiversity footprints and additional analyst coverage produces no incremental price discovery. Even regulatory penalties, when they occur, are priced only slowly.

6.1 The information gap

Firms rarely discuss biodiversity in their disclosures, and the highest-footprint firms disclose the least. Only 2% of 10-K filings contain biodiversity-related language (Giglio et al., 2026). For mandatory 10-K disclosure, BFI is not related to whether a firm uses biodiversity language in any specification (Internet Appendix IA.11).

The CDP operates the most widely used voluntary corporate environmental survey and has recently added a biodiversity module to its annual questionnaire (covering 2022–2025 for the firms in my sample). The module asks each firm whether it operates in or near biodiversity-sensitive areas, whether it monitors biodiversity indicators, whether it has taken biodiversity-related actions, whether it has made public biodiversity commitments, and whether it voluntarily publishes biodiversity information—the five indicators in Internet Appendix Table 9. Fewer than one in five sample firms respond at all, and the survey elicits only yes-or-no answers rather than any measured quantity of ecological impact. The contrast with carbon is sharp: across CDP's respondents, roughly 6,300 firms report a precise emissions figure, yet the biodiversity module collects no comparable quantity, and only about a third of respondents acknowledge even operating near biodiversity-sensitive areas.

For voluntary CDP participation, the relationship is inverted. Internet Appendix Table 9 shows that BFI is negative across all five CDP biodiversity indicators. A one-standard-deviation increase in BFI significantly lowers the probability of reporting biodiversity actions by 11.1 percentage points and of publishing voluntary information by 9.9 percentage points, so the firms with the largest biodiversity operating footprints disclose the least.

Public news is an additional channel that could reveal a firm's footprint, even when the firm and its raters do not. I use RepRisk data, which records biodiversity-related news incidents for each firm, to ask whether BFI is visible through news. Regressing a firm's annual incident count on BFI with firm and year fixed effects, I find that BFI does not predict news coverage, which instead correlates with toxic releases (Internet Appendix Table IA.8). I then test whether the return premium is concentrated in periods when biodiversity is more salient. I interact BFI with the New York Times biodiversity attention index of Giglio et al. (2026), which measures how often the Times discusses biodiversity, and find that the interaction is also insignificant (Internet Appendix Table IA.9).

Nor does more financial intermediation resolve the mispricing. Greater analyst coverage fails to attenuate the BFI premium (Internet Appendix Table IA.10), which points away from costly information acquisition as the main friction. BFI is built from public TRI facility locations and eBird counts, so the issue is not that the underlying data are expensive to produce, as in Veldkamp (2006). Instead, the issue is interpretive, as in the Edmans (2011) study of the employee-satisfaction premium, where investors fail to incorporate a public but nonstandard firm attribute into their valuation. The biodiversity operating footprint is omitted from earnings models (Hong et al., 2007), so additional coverage only implies that more forecasters are using the same incomplete framework, rather than correcting the error. Consistent with this interpretation, the premium is strongest among small and low-institutional-ownership firms (Internet Appendix Table IA.10), while short interest, liquidity, and idiosyncratic volatility do not explain the spread (Internet Appendix Table IA.7).

Under recognized ambiguity, investors identify the uncertain attribute and demand compensation for bearing it (Epstein and Schneider, 2008), but this channel is difficult to reconcile with three empirical patterns: additional analyst coverage does not attenuate the premium, high-BFI firms continue to generate positive earnings surprises across quarters, and the correction arrives with a lag rather than at once (the delayed return response to penalties, Internet Appendix IA.14). These patterns suggest that investors do not perceive the uncertainty. Instead, they observe aggregate costs without identifying the operational and ecological source, and the errors in expectations persist.

7 Conclusion

Green finance requires environmental externalities to be measured in a form that can be compared across firms and related to firm value. This problem is tractable when the externality can be expressed using common units. Greenhouse gas emissions, for example, can be converted into CO_2 equivalents, disclosed through formal channels, or imputed from firm fundamentals such as sales and industry. Biodiversity resists this treatment because each facility's ecological impact depends on local ecosystem conditions, while a firm's total footprint accumulates across the facilities it operates. Analysts do not incorporate this footprint into their forecasts, and firms with larger footprints provide less voluntary biodiversity information. Investors observe the higher production costs associated with this footprint but do not identify their operational and ecological source. Analysts misread these stable expenses as deteriorating fundamentals and issue pessimistic price targets, while high-footprint firms subsequently generate positive earnings surprises. Put differently, investors form expectations that omit an attribute they cannot easily evaluate, and that omission produces a return premium. In a panel of 1,762 U.S. firms from 1988 to 2024, I document the resulting return premium among firms that operate more ecologically consequential facilities relative to their asset base.

The evidence points to interpretation rather than data availability as the central obstacle. Emerging frameworks, such as the Taskforce on Nature-related Financial Disclosures, require location-specific ecological assessments, but voluntary reporting leaves coverage uneven. This limitation is especially important for small and mid-cap firms where the premium is largest. For these firms, the information is public, but translating facility locations and ecological data into a comparable, valuation-relevant signal remains specialized, unstandardized work. The findings therefore show that public disclosure alone is insufficient without measures that translate operational data into comparable firm-level externality intensities.

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," *The Journal of Finance*, 2021, *76* (1), 5–55.

Altman, Edward I, "Financial ratios, discriminant analysis and the prediction of corporate bankruptcy," *The 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 Mispricing?," *European Corporate Governance Institute – Finance Working Paper 940/2023*, 2023.

Bellon, Aymeric, "Fresh Start or Fresh Water: The Impact of Environmental Lender Liability," *Journal of Finance*, 2025. Forthcoming. Available at SSRN: https://ssrn.com/abstract=3877378.

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

Callaway, Brantly and Pedro HC Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of econometrics*, 2021, *225* (2), 200–230.

Chan, Louis KC, Narasimhan Jegadeesh, and Josef Lakonishok, "Momentum strategies," *The journal of Finance*, 1996, *51* (5), 1681–1713.

Chen, Jason, "Redeploying dirty assets: The impact of environmental liability," *Journal of Financial Economics*, 2025, *170*, 104070.

Chen, Luoye, Dragon Yongjun Tang, Liying Wang, and Yi Zhang, "Biodiversity Regulation and Corporate Adaptation: Evidence from the Endangered Species Act," 2024. Working paper.

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

Derrien, François, Philipp Krueger, Augustin Landier, and Tianhao Yao, "ESG news, future cash flows, and firm value," *The Journal of Finance*, 2025, *80* (6), 3499–3554.

Dimson, Elroy, Marcin Kacperczyk, and Laura Starks, "Biodiversity and natural resource finance," *Review of Finance*, 2026, *30* (1), 1–9.

Donangelo, Andres, "Labor mobility: Implications for asset pricing," *The Journal of Finance*, 2014, *69* (3), 1321–1346.

eBird, "eBird Basic Dataset. Version: EBD_relDec-2024," Cornell Lab of Ornithology, Ithaca, New York 2024.

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.

Fraixedas, Sara, Andreas Lindén, Markus Piha, Mar Cabeza, Richard Gregory, and Aleksi Lehikoinen, "A state-of-the-art review on birds as indicators of biodiversity: Advances, challenges, and future directions," *Ecological Indicators*, 2020, *118*, 106728.

Garel, Alexandre, Arthur Romec, Zacharias Sautner, and Alexander F Wagner, "Do investors care about biodiversity?," *Review of Finance*, 2024, *28* (4), 1151–1186.

_ , _ , _ , and Alexander F. Wagner, "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," Technical Report, National Bureau of Economic Research 2024.

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

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, Joséphine Gantois, and Anouch Missirian, "The Effectiveness of Local Conservation Ballots," in "AEA Papers and Proceedings," Vol. 115 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2025, pp. 403–408.

Hong, Harrison and Jeremy C Stein, "A unified theory of underreaction, momentum trading, and overreaction in asset markets," *The Journal of finance*, 1999, *54* (6), 2143–2184.

_ , _ , and Jialin Yu, "Simple forecasts and paradigm shifts," *The Journal of Finance*, 2007, *62* (3), 1207–1242.

Hou, Kewei, Chen Xue, and Lu Zhang, "Replicating Anomalies," *The Review of Financial Studies*, 2020, *33* (5), 2019–2133.

Hsu, Po-Hsuan, Kai Li, and Chi-Yang Tsou, "The pollution premium," *The Journal of Finance*, 2023, *78* (3), 1343–1392.

La Porta, Rafael, "Expectations and the Cross-Section of Stock Returns," *Journal of Finance*, 1997, *52* (5), 1715–1742.

Merton, Robert C., "A Simple Model of Capital Market Equilibrium with Incomplete Information," *Journal of Finance*, 1987, *42* (3), 483–510.

Novy-Marx, Robert, "Operating leverage," *Review of Finance*, 2011, *15* (1), 103–134.

Pedersen, Lasse Heje, "Carbon Pricing versus Green Finance," *Journal of Finance*, 2026, *81* (2), 561–602.

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), 1–36.

U.S. Environmental Protection Agency, "2022 TRI National Analysis," Technical Report, U.S. Environmental Protection Agency, Office of Pollution Prevention and Toxics March 2024.

Veldkamp, Laura L., "Media Frenzies in Markets for Financial Information," *American Economic Review*, 2006, *96* (3), 577–601.

Ľuboš Pástor, Robert F. Stambaugh, 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: Bird abundance and the location of industrial facilities

County-level bird abundance and the geographic distribution of the industrial facilities used to construct the Biodiversity Footprint Index. Each county is shaded by the geometric mean of total birds reported per eBird checklist, pooling all checklists within the county from 2010 through 2024; darker shades denote higher abundance. Black points mark the locations of TRI facilities.

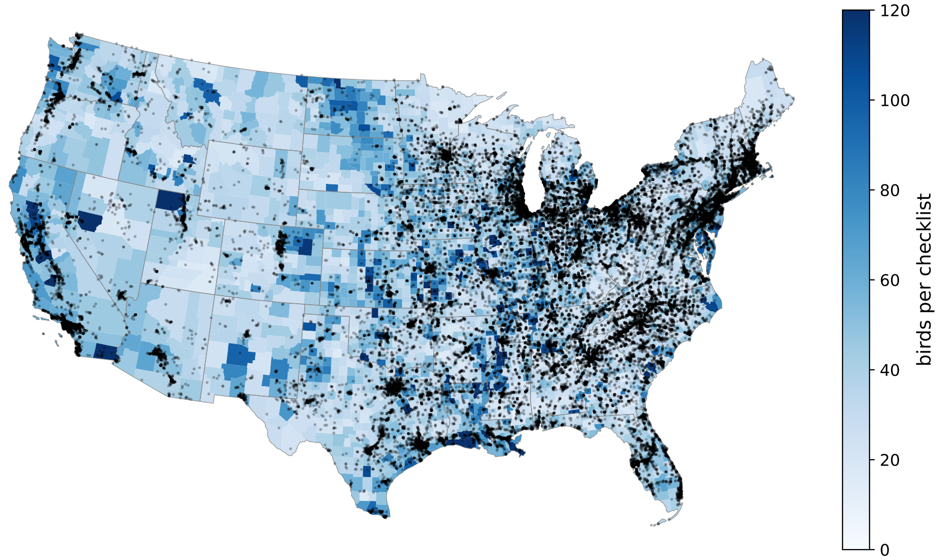

Figure 2: Visual validation of satellite-inferred construction dates

High-resolution (~1 m) aerial imagery from the USDA National Agriculture Imagery Program (NAIP) before and after the construction year detected by the NLCD impervious surface algorithm. The images are warped to an identical geographic grid so that landscape features align across years, covering approximately 1.8 km×1.8 km centered on the facility's TRI-reported coordinates. The site changes from a South Carolina field in 2011 to a tire manufacturing plant by 2015, consistent with the satellite-inferred construction year of approximately 2013.

Figure 3: 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

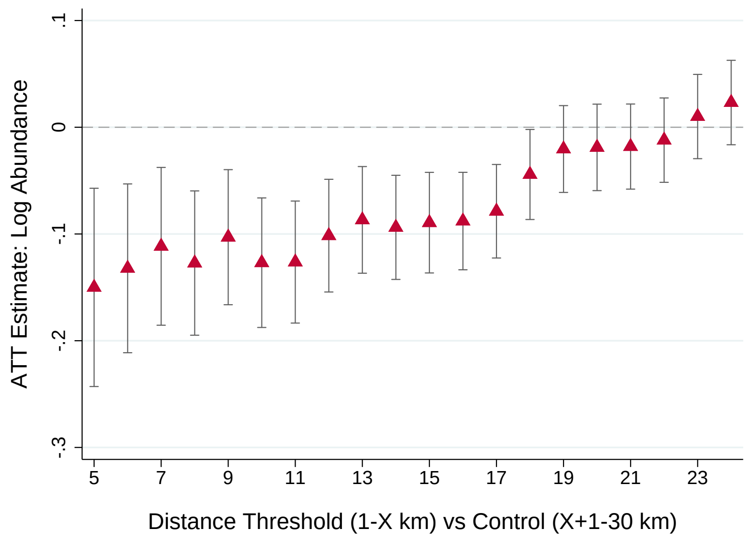

(b) Richness

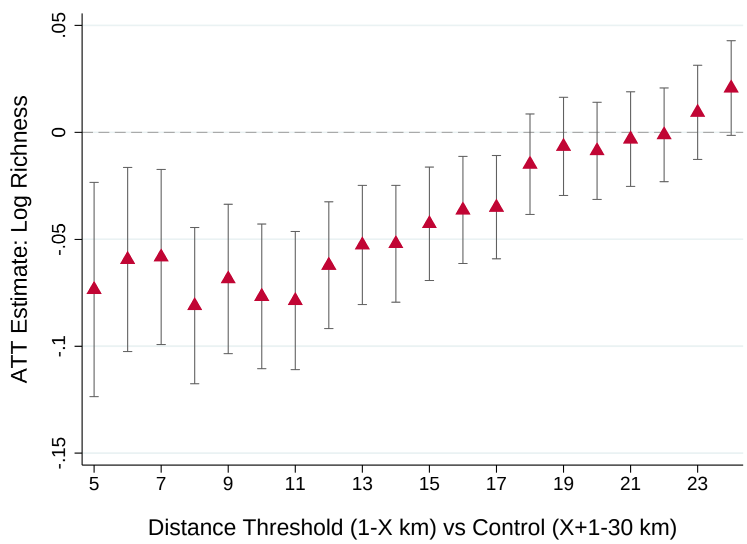

Figure 4: Ecological impact of industrial facility construction

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 the onset of construction, the boundary between the last pre-construction year ($t = -1$) and the construction year ($t = 0$). Panel (a): log bird abundance declines 6% to 9% after construction. Panel (b): log species richness declines 3% to 5%.

(a) Abundance

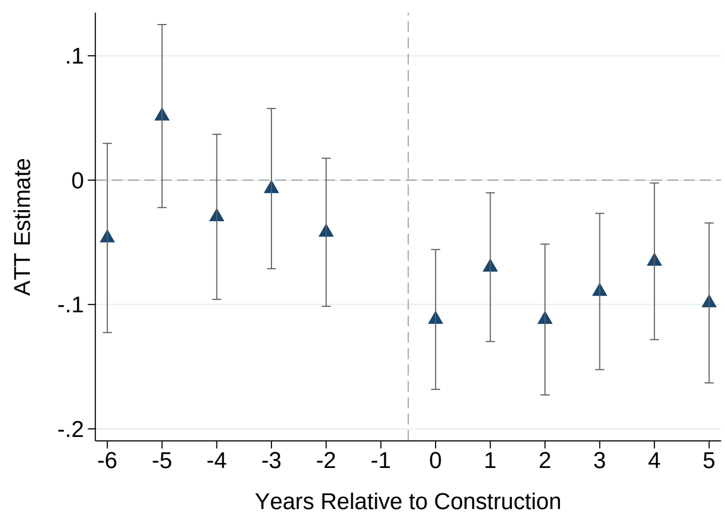

(b) Richness

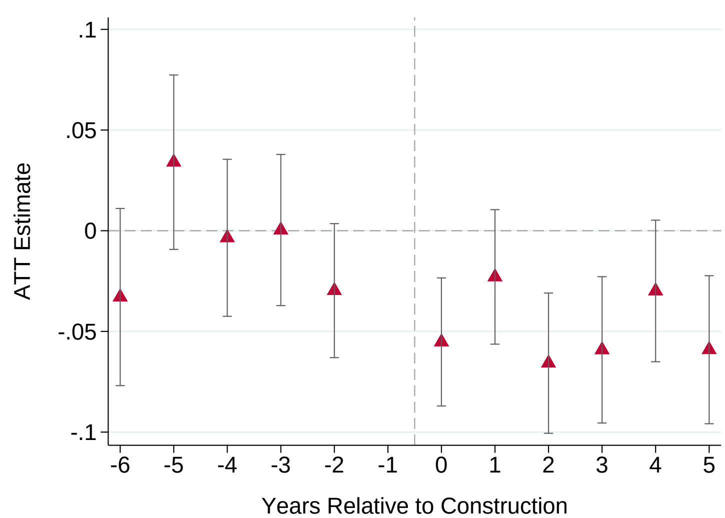

Figure 5: Placebo distributions with actual estimates

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

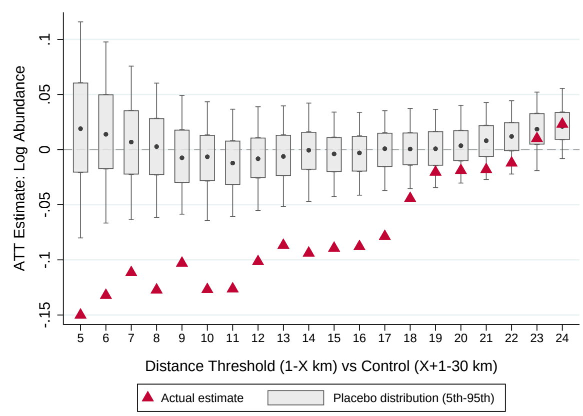

(b) Richness ATT estimates

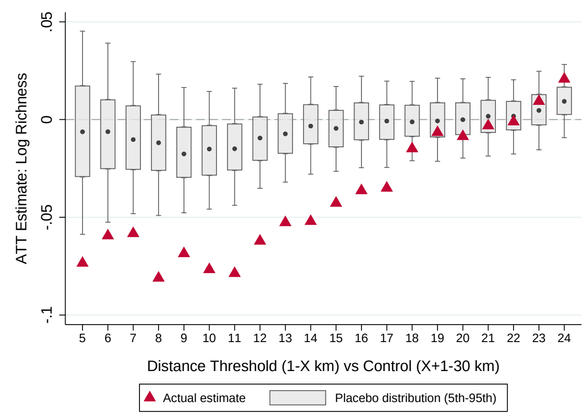

(c) Abundance z-statistics

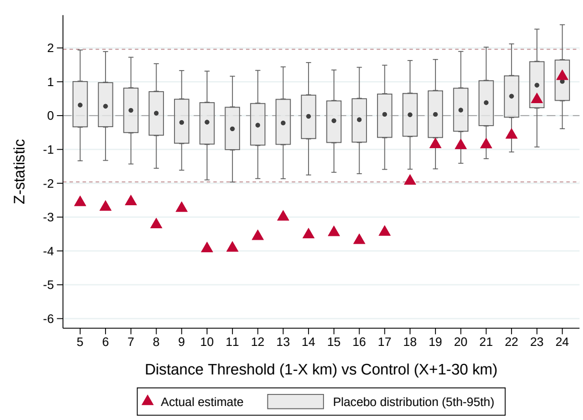

(d) Richness z-statistics

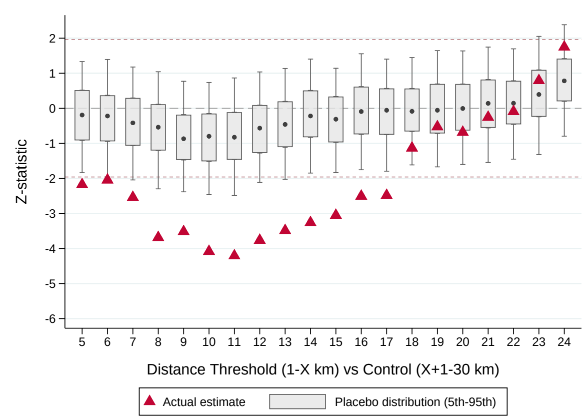

Tables

Table 1: Summary statistics

The sample comprises 1,762 TRI-linked firms from October 1988 through December 2024. Monthly variables are summarized over the 172,938 firm-months of the Fama-MacBeth sample. BFI 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; BFI and Releases follow the vintage-year timing rule in Section 2.1. All continuous variables are winsorized at the 1st and 99th percentiles.

	Mean	SD	p5	p25	Median	p75	p95	*N*
Panel A: Monthly variables								
Monthly return (%)	1.09	10.75	-16.29	-4.88	0.87	6.61	19.21	172,938
BFI	-3.59	1.77	-6.76	-4.75	-3.48	-2.34	-0.77	172,938
Releases	3.11	3.15	-3.13	1.47	3.63	5.35	7.31	172,938
ln(bird abundance)	4.02	1.30	1.95	3.07	4.00	4.95	6.25	172,938
Log market cap	14.32	2.08	10.71	12.92	14.37	15.74	17.73	172,938
Log book-to-market	-0.70	0.73	-1.99	-1.12	-0.66	-0.23	0.40	172,938
Momentum (-12, -2)	0.13	0.35	-0.43	-0.05	0.13	0.31	0.72	172,938
ROA	11.90	20.33	-19.45	4.64	12.65	21.69	40.70	172,938
Leverage	0.27	0.15	0.01	0.16	0.27	0.37	0.54	172,938
Investment rate (I/K)	0.15	0.19	0.02	0.04	0.09	0.17	0.52	172,938
Idiosyncratic volatility	0.02	0.01	0.01	0.01	0.02	0.03	0.05	172,938
Number of TRI facilities	6.65	9.07	1.00	2.00	3.00	8.00	25.00	172,938
Log total assets	7.60	1.87	4.52	6.26	7.58	8.93	10.70	172,938
Log tangibility (PPE/A)	-1.33	0.64	-2.49	-1.75	-1.27	-0.86	-0.35	172,868
Amihud illiquidity ($\times 10^6$)	0.17	0.80	0.00	0.00	0.00	0.01	0.66	170,786
Turnover	6.76	6.11	0.90	2.72	5.04	8.61	19.07	170,778
Analyst coverage	29.54	21.78	3.00	12.00	25.00	43.00	71.00	145,554
Panel B: Quarterly operating variables								
COGS/Sales (%)	68.48	16.34	36.43	60.01	70.76	79.62	90.70	69,936
SGA/Sales (%)	17.19	12.38	0.00	8.30	14.92	23.84	41.66	70,263
Sales growth (%, YoY)	7.14	22.34	-24.97	-3.67	5.05	14.77	45.97	59,515
Gross profitability (GP/A)	0.08	0.05	0.02	0.05	0.07	0.10	0.17	69,935
Asset turnover	0.28	0.14	0.08	0.18	0.26	0.35	0.54	70,263
Altman Z-score	10.02	35.96	0.64	1.40	2.55	5.00	29.99	64,232
Abatement indicator	0.52	0.50	0.00	0.00	1.00	1.00	1.00	57,845
Panel C: Earnings announcement variables								
SUE (seasonal random walk)	-0.01	1.64	-2.56	-0.50	0.07	0.63	2.48	68,904
SUE (analyst-based)	0.98	3.42	-3.80	-0.41	0.67	2.18	6.71	53,964
Long-term growth forecast	12.29	10.44	-0.31	7.08	11.00	15.00	30.00	30,743

Table 2: 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. Treated units are facility-ring-years for the inner ring (1–14 km). The control set in Panels A and B comprises all outer-ring (15–30 km) observations, which are never treated, plus inner-ring observations from facilities not yet constructed at time t. Panel C reports Post ATT under a more conservative specification that restricts the control set to never-treated outer-ring observations only. Both specifications use the same panel; they differ in the subset of comparisons computed at each cohort × event-time cell. The outcome is log adjusted abundance (Panel A, top of Panel C) or log adjusted species richness (Panel B, bottom of Panel C). 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)
Panel C: Robustness — never-treated controls only				
Abundance Post ATT	-0.060^{**}	-0.071^{***}	-0.069^{***}	-0.089^{***}
	(−2.56)	(−2.80)	(−2.73)	(−3.57)
Richness Post ATT	-0.031^{**}	-0.038^{***}	-0.037^{***}	-0.048^{***}
	(−2.35)	(−2.70)	(−2.62)	(−3.45)
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 3: Size × BFI double sort: multi-factor alphas

This table reports alphas of high-minus-low BFI spread portfolios from a conditional double sort on market equity terciles and BFI terciles. Each vintage year, firms are first sorted into size terciles, then into BFI terciles within each size group. Value-weighted portfolios are formed within each of the nine size–BFI cells. Each cell reports the alpha of the high-BFI minus low-BFI spread portfolio, 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); column (6) adds traded liquidity (Pástor and Stambaugh, 2003); and column (7) adds the green factor (Pástor, Stambaugh, and Taylor, 2021). 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
Small	0.862***	0.890***	0.814***	0.819***	0.739***	0.836***	0.822***
	(0.142)	(0.141)	(0.144)	(0.151)	(0.155)	(0.149)	(0.217)
Mid	0.524***	0.516***	0.472***	0.489***	0.397***	0.470***	0.565***
	(0.107)	(0.106)	(0.110)	(0.110)	(0.112)	(0.114)	(0.184)
Big	0.014	-0.031	-0.135	-0.101	-0.087	-0.155	0.063
	(0.146)	(0.128)	(0.132)	(0.140)	(0.127)	(0.132)	(0.170)
Pooled	0.467***	0.459***	0.384***	0.402***	0.350***	0.384***	0.483***
	(0.085)	(0.078)	(0.080)	(0.081)	(0.082)	(0.082)	(0.117)
N months	434	434	434	432	434	434	180

Table 4: Fama-MacBeth regressions: BFI and stock returns

This table reports monthly Fama-MacBeth cross-sectional regressions of excess stock returns on the Biodiversity Footprint Index (BFI) and progressively expanded control sets. The dependent variable is the monthly excess return in percent. BFI 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 returns in excess of the market, within each month), 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 BFI'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 172,938 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 BFI 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)
BFI (standardized)	0.139***	0.265***	0.250***	0.173***	0.143***	0.174***	0.155***	0.150***	0.165***
	(2.67)	(5.41)	(4.81)	(3.81)	(2.98)	(3.79)	(3.46)	(3.18)	(3.39)
Releases (standardized)					0.068**	0.057*	0.073**	0.049	0.047
					(2.14)	(1.89)	(2.40)	(1.56)	(1.51)
Log book-to-market		-0.915***	-0.927***	-0.727***	-0.743***	-0.741***	-0.739***	-0.762***	-0.882***
		(-7.89)	(-9.21)	(-7.04)	(-7.25)	(-7.43)	(-7.48)	(-7.37)	(-8.62)
Momentum ($t-12, t-2$)		-0.347	-0.384	-0.781***	-0.783***	-0.748***	-0.763***	-0.813***	-0.836***
		(-1.31)	(-1.44)	(-2.93)	(-2.95)	(-3.00)	(-3.07)	(-3.21)	(-3.29)
Log ME			-0.013	-0.093**	-0.100**	-0.005	0.010	0.013	-0.011
			(-0.28)	(-2.20)	(-2.32)	(-0.13)	(0.29)	(0.33)	(-0.28)
ROA				0.032***	0.032***	0.033***	0.032***	0.032***	0.033***
				(10.48)	(10.37)	(11.75)	(11.57)	(10.73)	(10.48)
Leverage				-0.085	-0.127	-0.320	-0.380	-0.349	-0.979***
				(-0.22)	(-0.33)	(-0.88)	(-1.06)	(-0.95)	(-2.60)
Idiosyncratic volatility						23.629***	25.036***	25.683***	26.485***
						(3.72)	(3.95)	(3.99)	(4.10)
Investment rate (I/K)							-1.613***	-1.654***	-1.719***
							(-5.34)	(-5.14)	(-5.30)
Asset turnover (S/A)								0.629*	0.441
								(1.72)	(1.15)
Gross profitability (GP/A)								-1.095	-0.498
								(-0.80)	(-0.36)
Tangibility (ln PPE/A)								0.084	0.139*
								(1.18)	(1.89)
Altman Z-score									-0.000***
									(-2.94)
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
N months	428	428	428	428	428	428	428	427	427
Avg. firms/month	404	404	404	404	404	404	404	397	387
Avg. R^2	0.08	0.11	0.12	0.14	0.14	0.15	0.16	0.17	0.18

Table 5: Biodiversity exposure, operating performance, and cost structure

This table relates biodiversity exposure to operating performance and cost structure on a quarterly panel of TRI-reporting firms, 1989–2024. Columns (1)–(3) report operating performance: income before extraordinary items over lagged assets, year-over-year sales growth, and operating margin ((Sales − XOPR)/Sales). Columns (4)–(7) decompose the cost structure: COGS/Sales and SG&A/Sales in levels, and their year-over-year changes. All dependent variables are in percentage points. Coefficients near zero on the change columns indicate that biodiversity exposure raises production costs in levels rather than as a trend, while administrative costs are unaffected in both levels and changes. BFI is standardized within TRI vintage year. Columns (1)–(5) 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. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

	Operating performance			Cost structure			
	(1) Inc/A	(2) Sales Gr.	(3) Oper. Margin	(4) COGS/Sales	(5) SG&A/Sales	(6) ΔCOGS/Sales	(7) ΔSG&A/Sales
BFI (standardized)	-0.064***	-1.194***	-0.429***	0.353***	-0.025	-0.056	0.002
	(-3.08)	(-6.39)	(-4.52)	(3.83)	(-0.49)	(-0.62)	(0.04)
Lagged DV	0.265***	0.638***	0.333***	0.579***	0.668***		
	(5.75)	(86.02)	(18.95)	(29.39)	(26.28)		
Log assets	-0.067**	-1.782***	1.180***	-0.167	-0.335***	0.435***	0.359***
	(-1.98)	(-7.60)	(7.00)	(-0.95)	(-3.30)	(3.89)	(4.94)
Log B/M	-1.036***	-2.847***	-2.205***	1.153***	0.581***	0.503***	0.232***
	(-27.51)	(-12.01)	(-15.01)	(10.62)	(9.93)	(4.92)	(3.98)
ROA	0.001	-0.012**	0.038***	0.003	0.010***	-0.031***	-0.010***
	(0.32)	(-2.35)	(9.12)	(0.60)	(5.11)	(-9.10)	(-6.82)
Leverage	-2.060***	2.490***	-0.611	0.802	0.269	-1.459***	-1.549***
	(-14.88)	(2.59)	(-1.09)	(1.57)	(0.88)	(-3.00)	(-4.60)
Log PPE/Assets	-0.038	-1.642***	-0.195	0.209	-0.095	0.282	-0.039
	(-0.82)	(-4.40)	(-0.71)	(0.78)	(-0.55)	(1.48)	(-0.32)
Firm FE	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓
Adj. R^2	0.445	0.547	0.752	0.912	0.932	0.040	0.073
Observations	62,962	56,016	64,931	64,981	65,589	57,416	57,907

Table 6: Analyst-implied returns and biodiversity footprint

This table reports panel regressions of analyst-implied 12-month returns on standardized BFI 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. BFI 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 price target dispersion (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. Column (5) interacts BFI with within-vintage size terciles in a saturated specification where every covariate, industry effect, and year effect is also interacted with the size tercile. Column (5) reports the marginal effect of BFI within each size tercile, −0.104, −0.033, and −0.010 for small, mid, and large firms respectively, and the cross-tercile equality of slopes is rejected at the 10% level. The sample is restricted to firms with non-missing BFI 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)	(5)
BFI (std.)	-0.027*	-0.047***	-0.046***	-0.045***	
	(0.014)	(0.015)	(0.016)	(0.017)	
BFI (Small)					-0.104***
					(0.029)
BFI (Mid)					-0.033*
					(0.017)
BFI (Large)					-0.010
					(0.032)
Releases (std.)				0.014	0.012
				(0.011)	(0.017)
ln(ME)	-0.039***	-0.050***	-0.044***	-0.044***	-0.188***
	(0.011)	(0.011)	(0.012)	(0.012)	(0.020)
ln(B/M)	0.094***	0.089***	0.109***	0.092***	0.052
	(0.018)	(0.018)	(0.020)	(0.020)	(0.032)
ln(# Analysts)	0.051**	0.034	0.024	0.015	0.047*
	(0.021)	(0.021)	(0.021)	(0.022)	(0.024)
PT Dispersion	0.403***	0.532***	0.483***	0.518***	0.328***
	(0.105)	(0.110)	(0.110)	(0.115)	(0.125)
Industry FE	×	×	✓	✓	✓
Year FE	×	✓	✓	✓	✓
Adj. R^2	0.052	0.099	0.109	0.107	0.141
N	115,219	115,219	115,219	108,099	108,099
Firms	1,069	1,069	1,069	1,012	1,012

Table 7: Earnings surprises and announcement returns

Panel A regresses earnings surprise measures on standardized BFI and controls using the quarterly earnings announcement sample. SUE1 is the seasonal random-walk surprise ($\text{EPS}_q - \text{EPS}_{q-4}$, scaled by rolling standard deviation), following Hou et al. (2020). SUE2 is the analyst-based surprise (actual − consensus mean, scaled by price) 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 BFI 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 BFI 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. LTG denotes the I/B/E/S long-term growth forecast. BFI 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: BFI predicts surprises				*Panel B: Announcement returns*			
	SUE1		SUE2		CAR $[-1, +1]$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
BFI (std.)	0.056***	0.063***	0.123***	0.085*	0.100**	0.063	0.040	0.059
	(0.012)	(0.012)	(0.042)	(0.044)	(0.043)	(0.043)	(0.044)	(0.056)
Releases (std.)	0.000	0.002	-0.105***	-0.062*				0.010
	(0.009)	(0.009)	(0.031)	(0.033)				(0.043)
SUE1						0.545***		0.118***
						(0.028)		(0.025)
SUE2							0.655***	0.608***
							(0.017)	(0.020)
LTG								0.021***
								(0.004)
ln(1+Analysts)	-0.068***	-0.067***	0.060	0.006				0.060
	(0.010)	(0.011)	(0.050)	(0.051)				(0.103)
ln(ME)	0.051***	0.051***	0.146***	0.180***	-0.025	-0.038*	-0.138***	-0.075
	(0.008)	(0.008)	(0.029)	(0.030)	(0.023)	(0.023)	(0.029)	(0.046)
ln(B/M)	-0.188***	-0.198***	-0.290***	-0.234***	-0.043	0.053	0.053	0.084
	(0.014)	(0.014)	(0.043)	(0.043)	(0.051)	(0.052)	(0.054)	(0.059)
Industry FE	×	✓	×	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Adj. R^2	0.038	0.039	0.048	0.052	0.002	0.017	0.104	0.096
N	61,807	61,799	49,196	49,196	66,392	65,823	51,911	26,824
Firms	1,586	1,585	1,358	1,358	1,697	1,677	1,434	1,157

Table 8: Earnings announcement realization and post-announcement drift

This table examines how much of the annual BFI premium is realized during earnings announcements and how the High−Low BFI spread evolves at extended post-announcement windows. For each size group and window, the High−Low spread is the difference in mean cumulative abnormal returns between the top and bottom BFI terciles across all announcements, with standard errors clustered by firm and announcement year-month; coefficient and t-statistic come from the same regression. Panel A divides the annualized announcement return (four times the $[-1,+1]$ High−Low spread) by the annual FF5+Mom portfolio alpha from Table 3 (small: $0.739\% \times 12 = 8.87\%$; mid: $0.397\% \times 12 = 4.76\%$; the large-firm alpha is negative, so no fraction is defined). Panel B reports the spread at the $[-1,+1]$, $[-1,+2]$, and $[-1,+3]$ announcement windows and the $[-1,+40]$ and $[-1,+60]$ post-announcement windows, for the full sample and by size tercile. CARs are market-adjusted. BFI terciles are formed within TRI vintage year; size terciles within calendar year using lagged market capitalization. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Panel A: EA Realization Fraction by Size Tercile

	T1 (Low)	T3 (High)	H−L	Annual α	Fraction
Small	−0.042	0.672	0.715**	8.87%	32.2%
			(2.54)		
Mid	0.240	0.376	0.136	4.76%	11.4%
			(1.00)		
Big	0.139	0.290	0.150	−1.04%	—
			(1.20)		

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

	$[-1,+1]$	$[-1,+2]$	$[-1,+3]$	$[-1,+40]$	$[-1,+60]$
Full sample	0.400***	0.481***	0.473***	1.405***	2.117***
	(4.39)	(4.92)	(4.55)	(4.83)	(5.86)
Small	0.715**	0.894***	0.820***	1.567**	2.049**
	(2.54)	(3.00)	(2.62)	(2.13)	(2.47)
Mid	0.136	0.158	0.165	−0.083	0.115
	(1.00)	(1.12)	(1.13)	(−0.28)	(0.32)
Big	0.150	0.179	0.229	0.107	0.353
	(1.20)	(1.29)	(1.59)	(0.43)	(0.99)
N (events)	46,046	46,044	46,041	45,680	45,376

Table 9: CDP biodiversity disclosure and continuous BFI

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). *BFI* (std) is ln(bird abundance/total assets), standardized cross-sectionally. *Releases* (std) is ln(TRI releases/total assets), standardized analogously. *Abatement* is an indicator equal to one if any of the firm's TRI facilities reported pollution prevention activities. *BFI* $\times$ *Abatement* is the product of standardized BFI 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. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	Near Bio	Indicators	Actions	Commitments	Vol. Disc.
BFI (std)	-0.045	-0.038	-0.111**	-0.074*	-0.099**
	(-0.97)	(-0.89)	(-2.44)	(-1.73)	(-2.49)
BFI $\times$ Abatement	0.053	0.096***	0.085**	0.043	0.063
	(1.39)	(2.85)	(2.19)	(0.99)	(1.61)
Abatement	-0.021	0.005	0.042	0.102**	0.043
	(-0.49)	(0.12)	(0.92)	(2.10)	(0.98)
Releases (std)	0.010	0.034	0.025	0.028	0.016
	(0.36)	(1.32)	(0.91)	(1.02)	(0.73)
ln(Assets)	0.022	0.055**	0.024	0.001	0.006
	(0.83)	(2.19)	(0.89)	(0.03)	(0.26)
ln(B/M)	-0.026	-0.027	-0.062*	-0.017	-0.062**
	(-0.74)	(-0.84)	(-1.70)	(-0.51)	(-2.13)
ROA	0.000	0.001	-0.002	0.000	-0.001
	(0.03)	(0.72)	(-1.57)	(0.17)	(-0.74)
ln(PPE/Assets)	0.102***	0.111***	0.156***	0.121***	0.148***
	(2.96)	(3.17)	(4.27)	(3.17)	(4.81)
Leverage	0.102	-0.009	-0.189	-0.154	-0.095
	(0.52)	(-0.06)	(-0.93)	(-0.94)	(-0.57)
N	595	804	797	864	942
Firms	263	285	286	297	307
R^2	0.206	0.222	0.232	0.164	0.156

Appendix

This appendix collects results that support the main analysis. Section A.1 decomposes BFI to rule out a mechanical size explanation. Section A.2 decomposes the return premium into its intensive and extensive margins, and Section A.3 applies the same decomposition to the cost penalty. Section A.4 shows that the regulatory pipeline responds to emissions rather than to ecological context. Ecological validation checks and other supporting evidence are in the Internet Appendix.

A.1 Alternative BFI constructions and the ecological numerator

One concern is that the return-predictive power of BFI derives from the asset-scaling denominator rather than the ecological numerator. BFI equals log(bird abundance/total assets), which decomposes additively into log(birds) − log(assets). Table A.2 addresses this by replacing BFI with its unscaled numerator, log(birds), in Fama-MacBeth regressions with full controls. With standard controls (Columns 1 and 2), log(birds) is statistically significant but economically weaker than BFI. Columns 3–4 add log(assets) as a separate control. log(birds) strengthens in Column 4, while log(assets) has a negative coefficient.

Three additional points rule out a mechanical size explanation. First, toxic releases intensity shares the same 1/assets denominator yet BFI retains its predictive power after controlling for releases in every specification (Table 4, columns 5–9). Second, the denominator captures exposure *intensity* rather than a disguised size effect. log(assets) is significant alongside log(birds) because it measures biodiversity richness per dollar of assets, not because smaller firms mechanically earn higher returns. Third, columns (6)–(7) of Table A.2 replace total assets with alternative denominators. BFI divided by either firm sales or plant, property, and equipment (PP&E) does not affect the significance of the return premium. 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.

Column (8) of Table A.2 shows that BFI survives a comprehensive set of size-related controls to the baseline specification: $\ln(\text{ME})^2$, Amihud illiquidity, share turnover, idiosyncratic volatility, and $\ln(1 + \text{Analysts})$. Column (9) adds facility-area controls drawn from EPA ECHO Facility Demographics (population density, minority share, and low-income share within a 5-mile radius) and from NLCD (log impervious surface area in the 15–30 km ring around facilities), each aggregated across a firm's TRI facilities. BFI remains significant after these controls, so the premium reflects the surrounding biodiversity of a firm's operations rather than the demographic or land-cover makeup of facility neighborhoods.

A.2 Decomposing the return premium: intensive versus extensive margin

Table A.3 decomposes the return premium using the identity $\text{BFI} = \ln(\bar{A}) + \ln(N) - \ln(\text{assets})$. The analysis distinguishes between abundance per facility, $\bar{A}$, and the number of facilities, N, relative to the firm's asset base.

Column (1) reports the baseline premium, while column (2) enters aggregate bird abundance and assets separately. Both predict returns, indicating that the premium does not arise solely from scaling by assets. Columns (3)–(5) then separate the numerator into its intensive and extensive margins. Facility count relative to assets predicts returns of 0.26% per month, whereas abundance per facility is indistinguishable from zero both separately and alongside facility intensity. The return premium therefore reflects the number of ecologically consequential facilities a firm operates relative to its asset base rather than the ecological richness of its average facility.

The extensive-margin result is consistent with the aggregate-externality structure in Pedersen (2026), in which externalities are summed across a firm's activities before being scaled by firm value. The difference-in-differences evidence in Section 3 shows that industrial facilities reduce local biodiversity on average, giving facility count a direct ecological interpretation. The importance of scaling appears in column (6), where unscaled facility count predicts lower returns when included alongside BFI, while the asset-scaled footprint remains positive and significant. BFI correlates 0.92 with facility count relative to assets, explaining why the extensive margin reproduces the footprint premium.

A.3 Decomposing the cost penalty: intensive versus extensive margin

BFI aggregates effort-adjusted bird abundance across all of a firm's facilities and scales by assets, so it decomposes additively into abundance per facility (the intensive margin), facility count (the extensive margin), and log assets: BFI $= \ln(\bar{A}) + \ln(N) - \ln(\text{assets})$. Table A.4 re-estimates the COGS/Sales regression of Table 5 across five specifications on a common sample. Column (1) is the baseline BFI specification. Column (2) replaces BFI with a single regressor for log total abundance while retaining log assets as a control, and recovers the same positive, significant cost effect. Columns (3) and (4) enter the two margins one at a time, and column (5) enters them together. The intensive margin is null in every specification, while the facility count carries the cost effect. The penalty therefore operates through the breadth of a firm's operating footprint rather than the biodiversity of any single site. The difference-in-differences evidence in Section 3 establishes that each facility causally degrades local biodiversity, so the facility count aggregates validated ecological impacts, and the asset scaling expresses this footprint per dollar of firm value rather than firm size.

A.4 Regulatory burden and ecological exposure

I next test whether facilities in ecologically richer locations attract more EPA regulatory attention, conditional on the toxic emissions they report. The sample is a facility-year panel of 13,731 TRI-reporting facilities and 236,325 facility-years over 1987–2024. I restrict the sample to facility-years with strictly positive total releases, so the log-releases regressor enters as ln(releases), mirroring the convention used throughout the firm-level Fama-MacBeth regressions in Section 4. Each facility-year is matched to the year's effort-adjusted log bird abundance in the 15–30 km ring around the facility, the same eBird-derived measure that enters BFI in the main analysis. Both regressors are time-varying at the facility-year level, and the analysis tests whether ecological richness predicts enforcement after partialling out year-by-year emissions.

Each pair of columns covers one stage of the regulatory pipeline in Table A.5. Columns 1–2 use the count of EPA inspections recorded for each facility-year in the EPA's Integrated Compliance Information System (ICIS-FE&C). Columns 3–4 use the count of informal enforcement actions such as warning letters and notices of violation. Columns 5–6 use an indicator for any formal penalty assessed in that fiscal year, scaled by 100 so coefficients read as percentage-point changes in the probability of penalty. Columns 7–8 use the dollar amount of penalties assessed. Within each pair, odd-numbered columns include only ecological richness; even-numbered columns add ln(releases) to isolate the conditional contribution of ecology beyond emissions.

Columns 1–4 and 7–8 estimate Poisson Pseudo-Maximum-Likelihood (PPML) regressions on raw counts and amounts, recovering proportional effects directly. Columns 5–6 estimate linear probability models on the binary penalty indicator. All specifications include EPA region, two-digit NAICS industry, and year fixed effects. Standard errors are clustered at the facility level. PPML drops fixed-effect groups in which the outcome is zero throughout. The reduction in the penalty-amount sample to roughly 78,000 observations reflects this mechanical drop: monetary penalties are rare, so several region, NAICS-2, and year groups contain no positive penalty and are excluded from the PPML sample.

Bird abundance has no predictive power for any of the four outcomes once emissions are controlled. The coefficient on biodiversity richness is economically small and statistically insignificant in every column. By contrast, ln(releases) is highly significant for all four outcomes. The result is robust to tighter fixed-effect structures which is not shown. In untabulated tests, replacing the additive region and year fixed effects with state-by-year and industry-by-year interactions, or adding facility fixed effects to identify off purely within-facility variation, leaves the ecological-richness coefficient economically negligible while ln(releases) remains significant. The regulatory apparatus responds to what firms emit, not to the ecological context in that they emit. This null reinforces the central mechanism. Because regulators do not single out ecologically rich locations, the higher costs of operating there surface as ordinary operating expense, indistinguishable to intermediaries from any other cost.

Appendix Tables

Table A.1: Variable definitions

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

Variable	Definition
Panel A: Ecological exposure variables	
BFI	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.
BFI$_{sales}$	ln(Bird Abundance/trailing twelve-month Sales), alternative scaling that replaces total assets with sales. Source: eBird, TRI, Compustat.
BFI$_{PP\&E}$	ln(Bird Abundance/net PP&E), alternative scaling that replaces total assets with net property, plant, and equipment. Source: eBird, TRI, Compustat.
ln(bird abundance)	Natural log of effort-adjusted bird abundance in the 15–30 km ring, before asset scaling. The ecological numerator of BFI. Source: eBird, TRI.
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.
Number of TRI facilities	Active TRI facilities operated by the firm per year. Source: 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 total assets	ln(ATQ), in millions. Source: Compustat.
Log book-to-market	ln(book equity/market equity). Source: Compustat, CRSP.

Continued on next page

Table A.1 continued from previous page

Variable	Definition
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.
Tobin's Q	(Total Assets − Common Equity + Market Equity)/Total Assets, computed monthly: (ATQ − CEQQ + |prc| × shrout/1,000)/ATQ, with the most recent quarterly accounting values carried forward and market equity updated each month. Source: Compustat, CRSP.
Investment rate (I/K)	Capital expenditures / lagged net PP&E. Source: Compustat.
Log tangibility	ln(net PP&E/total assets). Source: Compustat.
Asset age	1 − (net PP&E/gross PP&E), the depreciation fraction of fixed capital. Higher values indicate older fixed assets. Source: Compustat.
Idiosyncratic volatility	SD of daily returns in excess of the market, within each month. 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(WC/TA) + 3.3(EBIT/TA) + 0.6(MVE/TL) + 1.0(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.
Long-term growth forecast	Median analyst long-term EPS growth forecast (%). 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.
Short interest ratio (SIR)	Shares short / shares outstanding, monthly. Source: Compustat.
Amihud illiquidity	Monthly average of daily $\|r_d\|/\text{DolVol}_d$ ($\times 10^6$), following Amihud (2002). Source: CRSP.
Turnover	Monthly volume / shares outstanding (%). Source: CRSP.
MSCI environmental score	Industry-adjusted environmental pillar score. Source: MSCI ESG Ratings.
MSCI biodiversity sub-scores	Sub-scores comprising biodiversity exposure (BiodivExp), biodiversity management (BiodivMgmt), and biodiversity weight (BiodivWt) within MSCI's environmental pillar. 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: Governance	
Board size	Number of directors. Source: BoardEx.
Board independence	Fraction of non-executive directors. Source: BoardEx.
Director network size	Average director NetworkSize (log-transformed in regressions). Source: BoardEx.

Table A.2: Fama-MacBeth tests of alternative BFI constructions

This table reports Fama-MacBeth (1973) cross-sectional regressions of monthly stock returns on the indicated biodiversity measure and lagged controls. The table tests whether the return-predictive power of BFI depends on the choice of scaling denominator and whether it survives comprehensive size-related and facility-location controls. The sample is monthly from October 1988 to December 2024. Each column reports time-series averages of monthly coefficient estimates with Newey-West (1987) standard errors (six lags). Columns (1)–(4) replace BFI 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 BFI ≡ ln(birds/assets) specification on this sample for comparison. Column (6) replaces total assets with trailing twelve-month sales: $\text{BFI}_{sales} \equiv$ ln(birds/sales). Column (7) uses net PP&E: $\text{BFI}_{PP\&E} \equiv$ ln(birds/PP&E). Column (8) adds the full battery of size-related controls to the baseline specification: $\ln(\text{ME})^2$, Amihud illiquidity, share turnover, idiosyncratic volatility, and ln(1+Analysts). Column (9) adds facility-area location controls to the column (8) specification: log population density, minority share, and low-income share within a 5-mile radius from EPA ECHO Facility Demographics, and log built-up surface area in the 15–30 km ring around facilities from NLCD. Location controls are aggregated across each firm's TRI facilities. 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)	(9)
ln(birds) (std.)	0.044*	0.048*	0.075***	0.084***					
	(1.74)	(1.89)	(2.88)	(3.22)					
ln(assets)			-0.767***	-0.862***					
			(-4.46)	(-5.34)					
BFI_*assets* (std.)					0.184***			0.214***	0.201***
					(4.40)			(4.38)	(4.14)
BFI_*sales* (std.)						0.162***			
						(4.08)			
BFI_*PP&E* (std.)							0.164***		
							(3.56)		
Log book-to-market	✓	✓	✓	✓	✓	✓	✓	✓	✓
Momentum	✓	✓	✓	✓	✓	✓	✓	✓	✓
Log ME	✓	✓	✓	✓	✓	✓	✓	✓	✓
$(\text{Log ME})^2$	×	×	×	×	×	×	×	✓	✓
ROA, Leverage	✓	✓	✓	✓	✓	✓	✓	✓	✓
Idiosyncratic volatility	×	✓	×	✓	×	×	×	✓	✓
Amihud, Turnover	×	×	×	×	×	×	×	✓	✓
ln(1 + Analysts)	×	×	×	×	×	×	×	✓	✓
Facility-area controls	×	×	×	×	×	×	×	×	✓
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
N months	428	428	428	428	428	428	428	428	425
Avg. firms/month	436	436	436	436	436	436	436	403	389

Table A.3: Decomposing the return premium: intensive versus extensive margin

Monthly Fama-MacBeth regressions of excess returns on the biodiversity footprint and its components. The index decomposes as BFI = ln(birds) − ln(assets), and its numerator as ln(birds) = $\ln(N)+\ln(\bar{A})$, where N is a firm's facility count and $\bar{A}$ is bird abundance per facility. Column (1) is the baseline premium. Column (2) enters the numerator and denominator separately. Columns (3) and (4) enter the extensive margin, $\ln(N/\text{Assets})$, and the intensive margin, $\ln(\bar{A})$, one at a time, and column (5) enters both. Column (6) adds the unscaled facility count $\ln(N)$ alongside BFI. BFI correlates 0.92 with $\ln(N/\text{Assets})$, while $\ln(N/\text{Assets})$ and $\ln(\bar{A})$ are orthogonal (−0.01). All regressors are lagged one month, winsorized at the 1st and 99th percentiles, and standardized within month. Every column controls for log book-to-market, momentum (cumulative return from $t-12$ to $t-2$), log market equity, return on assets, and leverage, and includes Fama-French 12-industry fixed effects. Newey-West(6) t-statistics in parentheses. Coefficients are in percent per month. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)
BFI	0.185***					0.438***
	(4.06)					(4.10)
ln(Birds)		0.072**				
		(2.38)				
ln(Assets)		-1.379***				
		(-4.43)				
$\ln(N/\text{Assets})$ (extensive)			0.257***		0.256***	
			(4.29)		(4.19)	
$\ln(\bar{A})$ (intensive)				0.032	0.016	
				(1.09)	(0.52)	
$\ln(N)$, unscaled						-0.247***
						(-3.26)
Controls	✓	✓	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓	✓	✓
N months	428	428	428	428	428	428
Avg. firms/month	404	404	404	404	404	404

Table A.4: Decomposing the BFI cost penalty: intensive versus extensive margins

This Internet Appendix table decomposes the BFI cost penalty into an intensive margin, effort-adjusted bird abundance per facility ($\bar{A}$), and an extensive margin—the number of a firm's TRI facilities (N). The dependent variable is COGS/Sales in percent, and the sample is the quarterly panel of TRI-reporting firms. Since BFI equals log total abundance minus log assets, it splits into $\ln \bar{A}$ and $\ln N$ once assets are held in the control set. Column (1) is the baseline BFI specification and reproduces the COGS/Sales result of Table 5; column (2) replaces BFI with log total abundance; column (3) enters abundance per facility ($\ln \bar{A}$) alone and column (4) the facility count ($\ln N$) alone; and column (5) enters both together. The cost runs through the extensive margin: $\ln N$ is positive and significant, while $\ln \bar{A}$ is indistinguishable from zero, so the penalty reflects the breadth of a firm's operating footprint rather than the biodiversity of any single site. All columns share one common sample and include firm and year fixed effects, the one-quarter lagged dependent variable, and the control set (log assets, log book-to-market, ROA, leverage, and log PPE/Assets, each lagged one quarter). The biodiversity-margin regressors are standardized within TRI vintage, so their coefficients are percentage points of COGS/Sales per standard deviation; controls are in raw units. t-statistics from standard errors clustered by firm are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1)	(2)	(3)	(4)	(5)
BFI	0.353***				
	(3.83)				
Log total abundance		0.199***			
		(2.63)			
Abundance/facility $\ln \bar{A}$			-0.038		-0.054
			(-0.92)		(-1.31)
Facility count $\ln N$				0.425***	0.431***
				(4.00)	(4.04)
Lagged COGS/Sales	0.579***	0.580***	0.580***	0.578***	0.578***
	(29.39)	(29.41)	(29.36)	(29.57)	(29.59)
Log assets	-0.167	-0.349**	-0.301*	-0.416**	-0.414**
	(-0.95)	(-2.01)	(-1.72)	(-2.42)	(-2.40)
Log B/M	1.153***	1.159***	1.163***	1.160***	1.160***
	(10.62)	(10.67)	(10.69)	(10.69)	(10.70)
ROA	0.003	0.003	0.003	0.002	0.002
	(0.60)	(0.59)	(0.62)	(0.55)	(0.55)
Leverage	0.802	0.795	0.750	0.836*	0.825
	(1.57)	(1.56)	(1.47)	(1.66)	(1.63)
Log PPE/Assets	0.209	0.216	0.239	0.194	0.197
	(0.78)	(0.80)	(0.88)	(0.72)	(0.74)
Firm FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓
Adj. R^2	0.912	0.912	0.912	0.912	0.912
Observations	64,981	64,981	64,981	64,981	64,981

Table A.5: Regulatory burden and ecological exposure

This table tests whether facilities in ecologically richer areas face more EPA regulatory attention, conditional on toxic emissions. The sample is a facility-year panel of TRI-reporting facilities from 1987 to 2024, restricted to facility-years with positive total releases and non-missing 15–30 km control-ring bird abundance from eBird. The dependent variables span the 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 total penalty amount in dollars (columns 7–8). Columns 1–4 and 7–8 estimate Poisson Pseudo-Maximum-Likelihood (PPML) regressions; PPML coefficients are semi-elasticities (proportional change in expected outcome per unit of regressor). Columns 5–6 estimate linear probability models with the binary penalty outcome scaled by 100, so coefficients are in percentage points. Ecological richness is the effort-adjusted log bird abundance in the 15–30 km control ring for the facility-year, the same measure used to construct BFI. ln(releases) is the natural log of total toxic releases reported to the TRI for the facility-year. All specifications include EPA region, two-digit NAICS, and year fixed effects. PPML drops fixed-effect groups in which the outcome is zero throughout. The reduction in N for columns 7–8 reflects this mechanical drop: monetary penalties are rare in the sample, so several region, NAICS-2, and year groups contain no positive penalty and are excluded from the PPML sample. Standard errors clustered by facility appear in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

	Inspections (PPML)		Informal actions (PPML)		Pr(penalty) (LPM)		Penalty amount (PPML)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ecological richness	0.010	0.012	-0.028	-0.031	0.018	0.017	-0.061	-0.076
	(0.020)	(0.019)	(0.043)	(0.042)	(0.028)	(0.028)	(0.121)	(0.107)
ln(releases)		0.075***		0.043***		0.063***		0.455***
		(0.004)		(0.008)		(0.005)		(0.072)
EPA Region FE	✓	✓	✓	✓	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
N	206,991	206,991	182,241	182,241	236,324	236,324	78,393	78,393

Internet Appendix

Table IA.1: Variable definitions: Internet Appendix variables

Additional variables used in Internet Appendix robustness tests and facility-level analyses. Variables defined in the main paper appear in Table A.1 and are not repeated here.

Variable	Definition
Panel A: Location and channel controls	
SEPA share	Fraction of a firm's TRI facilities located in states with a mini-NEPA (State Environmental Policy Act) statute that mandates environmental review of state actions affecting protected resources. State assignment from the facility ZIP code. Source: TRI, state environmental policy statutes.
Offsite waste intensity	Total off-site waste transfers (sum of POTW, off-site disposal, off-site recycling, and off-site treatment, in pounds) divided by total assets, winsorized at the 1st and 99th percentiles by year. Source: TRI Section 6/8, Compustat.
CERCLA share	Share of a firm's TRI releases subject to Comprehensive Environmental Response, Compensation, and Liability Act (CERCLA) cleanup liability, computed as (water + underground injection + landfills + on-site surface disposal) / (CERCLA-relevant + air releases). Air pollution is excluded because it is not covered by CERCLA (Bellon, 2025). Source: TRI.
NLCD impervious surface	Log of built-up (impervious) surface area in the 15–30 km control ring around facilities, aggregated across a firm's TRI facilities. Source: National Land Cover Database.
Demographic controls	Log population density, minority share, and low-income share within a 5-mile radius of each facility, aggregated across a firm's TRI facilities. Source: EPA ECHO Facility Demographics.
Panel B: Regulatory outcomes	
EPA inspections	Count of EPA inspections of a facility in a given year. Source: EPA ECHO ICIS.
EPA informal enforcement actions	Count of informal EPA enforcement actions (e.g., notices of violation, warning letters) against a facility in a given year. Source: EPA ECHO.
Penalty indicator	Equal to one if a facility receives any formal monetary penalty in a given year. Source: EPA ECHO.
Penalty amount	Total dollars of formal monetary penalties assessed against a facility in a given year. Source: EPA ECHO.
Panel C: Production technology	
Log labor-to-capital ratio	Natural log of Compustat annual employees divided by net property, plant, and equipment (ln(emp/ppent)). Captures the production-function input mix. Source: Compustat.
Labor share of value added	Imputed annual staff expense divided by the sum of imputed staff expense, operating income after depreciation, and depreciation. Imputed staff expense is Compustat employees times the BLS Quarterly Census of Employment and Wages state by 6-digit NAICS average weekly wage, multiplied by 52. Follows Donangelo (2014). Source: Compustat, BLS QCEW.
Variable cost share	COGS divided by the sum of COGS, selling and administrative expense, and depreciation, all annualized. Follows Novy-Marx (2011). Source: Compustat.

Table IA.2: Summary statistics: Internet Appendix variables

Descriptive statistics for variables used in the Internet Appendix analyses. *Tobin's Q* is the market value of assets divided by book assets, entered as a control in the RepRisk news regression. *SEPA share* is the firm-year fraction of TRI facilities in mini-NEPA states. *Offsite waste intensity* is total off-site TRI transfers (POTW, off-site disposal, off-site recycling, and off-site treatment, in pounds) divided by total assets, winsorized at the 1st and 99th percentiles by year. *CERCLA share* is the firm-year share of TRI releases subject to CERCLA cleanup liability (water, underground injection, landfills, and on-site surface disposal); air pollution is excluded because it is not covered by CERCLA (Bellon, 2025). The demographic controls (log population density, minority share, low-income share) are averages within a 5-mile radius of each firm's TRI facilities, drawn from EPA ECHO. All continuous variables are winsorized at the 1st and 99th percentiles.

	Mean	SD	p5	p25	Median	p75	p95	*N*
Tobin's Q	1.68	0.86	0.90	1.15	1.42	1.92	3.40	18,672
SEPA share	0.39	0.37	0.00	0.00	0.33	0.67	1.00	41,529
Offsite waste intensity	1001.11	3140.95	0.00	4.58	66.83	463.36	4893.60	18,685
CERCLA share	0.08	0.22	0.00	0.00	0.00	0.02	0.68	38,857
Log population density	6.68	1.08	4.73	6.01	6.76	7.43	8.32	18,672
Minority share	0.38	0.16	0.13	0.27	0.37	0.48	0.69	18,672
Low-income share	0.31	0.07	0.18	0.27	0.31	0.35	0.43	18,672

Section A: Ecological Validation

IA.1 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 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.

Effort-adjusted abundance is a first-step residual used as the dependent variable in subsequent regressions, so I include the first-step regressors in those specifications to avoid the coefficient bias 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 2 shows that the core results are robust regardless of whether these controls are included (comparing columns 1 and 4).

IA.2 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 IA.1 validates this procedure across all facilities. By construction, average impervious surface growth peaks at the detected year (year 0); it is near zero in the surrounding years, so built-up area jumps in a single year rather than accumulating gradually, consistent with a discrete construction event.

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 IA.2 shows that 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.

Figure IA.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. Year 0 is the largest jump by construction, so the informative pattern is the near-zero growth in surrounding years, consistent with a discrete 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.

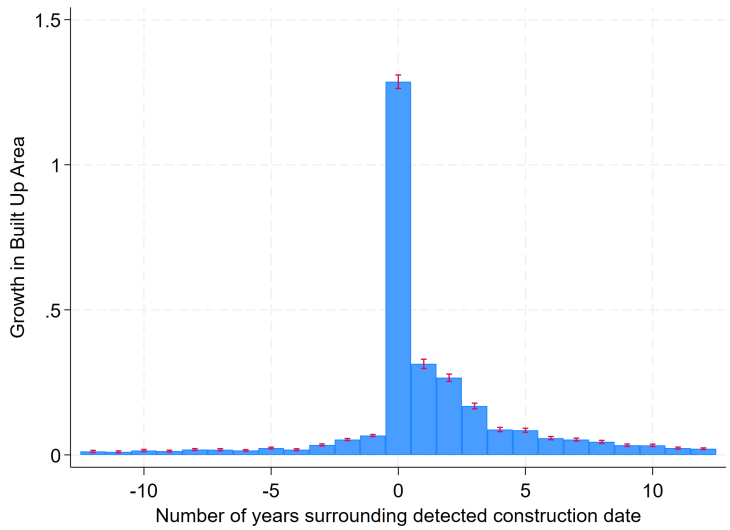

Figure IA.2: Spatial attenuation of ecological decline

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 tested ring shows a significant positive effect, so the near-field losses are not offset by gains in adjacent rings. $N = 3{,}623$ facilities with construction year ≥ 1990; event window $[-6, +5]$ years; clustering by facility.

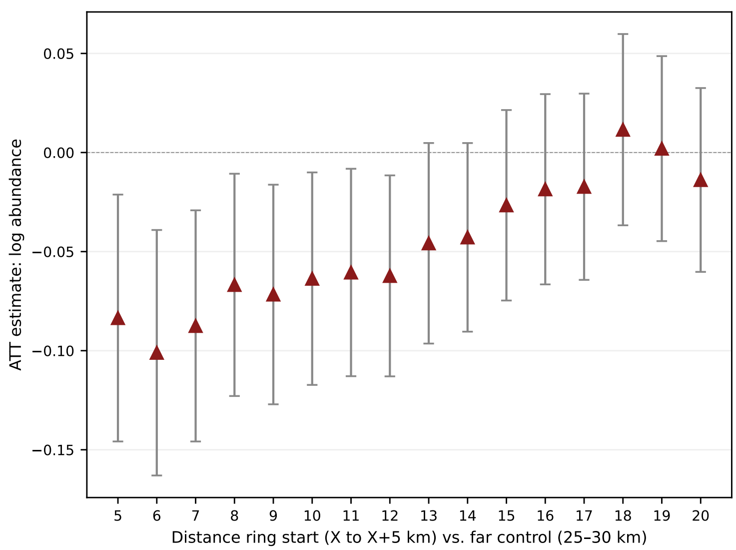

IA.3 Within-facility and pathway evidence

Three further pieces of evidence support the ecological-validation strategy. First, a within-facility dose-response test shows that years with higher toxic releases are associated with lower bird abundance at the same site (IA Figure IA.3). Second, ambient water-quality monitoring shows elevated lead concentrations near TRI facilities (IA Figure IA.4). Third, randomization inference confirms that the lead effects are correctly timed rather than artifacts of generic trends (IA Figure IA.5).

Figure IA.3: 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). The fitted line is the within-facility regression coefficient applied to the bin means; the vertical axis is zoomed to the range of conditional means, which amplifies the visual slope relative to the unconditional scale of log abundance. Ecological degradation scales with the intensity of industrial activity within a given facility over time.

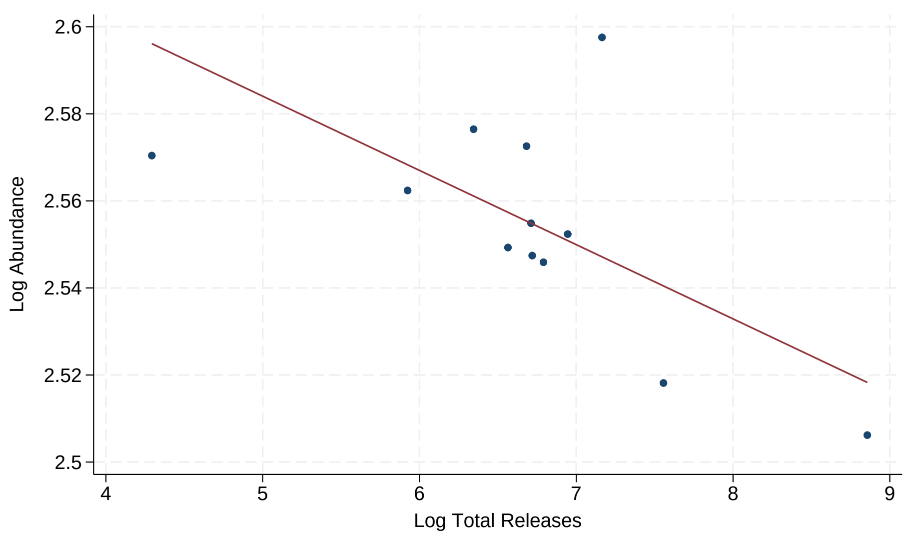

Lead contamination pathway. I use ambient water-quality monitoring data to test whether TRI facility construction raises local lead concentrations. I match each TRI facility to nearby monitoring stations, restrict to routine water samples, harmonize units to μg/L, and summarize lead within each station-year using the median to reduce sensitivity to extreme readings. I then implement the same ring-based staggered difference-in-differences design used for bird abundance in equation (3), replacing the outcome with median lead concentration (μg/L) and varying the treatment threshold X from 5 to 24 kilometers.

IA Figure IA.4 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.4: 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.

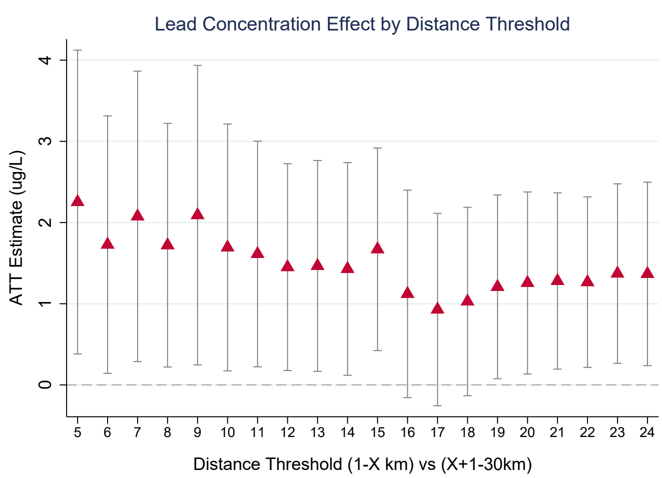

Lead placebo tests. I apply the same randomization inference used for the biodiversity placebos, keeping the spatial design fixed and permuting only facility construction years. Each of the 1,000 iterations randomly reassigns construction timing across facilities and re-estimates the staggered DiD under 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.5 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 footprint index captures.

Figure IA.5: 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

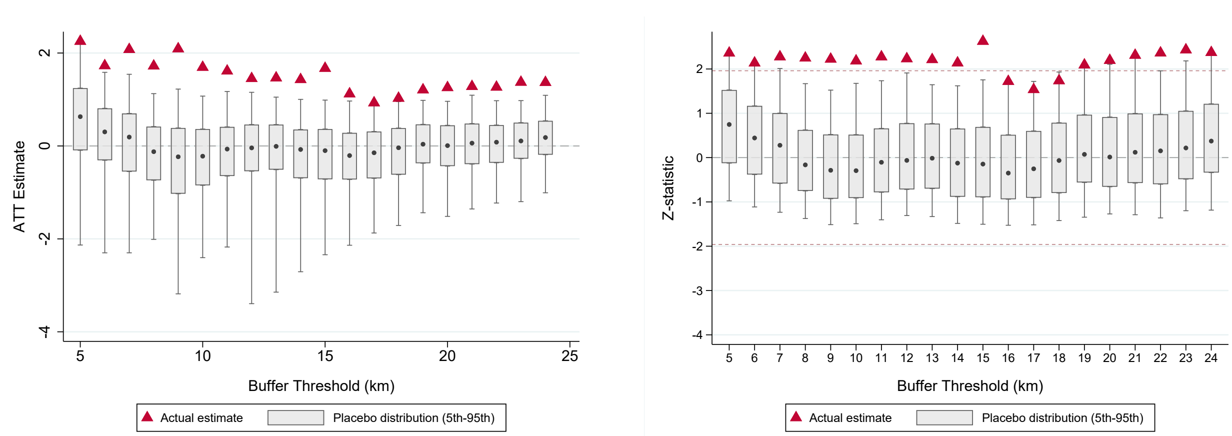

Section B: Robustness of the Biodiversity Premium

IA.4 BFI and toxic releases in Fama-MacBeth regressions

Table 4 in the main paper already estimates BFI and toxic releases jointly across nine specifications. Table IA.3 extends that joint specification with size and idiosyncratic-volatility interactions to test whether the BFI premium operates through the same channels as a pollution premium would. When both variables enter together (column 3), BFI retains a significant coefficient, indicating that the ecological dimension adds incremental pricing information beyond emissions intensity. Column 4 adds interactions with log market equity; the BFI and Releases size-interaction coefficients are each indistinguishable from zero. Column 5 adds idiosyncratic volatility interactions, and BFI carries incremental return-predictive power beyond releases in all specifications.

Table IA.3: BFI and 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. Newey-West(6) t-statistics in parentheses. Coefficients are in percent per month. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1)	(2)	(3)	(4)	(5)
BFI	0.279***		0.255***	0.241***	0.261***
	(6.06)		(5.20)	(4.28)	(5.08)
Releases		0.115***	0.045	0.052	0.060*
		(3.46)	(1.40)	(1.55)	(1.86)
Log ME				-0.015	0.138*
				(-0.16)	(1.84)
Idiovol					0.247**
					(2.52)
Momentum	-0.356	-0.319	-0.356	-0.402	-0.318
	(-1.37)	(-1.21)	(-1.34)	(-1.54)	(-1.31)
BFI × Size				-0.009	0.017
				(-0.23)	(0.39)
Rel × Size				-0.001	0.053
				(-0.02)	(1.39)
BFI × Idiovol					-0.001
					(-0.02)
Rel × Idiovol					0.106*
					(1.75)
N	186,528	174,556	174,556	174,556	174,555
Avg. R^2	0.11	0.11	0.12	0.14	0.16

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

Table IA.4 reports high-minus-low BFI spread alphas from a 10×5 conditional double sort. Each vintage year, firms are assigned to ten size deciles using unconditional size breakpoints computed across all sample firms, then to five BFI quintiles within each decile. Under FF3 controls, the spread is positive and significant through the ninth decile, and it is significant at the sixth decile under all six factor models. Only the largest decile is indistinguishable from zero across every specification, while the ninth remains significant under the CAPM and three-factor models. 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.4: Size × BFI double sort: unconditional breakpoints (10 × 5)

Each vintage year, firms are sorted into ten size deciles by market equity, then into five BFI quintiles within each decile. Size breakpoints are unconditional, computed across all sample firms. The table reports the high-minus-low BFI 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
D1 (Small)	1.639***	1.711***	1.644***	1.474***	1.413***	1.631***
	(3.97)	(4.40)	(4.05)	(3.50)	(3.46)	(3.91)
D2	1.278***	1.280***	1.033***	0.952***	0.851***	1.084***
	(4.34)	(4.36)	(3.58)	(2.90)	(2.92)	(3.77)
D3	1.204***	1.267***	1.231***	1.297***	1.115***	1.243***
	(4.63)	(4.68)	(4.43)	(4.59)	(3.77)	(4.37)
D4	0.668***	0.710***	0.832***	0.844***	0.844***	0.853***
	(2.73)	(2.81)	(3.33)	(3.28)	(3.31)	(3.37)
D5	0.748***	0.733***	0.797***	0.818***	0.660***	0.820***
	(3.37)	(3.19)	(3.23)	(3.13)	(2.89)	(3.18)
D6	0.714***	0.714***	0.558***	0.605***	0.473**	0.554***
	(3.69)	(3.79)	(2.80)	(2.81)	(2.21)	(2.73)
D7	0.564***	0.564***	0.501**	0.527***	0.410**	0.492**
	(2.76)	(2.77)	(2.56)	(2.71)	(2.08)	(2.48)
D8	0.411**	0.409**	0.211	0.248	0.062	0.223
	(2.04)	(2.10)	(1.10)	(1.14)	(0.27)	(1.15)
D9	0.321*	0.310*	0.184	0.244	0.151	0.151
	(1.78)	(1.71)	(1.16)	(1.52)	(0.92)	(0.90)
D10 (Big)	0.066	0.021	-0.158	-0.085	-0.157	-0.175
	(0.33)	(0.12)	(-0.82)	(-0.42)	(-0.81)	(-0.91)
N months	434	434	434	434	434	434

IA.6 Industry × size × BFI triple sort

Table IA.5 controls for industry composition within size bins using a triple sort on Fama-French 12 industries, market equity terciles within industry, and BFI terciles within each industry-size cell. The small-firm BFI 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 concentration in small firms documented in Table 3.

Table IA.5: Industry × size × BFI 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 BFI spread portfolios from a triple sort on Fama-French 12 industry, market equity terciles within industry, and BFI terciles within each industry–size cell. Each vintage year, firms are first allocated to FF12 industries and size terciles within industry, then sorted into BFI terciles within each of the 36 industry–size bins. Value-weighted portfolios are formed by pooling across industries within each size–BFI cell. Each cell reports the alpha of the high-BFI minus low-BFI 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) FF3	(2) FF5	(3) FF5+BAB	(4) FF5+Mom	(5) FF5+Liq
H–L BFI α within industry–size cell (%/month):					
Small	0.816***	0.882***	0.906***	0.858***	0.918***
	(3.72)	(4.54)	(4.59)	(4.30)	(4.64)
Mid	0.541***	0.431***	0.437***	0.349***	0.437***
	(4.86)	(3.88)	(3.76)	(3.08)	(3.92)
Big	0.016	-0.116	-0.125	-0.132	-0.110
	(0.16)	(-1.07)	(-1.21)	(-1.26)	(-1.02)
Pooled	0.458***	0.399***	0.406***	0.358***	0.415***
	(5.41)	(5.44)	(5.26)	(4.54)	(5.59)
N months	434	434	432	434	434

IA.7 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 BFI 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.6 reports Fama-MacBeth regressions. BFI enters with a coefficient of 0.182% per month, statistically significant, in the univariate specification and remains significant in the joint 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 4.

Panel B reports double-sort alphas across five factor models. The small-firm alpha ranges from 85 to 94 basis points per month and is significant at the 1% level under every model. The mid-cap alpha ranges from 44 to 52 basis points per month and is significant at the 1% level under every model. The large-cap alpha is statistically zero under every model. The concentration in small firms documented in Section 4 is, if anything, sharper in the post-2002 period.

Table IA.6: Post-2002 eBird subsample robustness

The sample is restricted to observations with tri_vintage $\geq$ 2002 (October 2002 through December 2024; 267 months, 416 firms per month on average). *Panel A*: Monthly Fama-MacBeth regressions of excess returns on lagged firm characteristics. BFI 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 BFI spreads within size terciles (size terciles, then BFI terciles within each size group, by vintage year; value-weighted). Alphas in basis points per month from Newey-West(12) regressions of the H-L spread on each of five factor models, with standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Panel A: Post-2002 Fama-MacBeth regressions

	(1)	(2)	(3)	(4)
BFI (standardized)	0.182***	0.259***	0.167***	0.194***
	(2.78)	(3.82)	(2.62)	(3.29)
Releases (standardized)			0.037	0.033
			(0.83)	(0.80)
Log book-to-market		-0.862***	-0.652***	-0.661***
		(-7.24)	(-5.63)	(-5.78)
Log market equity		-0.039	-0.125**	-0.014
		(-0.64)	(-2.34)	(-0.35)
ROA			0.031***	0.032***
			(7.37)	(8.37)
Leverage			-0.010	-0.254
			(-0.02)	(-0.63)
Idiovol				28.940***
				(3.21)
Investment rate				-0.891***
				(-3.00)
Industry FE	Yes	Yes	Yes	Yes
N months	267	267	267	267
N observations	110,994	110,994	110,994	110,994

Panel B: Size × BFI H−L spread α (bps/month)

	FF3	FF5	FF5+BAB	FF5+Mom	FF5+Liq
Small	94.3***	88.0***	89.0***	84.9***	88.7***
	(15.4)	(16.1)	(17.1)	(16.6)	(16.5)
Mid	47.3***	45.2***	52.0***	43.7***	45.8***
	(13.2)	(14.5)	(14.6)	(14.4)	(14.5)
Large	4.0	0.8	4.6	0.8	0.3
	(13.9)	(13.4)	(14.2)	(13.2)	(13.2)

IA.8 Size, liquidity, and idiosyncratic volatility interactions

Table IA.7 asks whether any standard limits-to-arbitrage proxy absorbs the BFI premium, interacting BFI with one proxy at a time across three Fama-MacBeth specifications. The baseline BFI coefficient is economically and statistically significant. The interactions with log market equity and idiosyncratic volatility are both insignificant (Columns 1 and 2); the premium's concentration in small firms appears in the portfolio sorts (Tables 3 and IA.4) rather than in the continuous interaction. Using the short interest ratio (SIR) provides a different test. If the premium reflected limits to arbitrage from short-selling frictions, it should be larger among high-SIR stocks where short positions are already stretched and further shorting is difficult. Column 3 instead shows that the BFI $\times$ SIR interaction is negative, so the premium is weaker rather than stronger among heavily shorted stocks, inconsistent with short-selling frictions. BFI retains its predictive power across all three specifications, confirming that the ecological attribute contains independent pricing information. Each specification controls for book-to-market and momentum, and includes Fama-French 12-industry fixed effects.

Table IA.7: Limits to arbitrage: size, idiosyncratic volatility, and short interest interactions

Monthly Fama-MacBeth regressions of excess returns on BFI and its interaction with one limits-to-arbitrage proxy at a time: firm size (log market equity, column 1), idiosyncratic volatility (column 2), and short interest (the ratio of shares sold short to shares outstanding, column 3). Estimating one interaction per column avoids collinearity among the proxies. BFI and the moderators are cross-sectionally standardized each month, and interaction terms are products of standardized constituents. Log market equity is included in every column (interacted with BFI in column 1); log book-to-market and momentum (cumulative returns $t-12$ to $t-2$) enter as levels, and all columns include Fama-French 12-industry fixed effects. Column 3 is restricted to firm-months with short-interest data. Newey-West(6) t-statistics in parentheses. Coefficients in percent per month. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	Size	Idio. vol.	Short int.
BFI (std)	0.258***	0.276***	0.222***
	(5.24)	(5.95)	(4.34)
BFI × Log ME	-0.005		
	(-0.13)		
BFI × Idiosyncratic vol.		0.013	
		(0.21)	
BFI × Short interest			-0.118**
			(-2.40)
Log ME	-0.035	0.096	-0.100
	(-0.39)	(1.32)	(-1.05)
Idiosyncratic vol.		0.289**	
		(2.52)	
Short interest			0.041
			(0.62)
Log book-to-market	-0.960***	-0.951***	-0.904***
	(-9.65)	(-9.69)	(-9.32)
Industry FE	✓	✓	✓
N months	428	428	384
Avg. R^2	0.12	0.14	0.13

Section C: Mechanisms of Information Failure

IA.9 Media coverage and biodiversity exposure

Table IA.8 examines whether the biodiversity footprint draws media coverage. I regress the annual count of biodiversity-related news incidents from RepRisk on lagged BFI, lagged toxic releases intensity, and controls, with firm and year fixed effects, estimating the count by Poisson pseudo-maximum-likelihood and the incidence indicator by a linear probability model. The BFI coefficient is small and statistically indistinguishable from zero in every column. Toxic releases intensity, by contrast, predicts more coverage, and a standardized log facility count is unrelated to news. Media coverage tracks the measurable quantity of pollution a firm reports rather than the ecological richness of the area around its operations. Investors observe the elevated costs in financial statements, but no coverage links them to their ecological source.

Table IA.9 examines whether the BFI return premium varies with aggregate biodiversity attention. Column (1) reports the BFI baseline on the 2000–2023 NYT-coverage window: BFI carries a coefficient of 0.231 percent per month. Columns (2)–(5) add interactions of BFI with the standardized monthly NYT biodiversity index and a climate-attention placebo. Across all specifications, the BFI $\times$ NYT_{bio} interaction is small and statistically indistinguishable from zero, and the climate-attention placebo is similarly null. The BFI coefficient itself remains positive and significant in every column, including the saturated specification with the full control set. The premium does not depend on aggregate attention cycles, and biodiversity news coverage is not the mechanism through which prices adjust.

Table IA.8: Biodiversity footprint, toxic releases, and biodiversity news

This table tests whether the biodiversity footprint predicts the incidence of biodiversity-related news. In columns (1)–(2) the dependent variable is the annual count of biodiversity news incidents recorded by RepRisk, estimated by Poisson pseudo-maximum-likelihood (PPML); in columns (3)–(4) it is an indicator for any such incident, estimated by a linear probability model (LPM). The sample covers the RepRisk coverage window (2007 onward). BFI is the standardized biodiversity footprint index. ln(Facilities) is the standardized log facility count. All regressors are lagged one year. Every specification controls for lagged toxic releases intensity, log assets, leverage, return on assets, Tobin's Q, tangibility, and log book-to-market, and includes firm and year fixed effects. Standard errors clustered by firm appear in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	PPML: news count		LPM: any news	
	BFI	ln(Fac)	BFI	ln(Fac)
BFI_{t-1} (std)	0.028		0.001	
	(0.090)		(0.012)	
Releases_{t-1} (std)	0.271**	0.223*	0.013*	0.010
	(0.112)	(0.121)	(0.007)	(0.007)
$\ln(\text{Assets})_{t-1}$	0.442**	0.415***	0.041***	0.039***
	(0.181)	(0.155)	(0.015)	(0.015)
Leverage_{t-1}	-1.063	-1.065	-0.097**	-0.096**
	(0.662)	(0.662)	(0.044)	(0.044)
ROA_{t-1}	0.001	0.001	0.000	0.000
	(0.001)	(0.001)	(0.000)	(0.000)
$\ln(\text{Facilities})_{t-1}$ (std)		0.074		0.008
		(0.073)		(0.010)
Firms	327	327	754	754
Observations	4,104	4,104	7,787	7,787

Table IA.9: Fama-MacBeth tests of BFI and biodiversity attention

This table reports Fama-MacBeth (1973) cross-sectional regressions of monthly stock returns on BFI and its interactions with the New York Times biodiversity attention index. The sample runs from January 2000 through December 2023, the period covered by the NYT biodiversity index. NYT_{bio} is the monthly mean of a daily indicator that takes value one when a NYT article that day discusses biodiversity, standardized to mean zero and unit variance over the sample. $\text{NYT}_{\text{climate}}$ is the analogous index built from climate-themed daily coverage. Both interactions multiply standardized BFI by the standardized monthly NYT index. Column (1) reports the BFI baseline on the NYT-coverage window. Columns (2) and (3) add the BFI $\times$ NYT_{bio} and BFI $\times$ $\text{NYT}_{\text{climate}}$ interactions one at a time. Column (4) adds both interactions jointly. Column (5) adds both interactions and the full standard control set (log book-to-market, momentum, log market equity, ROA, leverage). All specifications include 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)
BFI (std.)	0.231***	0.196***	0.216***	0.198***	0.160***
	(3.51)	(3.28)	(3.66)	(3.72)	(3.14)
BFI x NYTbio		0.003		-0.001	-0.016
		(0.10)		(-0.03)	(-0.70)
BFI x NYTclim			0.012	0.006	0.004
			(0.54)	(0.29)	(0.36)
Industry FE	✓	✓	✓	✓	✓
Size, B/M, Momentum	×	×	×	×	✓
ROA, Leverage	×	×	×	×	✓
Months	288	288	288	288	288
Avg firms/month	413	413	413	413	413

IA.10 Institutional ownership and analyst coverage

Table IA.10 asks whether institutional ownership or analyst coverage moderates the BFI premium, interacting BFI with an indicator for each. Both specifications control for size, book-to-market, and momentum. Institutional ownership matters. The BFI $\times$ Low IO interaction is 0.34% per month and significant, and the premium is small among the remaining firms, so it concentrates in the bottom tercile of ownership. Analyst coverage shows a different pattern. Estimated across the full panel, including the 18% of firm-months with no coverage, the BFI $\times$ No coverage interaction is -0.03% and insignificant, so covered and uncovered firms earn the same premium. The evidence therefore shows that the premium persists where sophisticated capital is scarce, regardless of analyst coverage.

I show earlier that the premium arises from forecast errors that correct when earnings are released, so the announcement window is a more direct test of whether greater analyst coverage resolves the errors. If it did, the announcement return would be smaller for the firms analysts follow. Table IA.11 regresses the three-day announcement cumulative abnormal return on BFI for firms above and below the median coverage split. BFI predicts the announcement return in both groups, so higher analyst coverage does not attenuate the announcement correction.

Table IA.10: Institutional ownership, analyst coverage, and the premium: Fama-MacBeth interactions

Monthly Fama-MacBeth regressions of excess returns on BFI and its interaction with one neglect indicator. Column 1 interacts BFI with *Low IO*, an indicator for the bottom tercile of institutional ownership formed within TRI vintage, on the sample of firms with 13-F ownership data (1999–2024). Column 2 interacts BFI with *No coverage*, an indicator for firms with zero I/B/E/S analysts, on the full panel (18% of firm-months have no coverage). With the dummy interaction, the BFI row is the premium in the reference group (non-low-IO firms in column 1, covered firms in column 2), and BFI plus the interaction is the premium among low-IO or uncovered firms. BFI is cross-sectionally standardized each month. Log market equity, log book-to-market, and momentum (cumulative returns $t-12$ to $t-2$) enter as controls, and both columns include Fama-French 12-industry fixed effects. Newey-West(6) t-statistics in parentheses. Coefficients in percent per month. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	Inst. ownership	Analyst coverage
BFI (std)	0.106	0.276***
	(0.76)	(5.38)
BFI × Low IO	0.335**	
	(2.18)	
Low IO	-0.071	
	(-0.32)	
BFI × No coverage		-0.034
		(-0.18)
No coverage		0.079
		(0.25)
Log ME	-0.020	0.011
	(-0.37)	(0.23)
Log book-to-market	-0.905***	-0.967***
	(-7.92)	(-9.79)
Momentum ($t-12, t-2$)	-0.623*	-0.396
	(-1.85)	(-1.53)
Industry FE	✓	✓
N months	312	428
Avg. firms/month	427	436
Avg. R^2	0.13	0.13

Table IA.11: Earnings announcement CARs by analyst coverage

This table reports earnings-announcement CAR$[-1, +1]$ regressions (market-adjusted, in percent) on standardized BFI by analyst coverage. High Coverage equals one for firms with above-median analyst coverage (median = 17), and firms without I/B/E/S coverage are coded as zero analysts. BFI is standardized within TRI vintage year. All specifications include Fama-French 12-industry dummies (coefficients suppressed) and year fixed effects. t-statistics, based on standard errors clustered by firm and year-month, are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Full Sample	Below Median	Above Median	Interaction
BFI (std.)	0.100**	0.114*	0.112**	0.062
	(2.31)	(1.66)	(2.19)	(0.97)
High Coverage				0.126
				(1.57)
BFI × High Coverage				0.075
				(1.09)
ln(ME)	-0.025	0.011	-0.094***	-0.047*
	(-1.08)	(0.28)	(-2.77)	(-1.73)
ln(B/M)	-0.043	0.026	-0.108*	-0.046
	(-0.83)	(0.28)	(-1.84)	(-0.90)
N	66,392	32,433	33,959	66,392
Firms	1,697	1,530	948	1,697
Adj. R^2	0.002	0.001	0.002	0.002

IA.11 Biodiversity disclosure

Table IA.12 tests whether BFI or EPA enforcement 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. Both BFI and the lagged enforcement indicator are statistically indistinguishable from zero in all three specifications, confirming that neither ecological exposure nor enforcement shapes voluntary disclosure.

Firms also report biodiversity practices voluntarily through the Carbon Disclosure Project (CDP), and here biodiversity exposure does affect disclosure. The CDP asks firms five questions: whether a firm operates near areas important for biodiversity, uses biodiversity indicators to monitor performance, has taken conservation actions, has made a public commitment, and has published biodiversity information in voluntary reports. Table 9 regresses each indicator, coded one for a "Yes" response, on BFI with industry and year fixed effects. BFI is negative for all five, reaching -0.111 for biodiversity actions and -0.099 for voluntary publication, so firms with greater biodiversity exposure are less likely to report actions or to disclose at all. Disclosure is lowest where exposure is highest. Abatement offsets this avoidance. Among high-exposure firms, those that report pollution-prevention activity disclose more, and the BFI $\times$ Abatement interaction is positive for biodiversity indicators and actions.

Table IA.12: BFI, EPA enforcement, and corporate biodiversity disclosure

This table tests whether BFI or EPA enforcement predicts corporate biodiversity disclosure in 10-K filings. 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. *BFI* (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. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	Aggregate	Negative Sentiment	Regulation
Enforcement $(t-1)$	-0.016	0.011	-0.002
	(-1.54)	(1.10)	(-0.32)
BFI (std)	-0.008	-0.009	-0.009
	(-1.37)	(-1.40)	(-1.62)
Releases (std)	0.006	0.006	0.004
	(1.28)	(1.03)	(0.97)
ln(Assets)	0.007	-0.020	0.002
	(0.73)	(-1.27)	(0.23)
ln(B/M)	-0.000	0.003	-0.002
	(-0.04)	(0.41)	(-0.38)
ROA	0.000	-0.000	-0.000
	(0.46)	(-0.00)	(-0.28)
ln(PPE/Assets)	0.018	0.026	0.003
	(1.28)	(1.33)	(0.26)
Leverage	-0.038	-0.014	-0.039
	(-1.34)	(-0.58)	(-1.61)
N	7,063	7,063	7,063
Firms	611	611	611
R^2	0.544	0.450	0.580

Section D: Production Cost Channels

IA.12 Robustness to candidate cost-channel controls

Table IA.13 adds four candidate cost-channel variables to the COGS/Sales regression of Table 5 to test whether the BFI effect on production costs is absorbed by observable proxies for the operating environment. The four channels are the share of a firm's TRI facilities located in mini-NEPA (State Environmental Policy Act) states; an indicator for water-intensive industries; the intensity of off-site waste transfers (POTW plus disposal, recycling, and treatment) scaled by total assets; and the share of a firm's facilities subject to CERCLA (Superfund) cleanup liability. Columns (2)–(5) add each candidate channel one at a time, and column (6) adds all four jointly. Capital intensity (ln(PPE/Assets)) is in the standard control set in every column.

The BFI coefficient is stable across all six specifications, moving from 0.353% in the baseline to 0.376% when all four channels enter jointly. The t-statistic on BFI stays above 3.8 throughout. Each candidate channel is statistically insignificant on its own and remains insignificant when entered jointly. The COGS premium for high-BFI firms is therefore not concentrated in firms with the obvious legal, industrial, or waste-handling exposures one would expect to absorb a location-bound friction. The cost arises from a set of smaller, less directly observable frictions that no single proxy captures.

Table IA.13: Candidate cost channels and the BFI premium

This table examines whether four candidate cost channels absorb the BFI effect on COGS/Sales. The sample matches the quarterly panel of Table 5. *SEPA share* is the share of a firm's TRI facilities located in states with a mini-NEPA (State Environmental Policy Act) statute that mandates environmental review of state actions affecting protected resources. *Water-intensive industry* is an indicator equal to one for industries with documented high water dependence (food, textiles, paper, chemicals, petroleum refining, primary metals, electric utilities; SIC codes 2000–2099, 2200–2299, 2600–2699, 2800–2999, 3300–3399, 4900–4999). *Offsite waste intensity* is the firm's total off-site waste transfers (POTW, off-site disposal, off-site recycling, and off-site treatment, in pounds) divided by total assets, winsorized at the 1st and 99th percentiles by year. *CERCLA share* is the share of a firm's TRI releases subject to CERCLA cleanup liability (water, underground injection, landfills, and on-site surface disposal), excluding air pollution which is not covered by CERCLA (Bellon, 2025). Column (1) reproduces the baseline COGS specification. Columns (2)–(5) add each candidate channel one at a time. Column (6) adds all four jointly. Capital intensity (log PPE/Assets) is in the standard control set in every column and is reported for reference. Controls (log assets, log book-to-market, ROA, leverage, log PPE/Assets), each lagged one quarter, and the one-quarter lagged dependent variable are included throughout. 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)	(2)	(3)	(4)	(5)	(6)
BFI (standardized)	0.353***	0.381***	0.353***	0.380***	0.356***	0.376***
	(3.83)	(4.02)	(3.83)	(4.04)	(3.80)	(3.90)
SEPA share		-0.112				0.019
		(-0.30)				(0.05)
Water-intensive industry			0.206			0.132
			(0.70)			(0.43)
Offsite waste intensity				0.000		0.000
				(0.84)		(0.94)
CERCLA share					0.525	0.544
					(1.41)	(1.41)
Log PPE/Assets	0.209	0.251	0.213	0.251	0.270	0.308
	(0.78)	(0.85)	(0.79)	(0.86)	(0.96)	(1.01)
Firm FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Observations	64,981	60,788	64,981	60,962	60,090	57,762
Adj. R^2	0.912	0.916	0.912	0.915	0.912	0.915

IA.13 Production technology and the cost penalty

Table IA.14 adds three production-technology measures to the COGS/Sales regression of Table 5 to test which features of the firm's input mix absorb the BFI effect. The log labor-to-capital ratio is the natural log of Compustat annual employees divided by net property, plant, and equipment, capturing the production-function input mix. The Donangelo (2014) labor share of value added is imputed annual staff expense (Compustat employees times the BLS state by 6-digit NAICS average weekly wage, multiplied by 52) divided by the sum of imputed staff expense, operating income after depreciation, and depreciation. The Novy-Marx (2011) variable cost share is COGS divided by the sum of COGS, selling and administrative expense, and depreciation, all annualized. All specifications run on the common subsample of 34,008 firm-quarters where the three measures are jointly available. Columns (2)–(4) add each measure one at a time, and column (5) adds all three jointly.

Each measure absorbs a portion of the BFI cost effect. The log labor-to-capital ratio shrinks the BFI coefficient from 0.585 to 0.450, an attenuation of roughly one-fourth. On the natural log labor-to-capital sample of 72,168 firm-quarters, where the Donangelo and variable cost share measures are not required, BFI falls from 0.405 to 0.321 when the ratio is added under firm and year fixed effects, the same one-fourth attenuation. The joint-sample baseline of 0.585 is larger because requiring both the imputed labor share and the variable cost share restricts the panel to firm-quarters with complete operating-cost accounting, a subsample with greater BFI dispersion. The Donangelo labor share and the variable cost share enter with much larger coefficients because their numerators overlap with the COGS measurement, and each in isolation attenuates the BFI coefficient by a comparable amount. In column (5), the two ratio measures absorb the variation in production cost structure. The labor-to-capital ratio flips sign, and the joint BFI coefficient falls to 0.182. The cleanest economic statement uses the log labor-to-capital ratio, which has no accounting overlap with the dependent variable. High-BFI firms run more labor-intensive production technologies, and this technology choice accounts for about one-fourth of the BFI cost penalty.

Table IA.14: Production technology and the BFI cost penalty

This table adds three production-technology measures to the COGS/Sales regression of Table 5 to test which features of the firm's input mix absorb the BFI effect. *Log labor-to-capital ratio* is the natural log of Compustat annual employees divided by net property, plant, and equipment, scaled to the within-vintage cross-section. *Labor share of value added* follows Donangelo (2014): imputed annual staff expense (Compustat employees times the BLS Quarterly Census of Employment and Wages state by 6-digit NAICS average weekly wage, multiplied by 52) divided by the sum of imputed staff expense, operating income after depreciation, and depreciation. *Variable cost share* follows Novy-Marx (2011): COGS divided by the sum of COGS, selling and administrative expense, and depreciation, all annualized. All specifications are estimated on the common subsample where the three measures are jointly available, and include firm and year fixed effects, the one-quarter lagged dependent variable, and the standard control set (log assets, log book-to-market, ROA, leverage, log PPE/Assets, each lagged one quarter). Column (1) reproduces the baseline COGS specification on this subsample. Columns (2)–(4) add each production-technology measure one at a time. Column (5) adds all three jointly. t-statistics based on standard errors clustered by firm appear in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1)	(2)	(3)	(4)	(5)
BFI (standardized)	0.585***	0.450***	0.442***	0.413***	0.182*
	(4.72)	(4.03)	(3.55)	(4.15)	(1.69)
Log labor-to-capital ratio		2.356***			-0.309
		(5.45)			(-1.36)
Labor share of value added			14.584***		19.029***
			(20.54)		(26.39)
Variable cost share				50.148***	68.294***
				(16.73)	(26.61)
Log PPE/Assets	0.055	1.892***	-0.195	0.542**	0.150
	(0.20)	(5.00)	(-0.68)	(2.53)	(0.56)
Firm FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓
Observations	34,008	34,008	34,008	34,008	34,008
Adj. R^2	0.942	0.943	0.955	0.952	0.974

IA.14 Enforcement penalties and delayed price discovery

Another channel through which biodiversity information affects returns is through enforcement action. Regulatory penalties provide a firm-specific signal about the cost implications of operating in biodiverse locations, complementing the errors in expectations channel.

The market absorbs this information slowly, and the reaction is concentrated in enforcement events that carry monetary penalties. I compute Fama-French three-factor cumulative abnormal returns around EPA enforcement conclusions, splitting events by whether a monetary penalty was recorded. Table IA.15 shows that the immediate three-day CAR is statistically zero, but a negative drift emerges over the following weeks, significant at the $[-10, +30]$ window where the high-minus-low BFI spread reaches -4.94%. Over the longer $[-20, +60]$ window the spread widens to -6.45%. This spread represents an equity value loss of approximately \$21 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 delayed price response suggests that enforcement penalties reveal new information about ecological costs rather than simply realize risks already reflected in prices.

The drift is specific to penalties rather than regulatory attention. Panel B of Table IA.15 reports that non-penalty actions produce a $[-20, +60]$ spread of only -1.06%—less than one-sixth the -6.45% spread on penalty events in Panel A. Investors therefore respond to concrete cost realizations rather than to regulatory attention alone.

The enforcement drift and the earnings-surprise mechanism are complementary rather than contradictory. The positive earnings surprises arise because investors misinterpret existing ecological costs as operational decline. The enforcement drift instead reflects new information about future compliance costs and heightened regulatory scrutiny that a penalty reveals for the first time. Enforcement is rare, with 341 matched monetary-penalty events affecting 181 firms in the post-2000 window, of which 331 enter the tercile drift estimation (Panel A of Table IA.15), so this correction does not offset the return premium.

Table IA.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 residuals from a Fama–French three-factor 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 BFI terciles based on their lagged (year $-$ 1) biodiversity footprint index. The event and firm counts in each panel header are the matched enforcement events; events without a non-missing within-year footprint tercile are excluded from the tercile columns and the H–L spread, so the per-tercile event counts sum to fewer than the header. 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 (341 events, 181 firms)					
CAR $[-1, +1]$	-0.36	-0.36	-0.05	0.31	(0.37)
CAR $[-5, +5]$	-0.45	1.15	-2.33*	-1.88	(-1.28)
CAR $[-10, +30]$	-0.98	1.26	-5.93**	-4.94*	(-1.74)
CAR $[-20, +45]$	-0.60	1.06	-5.97*	-5.37	(-1.44)
CAR $[-20, +60]$	-0.32	2.45	-6.77*	-6.45	(-1.57)
N events	134	108	89		
Panel B: Post-2000, without monetary penalty (3002 events, 522 firms)					
CAR $[-1, +1]$	0.11	0.01	0.05	-0.07	(-0.33)
CAR $[-5, +5]$	0.25	-0.08	-0.37	-0.63*	(-1.85)
CAR $[-10, +30]$	-0.18	-0.90**	-0.73	-0.55	(-0.82)
CAR $[-20, +45]$	-0.12	-1.04*	-1.06	-0.94	(-0.95)
CAR $[-20, +60]$	-0.33	-0.74	-1.39*	-1.06	(-0.97)
N events	925	1120	876		
Panel C: Full sample, with monetary penalty (961 events, 322 firms)					
CAR $[-1, +1]$	-0.20	-0.01	-0.12	0.08	(0.19)
CAR $[-5, +5]$	-0.19	0.53	-0.94*	-0.75	(-1.18)
CAR $[-10, +30]$	0.27	1.18	-1.55	-1.81	(-1.44)
CAR $[-20, +45]$	0.08	0.71	-2.09	-2.17	(-1.31)
CAR $[-20, +60]$	0.24	0.88	-2.11	-2.35	(-1.30)
N events	361	326	255		
Panel D: Full sample, without monetary penalty (3086 events, 537 firms)					
CAR $[-1, +1]$	0.09	0.01	0.09	0.00	(0.01)
CAR $[-5, +5]$	0.27	-0.06	-0.36	-0.63*	(-1.86)
CAR $[-10, +30]$	-0.16	-0.90**	-0.66	-0.49	(-0.74)
CAR $[-20, +45]$	-0.19	-1.06*	-0.98	-0.79	(-0.82)
CAR $[-20, +60]$	-0.41	-0.79	-1.38*	-0.97	(-0.90)
N events	962	1148	894		