



The Center for Economic Research
Columbian College of Arts & Sciences

WORKING PAPER SERIES

Where Firms Retire Carbon Offsets: Operational Footprints and Offset Quality

Alvaro Pedraza

Tomas Williams

Federica Zeni

Working Paper No. 2026-010

May, 2026

THE CENTER FOR ECONOMIC RESEARCH

Department of Economics
Columbian College of Arts & Sciences
The George Washington University
Washington, DC 20052
<https://cer.columbian.gwu.edu/>

Where Firms Retire Carbon Offsets: Operational Footprints and Offset Quality*

Alvaro Pedraza[†]

Tomas Williams[‡]

Federica Zeni[§]

Abstract

Although the climate benefit of carbon abatement does not depend on where it occurs, firms do not treat carbon offsets as geographically fungible. Using transaction-level data on corporate retirements in voluntary carbon markets, we show that foreign firms retire disproportionately more offsets in countries where they operate, but these local retirements are systematically lower quality than the same firms' retirements elsewhere. We distinguish two mechanisms. Operational presence may improve screening of local projects, or it may raise the private value of visible local sourcing. The evidence supports the second mechanism. The quality gap declines with firm experience, consistent with learning, but is strongest in settings where local visibility is likely to matter most, including countries with higher climate ambition and weaker governance. Subsidiary-entry events corroborate the within-firm pattern: when a firm establishes a new subsidiary in a country, retirements there rise sharply, but these offsets come from lower-quality projects than the firm's retirements in other countries. Finally, in project segments where demand is dominated by firms with local operations, prices are less closely aligned with project integrity. The results reveal a demand-side distortion in markets for global public goods, whereby location-specific reputational benefits shift demand toward lower-quality supply and weaken the informational content of prices.

JEL Codes: F18, L14, Q54, Q58, G32.

Keywords: Voluntary carbon markets, Operational footprint, Local Goodwill, Carbon offset quality.

1 Introduction

Carbon dioxide is the canonical global externality: the climate impact of a ton of avoided emissions is invariant to where the abatement occurs. In a frictionless market for carbon offsets, a verified ton in Brazil should be a perfect substitute for one in Indonesia. Conditional on seeking credible climate mitigation, corporate buyers should allocate spending toward the highest-integrity projects rather than

*We would like to thank Deeksha Gupta and Tarun Ramadorai for valuable comments.

[†]World Bank. Email: apedrazamorales@worldbank.org

[‡]George Washington University. Email: tomaswilliams@gwu.edu

[§]Ecole Polytechnique Federale de Lausanne and Swiss Finance Institute (SFI). Email: federica.zeni@epfl.ch

toward projects located where they operate. Using transaction-level data from the global voluntary carbon market (VCM), this paper shows that corporate demand departs sharply from this benchmark on both dimensions: firms tilt retirements — permanent cancellations of offset credits claimed against their own emissions — toward countries where they have an operational presence, and these locally-sourced credits are systematically lower in environmental integrity than the same firms’ retirements elsewhere. The pattern reveals a previously undocumented friction in markets for global public goods: when corporate buyers receive reputational benefits that are local to where they operate, demand sorts away from quality rather than toward it.

The voluntary carbon market has emerged as a key vehicle for corporate decarbonization, yet unlike compliance carbon markets, where firm-level emissions reductions are increasingly well-documented (Colmer et al., 2025), the voluntary market is characterized by pervasive quality concerns on the supply side: many projects fail standard additionality and permanence tests, raising the risk that issued credits do not reflect genuine abatement (Schneider et al., 2019; West et al., 2023; Gill-Wiehl et al., 2024; Calel et al., 2025).¹ Much less is known about how corporate demand responds to these differences. If buyers actively screen for integrity, the geographic composition of their portfolios should reflect that screening, and demand could discipline supply toward higher-quality projects. We examine this conjecture and find the opposite.

Figure 1 documents the basic fact: firms allocate a disproportionate share of their retirements to countries where they maintain an operational presence. Because these retirements occur in the voluntary market rather than under compliance regimes, the concentration is not driven by regulatory mandates. Unlike cap-and-trade systems that may require domestic sourcing (see, e.g., Shapiro and Walker, 2025), firms in our sample face no formal obligation to align offset portfolios with operational geography. The observed spatial distribution therefore reflects a strategic choice.²

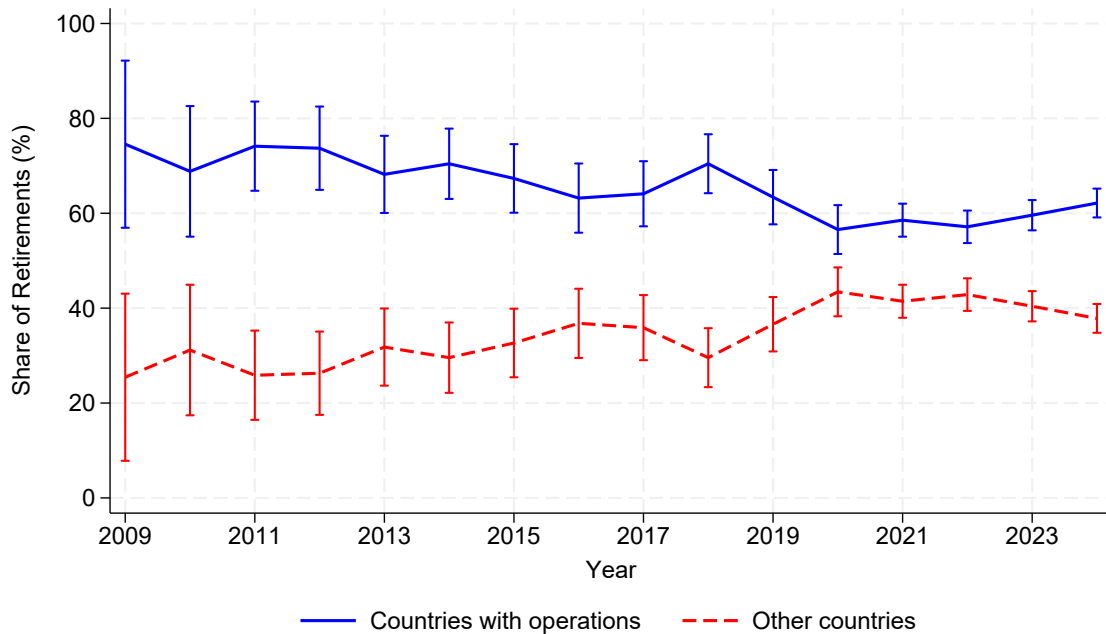
Two channels can rationalize this local bias. Under the *information channel*, operational presence enhances firms’ ability to evaluate project quality through local knowledge and proximity. Under the *goodwill channel*, firms instead derive non-pecuniary benefits from the *local visibility* of their sourcing — securing a “social license to operate,” building stakeholder relationships, or earning a community-engagement payoff that accrues only where local audiences can observe the action.³ Both channels predict local concentration. They differ in their predictions for the *quality* of local offsets: the information channel predicts a positive quality gradient — better information leads to better projects — while

¹A related literature on compliance carbon markets examines how emissions leakage between covered and uncovered sectors can also undermine the climate impact of carbon-pricing instruments (Fowle and Reguant, 2022).

²Some jurisdictions admit hybrid systems — most notably Colombia’s 2017 carbon tax, which permits tax neutralization via domestic offsets — we exclude such compliance-linked transactions from our primary analysis.

³The goodwill channel draws on a substantial non-market-strategy literature: locally salient environmental initiatives build legitimacy and reduce operational and political risks (Gunningham et al., 2004; Henisz et al., 2014); they provide insurance-like protection against reputational downside (Godfrey et al., 2009; Lins et al., 2017; Albuquerque et al., 2019); and multinationals target philanthropic giving and pollution geographically for strategic purposes (Bertrand et al., 2020; Ben-David et al., 2021; Broccardo et al., 2022).

Figure 1: Retirements by Geographical Presence



Note: The figure shows the annual mean share of carbon offsets retired in countries where the firm operates (solid blue line) and in countries where it does not (dashed red line). The two series partition each firm's retirements within a year and sum to 100% by construction. Error bars denote 95% confidence intervals around the cross-firm means.

the goodwill channel predicts a negative gradient, since the reputational payoff from local sourcing reduces the marginal value of project integrity.

We test these competing predictions by merging the universe of retirement records from major VCM registries with firm-level geographic revenue data and ex-post project quality ratings (MSCI, BeZero, and Sylvera). Identification uses firm-by-year fixed effects together with country-of-retirement-by-year fixed effects: we compare the quality of offsets the same firm retires in countries where it operates to those retired elsewhere in the same year, holding fixed both time-varying firm characteristics and country-year supply conditions.

Three sets of results emerge. First, firms do not allocate offset retirements as if geography were irrelevant. In our baseline specification, foreign firms with substantial operational presence in a project country tilt their retirement portfolios toward that country, but the tilt is concentrated in lower-rated projects. Within the same firm-year and conditional on country-year supply conditions, operational exposure is associated with higher retirement shares for lower-rated local projects and lower shares for higher-rated local projects. Thus, local sourcing is not a uniform preference for projects where firms operate; it is concentrated among lower-quality offsets. The pattern is robust to accounting for the quality of projects available in each country-year, to an ex-ante quality proxy based on observable project characteristics, and to aggregating project-level choices to firm-year exposure-quality cells.

Second, this pattern is difficult to reconcile with a pure information channel. If operational presence mainly improved local screening, firms should select better projects where they operate. Instead, the negative quality gradient is strongest where local visibility is likely to be most valuable. In countries with higher climate ambition, measured by the Climate Change Performance Index (CCPI), firms both concentrate retirements more strongly in operational markets and tilt those retirements toward lower-quality projects. Subsidiary-entry events sharpen this interpretation. After a U.S.-listed firm establishes a subsidiary in a country, its cumulative probability of retiring offsets there rises by 4.2 percentage points over three years, a 70% increase relative to the pre-entry baseline. The response reaches 9.6 percentage points in high-CCPI destinations and is essentially zero in low-CCPI destinations. The quality composition moves in the same direction: offsets retired in entry countries are lower quality than the same firms' retirements elsewhere. Although the gradient attenuates with country-specific firm experience, suggesting that information frictions matter, the dominant heterogeneity points to goodwill: the effect is sharpest in high-climate-ambition and weak-governance environments, where local visibility and stakeholder relationships are likely to be especially valuable.

Third, these demand distortions have equilibrium consequences for prices. In project segments where buyers with operational presence dominate, the relationship between average prices and project quality is sharply attenuated relative to segments dominated by other foreign buyers: prices reward integrity in the latter segment but become uninformative about quality in the former. This *price-quality decoupling* is, to our knowledge, the first documented equilibrium price signature of strategic non-pecuniary demand in an environmental market. It implies that as firms with strong local visibility incentives become a larger share of demand, the market's ability to communicate project integrity through prices weakens.

Taken together, the evidence is most consistent with the view that corporate visibility motives help organize demand spatially in a way that the supply-side integrity literature alone cannot detect, and that this reallocation may have measurable price-formation consequences. If the social cost of carbon exceeds the private reputational gains from local offsetting, the resulting allocation may be inefficient: capital may be diverted away from regions with the highest abatement potential and toward those with greater multinational presence, while prices may become less tightly aligned with environmental integrity.⁴ More broadly, the visibility versus integrity tradeoff we document may extend to other settings in which corporate buyers receive reputational credit for visible local environmental actions, with potential applications to green bonds, sustainability-linked debt, and corporate climate disclosures.

Contribution to the Literature. The paper makes three contributions to the existing literature. First, we provide the first systematic empirical decomposition of demand-side from supply-side integrity problems in the voluntary carbon market. An influential supply-side literature documents that many projects fail additionality and permanence tests (Schneider et al., 2019; West et al., 2023; Gill-Wiehl

⁴Related theoretical work shows that with imperfect certification, adverse selection can weaken voluntary carbon markets. When some buyers value non-carbon attributes, trade may persist despite quality uncertainty, weakening the link between transactions and climate effectiveness (Manshadi et al., 2025).

et al., 2024; Calel et al., 2025). We show that even within the set of issued credits — that is, holding supply fixed — demand sorts toward the lower-integrity end of the quality distribution in segments where buyers earn local visibility from the retirement. Improving supply-side standards alone will not eliminate the integrity problem in this market: demand-side governance must be part of the fix.

Second, we identify geography as a unifying variable connecting industrial organization (where firms locate) to climate-finance allocation (where mitigation dollars go). Recent work in corporate ESG and offset markets has emphasized capital-market and product-market pressures (Kim et al., 2023; Beck and Pedraza, 2026) and the role of non-pecuniary credit attributes in pricing and willingness-to-pay (Berg et al., 2025; Heeb et al., 2023). Relative to those studies, which document cross-sectional taste heterogeneity or aggregate willingness-to-pay, we identify a *geographic* organization of those preferences: firms align offset demand with their operational footprints, and that alignment has welfare consequences for the international allocation of climate finance. This connects to evidence that multinationals strategically allocate philanthropy and pollution across space (Bertrand et al., 2020, 2021; Ben-David et al., 2021).

Third, we provide the first evidence that strategic non-pecuniary demand has an *equilibrium price signature* in an environmental market. More broadly, this contribution connects to work on how environmental-market prices transmit underlying cost and policy shocks (Fabra and Reguant, 2014), but highlights a distinct mechanism: buyer-side visibility motives weaken the link between prices and environmental integrity. The price-quality decoupling we document complements the theoretical and empirical literature on equilibrium price effects of sustainable preferences (Oehmke and Opp, 2025; Pástor et al., 2025; Hartzmark and Shue, 2025) by adding a spatial dimension: prices become less informative about quality where local visibility substitutes for environmental efficacy. The two-channel framework underlying our results — an information channel rooted in proximity-based screening (in the tradition of Coval and Moskowitz, 1999, 2001) versus a goodwill channel rooted in social-license-to-operate considerations (Gunningham et al., 2004; Lins et al., 2017) — ties the paper to long-standing themes in non-market strategy while moving the literature toward an empirical test in which the two channels generate opposite predictions for portfolio composition.

2 Conceptual Framework

In this section, we present a stylized model that disciplines our empirical tests.

Information structure and timeline. The economy features three periods, $t \in \{0, 1, 2\}$, two carbon offset-producing countries, and two otherwise identical firms, each operating in one of the two countries. In each country, there are two offset projects that differ in quality: a high-quality project, indexed by $k = H$, and a low-quality project, indexed by $k = L$. Each project $k = \{H, L\}$ delivers a benefit θ_k from each unit of offsets q_k retired. For simplicity, we set the low quality to $\theta_L = 0$ and the high


quality to $\theta_H = \theta$. In addition, firms receive a local marginal benefit $\beta \geq 0$ when purchasing offsets in the country where they operate ($j = l$). This benefit is independent of the project quality and reflects reputational or stakeholder advantages associated with sourcing offsets domestically.

Firms differ in their information about offsets quality depending on their operational exposure to the country where offsets are produced. Firms operating in the same country as the offset projects (*local firms, l*) are differently informed about which project is high quality than firms operating abroad (*foreign firms, f*).

Formally, at time $t = 0$, in each country, each firm $j \in \{l, f\}$ receives a private signal about the project believed to be high quality, denoted by $s_j \in \{H, L\}$. The signal is correct with probability

$$\Pr(k = H | s_j = H) = 1 - \Pr(k = L | s_j = H) = \rho_j \in (1/2, 1).$$

After observing their signals, firms simultaneously choose the quantity of offsets to purchase from each project, taking prices as given. At time $t = 1$, each offset market clears and quantities are allocated. Finally, at time $t = 2$, the true quality of each project is realized. The model timeline is illustrated below.



Firm problem. Conditional on signal $s_j \in \{H, L\}$, firm $j \in \{l, f\}$ chooses quantities $\{q_{k,j}(s_j)\}_{k \in \{H, L\}}$ to maximize:

$$\max_{\{q_{k,j}(s_j)\}_{k \in \{H, L\}}} \left[B_{k,j}(s_j) q_{k,j}(s_j) - p_{k,j} q_{k,j}(s_j) - \frac{\gamma}{2} \text{Var}(\theta_k | s_j) q_{k,j}(s_j)^2 \right], \quad (1)$$

where beliefs are given by:

$$\mathbb{E}[\theta_k | s_j] = \begin{cases} \rho_j \theta & \text{if } k = s_j, \\ (1 - \rho_j) \theta & \text{if } k \neq s_j, \end{cases} \quad (2)$$

and variance:

$$\text{Var}(\theta_k | s_j) = \rho_j(1 - \rho_j)\theta^2, \quad \forall k \in \{H, L\}. \quad (3)$$

The expected benefits per unit of purchased offset are:

$$B_{k,j}(s_j) = [(\beta \mathbf{1}\{j = l\})^\sigma + (\mathbb{E}[\theta_k | s_j])^\sigma]^{1/\sigma}, \quad \sigma \leq 1. \quad (4)$$

Thus, a firm's marginal willingness to pay for an offset reflects two components: (i) the expected climate benefit implied by its private signal, and (ii) any non-pecuniary benefit associated with sourcing offsets locally. The relative weight of these two components determines how sensitive demand is to project quality. The CES specification captures diminishing marginal gains from retiring additional offsets, consistent with the idea that regulatory or reputational advantages may saturate beyond a certain level.

From the first order conditions, the optimal demand for each project satisfies:

$$q_{k,j}^*(s_j) = \frac{B_{k,j}(s_j) - p_{k,j}}{\gamma \text{Var}(\theta_k | s_j)}, \quad k \in \{H, L\}. \quad (5)$$

Offset producers. In each country, offset producers of type $k \in \{H, L\}$ sell competitively to both firms $j \in \{l, f\}$ and face convex production costs. Each offset market is segmented by the signal received by the local and foreign firm, (s_l, s_f) , with prices adjusting so that supply equals demand in each segment. Specifically, each producer k chooses quantities to sell to foreign and local firms, $q_{k,j}$, to maximize profits:

$$\max_{q_{k,j}} \sum_{j \in \{l, f\}} \left[p_{k,j} q_{k,j} - \frac{\phi_k}{2} q_{k,j}^2 \right], \quad (6)$$

where ϕ_k is a cost parameter, which we assume $\phi_L = \phi_H = \phi$. Worth noting is that while the physical generation of carbon credits from a given project is a joint process, voluntary carbon markets are highly decentralized and rely on bilateral, over-the-counter transactions.⁵ Consequently, we model ϕ_k not merely as a physical production cost, but as capturing segment-specific transaction, search, and contracting frictions. Because developers must independently navigate different broker networks and buyer relationships to clear quantities in local versus foreign markets, these marketing costs are separable. The convex formulation implies that while initial relationship-driven sales within a segment are relatively low-cost, the marginal cost of pushing larger volumes into a specific segment increases as developers must expend more effort to find willing counterparties.

The first-order condition with respect to each $q_{k,j}$ yields

$$q_{k,j}^* = \frac{p_{k,j}}{\phi}, \quad j \in \{l, f\}, \quad k = \{H, L\}. \quad (7)$$

⁵See, among others, [Hamilton et al. \(2010\)](#) and [Hamilton et al. \(2012\)](#), State of the Voluntary Carbon Markets (Ecosystem Marketplace), which documents the OTC and project-specific pricing structure of voluntary offsets.

Equilibrium. Equilibrium prices and quantities are determined by market clearing.

Formally, an equilibrium is a collection of prices and quantities such that: (i) firms choose offset quantities to maximize expected utility (1) given prices and their signal; (ii) producers choose quantities to maximize profits (6) given prices; and (iii) markets clear, so that for each project type $k \in \{H, L\}$ and firm $j \in \{l, f\}$, supply equals demand.

Given the linear-quadratic structure, equilibrium prices are signal-contingent and given by

$$p_{k,j}^*(s_j) = \frac{\phi}{\phi + \gamma \text{Var}(\theta_k | s_j)} B_{k,j}(s_j), \quad k \in \{H, L\}, j \in \{l, f\}. \quad (8)$$

2.1 Model Implications

Since the signal is correct with probability ρ_j , ex-post prices are obtained by averaging over signal realizations. For each project type $k \in \{H, L\}$, the expected price charged to firm j is:

$$p_{H,j}^* = \rho_j p_{H,j}^*(H) + (1 - \rho_j) p_{H,j}^*(L), \quad p_{L,j}^* = \rho_j p_{L,j}^*(L) + (1 - \rho_j) p_{L,j}^*(H), \quad (9)$$

with $p_{k,j}^*(s_j)$ as in (8). Furthermore, using the supply relation in (7), expected quantities are:

$$q_{H,j}^* = \rho_j \frac{p_{H,j}^*(H)}{\phi} + (1 - \rho_j) \frac{p_{H,j}^*(L)}{\phi}, \quad q_{L,j}^* = \rho_j \frac{p_{L,j}^*(L)}{\phi} + (1 - \rho_j) \frac{p_{L,j}^*(H)}{\phi}. \quad (10)$$

We prove the model's testable implications in Appendix B.

Proposition 1 (Local bias). *Let $\phi > 0$, $\gamma > 0$, $\theta > 0$, $\sigma \in (0, 1]$, and $\rho_f, \rho_l \in (0.5, 1)$. If the local benefit β satisfies*

$$\beta > \kappa(\rho_f, \rho_l) \theta, \quad (11)$$

where $\kappa_{\rho_f}(\rho_f, \rho_l) > 0$ and $\kappa_{\rho_l}(\rho_f, \rho_l) < 0$, then the total share of offset retirements exhibits a local bias:

$$\frac{Q_l^*}{Q_l^* + Q_f^*} > \frac{1}{2}, \quad (12)$$

with total average quantities in each country defined as

$$Q_j^* = q_{H,j}^* + q_{L,j}^*, \quad j \in \{l, f\}.$$

Proposition 1 provides a sufficient condition for the equilibrium share of surrendered offsets to exhibit a local bias. In particular, condition (11) sets a lower bound on the local benefit β necessary to generate this outcome in equilibrium.

Since $\kappa_{\rho_f}(\rho_f, \rho_l) > 0$ and $\kappa_{\rho_l}(\rho_f, \rho_l) < 0$, the larger the information gap between local and foreign firms (i.e., $\rho_l - \rho_f > 0$), the more slack the condition becomes. Intuitively, when firms are more informed about the quality of local offsets, they are more inclined to purchase domestically even without additional non-pecuniary gains.

Proposition 2 (Quality gradient). *Let $\phi > 0$, $\gamma > 0$, $\theta > 0$, $\sigma \in (0, 1]$, and $\rho_f, \rho_l \in (0.5, 1)$. If the local benefit β satisfies*

$$\beta > v(\rho_f, \rho_l) \theta, \quad (13)$$

where $v_{\rho_f}(\rho_f, \rho_l) < 0$ and $v_{\rho_l}(\rho_f, \rho_l) > 0$, then the fraction of high-quality offsets retired from local countries is strictly lower than that from foreign countries:

$$\frac{q_{H,l}^*}{Q_l^*} < \frac{q_{H,f}^*}{Q_f^*}. \quad (14)$$

Proposition 2 characterizes how local benefits distort the composition of retirements. When β exceeds the threshold in (13), the marginal valuation of local offsets becomes less sensitive to expected quality. As a result, the equilibrium share of high-quality offsets is lower in local markets than in foreign markets.

Unlike Proposition 1, the information gap now works in the opposite direction. Since $v_{\rho_f}(\rho_f, \rho_l) < 0$ and $v_{\rho_l}(\rho_f, \rho_l) > 0$, the condition becomes more binding as local firms become better informed relative to foreign firms (i.e., as $\rho_l - \rho_f$ increases). Intuitively, when local firms have more precise information about project quality, quality differences are more salient in their marginal valuation. In that case, a larger non-pecuniary benefit β is required to offset the importance of quality and generate a lower share of high-quality offsets locally. Conversely, when informational frictions are more severe or foreign firms are relatively better informed, non-pecuniary motives more easily dominate quality considerations, and the condition becomes slack.

Proposition 3 (Price-quality decoupling). *Let $\phi > 0$, $\gamma > 0$, $\theta > 0$, $\sigma \in (0, 1]$, and $\rho_f, \rho_l \in (0.5, 1)$. There always exists a threshold $f(\rho_l, \rho_f, \sigma, \theta) > 0$ such that, if the local benefit β satisfies*

$$\beta > f(\rho_l, \rho_f, \sigma, \theta), \quad (15)$$

then the equilibrium local price becomes less sensitive to expected project quality than the foreign price, that is

$$\frac{\partial p_l^*}{\partial \theta} < \frac{\partial p_f^*}{\partial \theta}. \quad (16)$$

Furthermore, the threshold satisfies $f_{\rho_l}(\rho_l, \rho_f, \sigma, \theta) > 0$ and $f_{\rho_f}(\rho_l, \rho_f, \sigma, \theta) < 0$.

Proposition 3 links buyer composition to equilibrium price formation. Quality-driven climate benefits and local non-pecuniary gains are imperfect substitutes. Thus, an increase in β reduces the marginal contribution of quality to total willingness to pay in local markets. As a result, equilibrium prices in segments with stronger local demand respond less to improvements in underlying project quality. In other words, the sensitivity of prices to θ is lower in markets where non-pecuniary motives constitute a larger share of marginal valuation.

When preferences are strictly convex in aggregation ($\sigma \in (0, 1)$), this dampening effect arises through CES substitution: local non-pecuniary benefits directly crowd out the marginal weight of quality in willingness to pay, generating a nonlinear attenuation of price sensitivity that strengthens in β . In the limiting linear case ($\sigma = 1$), this substitution channel disappears and β affects equilibrium prices only as a level shifter. Nonetheless, even in this benchmark, sufficiently large local benefits still reduce the relative responsiveness of local prices to quality changes by amplifying the baseline wedge between local and foreign demand compositions.

3 Testable Hypotheses

The conceptual framework yields three predictions that guide the empirical analysis. We examine whether the data are consistent with an environment in which local non-pecuniary benefits meaningfully influence firms' offset demand, even after accounting for informational differences across markets.

Specifically, we test the following three hypotheses:

Hypothesis 1 (Local Bias). *Firms allocate a larger share of their carbon offset retirements to countries where they maintain an operational footprint.*

Under geographic fungibility, a firm's offset portfolio should be independent of its revenue geography. Systematic concentration of retirements in operationally relevant countries indicates that location-specific factors enter firms' demand.

Hypothesis 2 (Negative Quality Gradient). *Offsets retired in a firm's country of operations exhibit lower average environmental integrity than those retired in foreign jurisdictions.*

If local non-pecuniary benefits are independent of project integrity, operational presence reduces the marginal valuation of quality in domestic markets, generating a negative quality gradient within firms.

The framework further implies that the magnitude of this gradient depends on the relative importance of environmental integrity and local signaling in firms' objective. When the marginal value of true emissions reductions is high, or when informational frictions are limited, the negative quality gradient

should attenuate. Conversely, when local signaling benefits are relatively more important, the gradient should strengthen.

Hypothesis 3 (Price–Quality Decoupling). *The relationship between offset prices and environmental integrity is weaker in jurisdictions with greater local operational demand.*

When non-pecuniary motives constitute a larger share of marginal willingness to pay, prices become less responsive to differences in project quality. As a result, prices in markets with stronger local operational demand are less tightly aligned with underlying environmental integrity.

4 Data and aggregate patterns

To examine how corporate participation in voluntary carbon markets is linked to firms’ operational footprint, we construct a comprehensive database that combines project-level offset retirement data, firm-level operational and emissions data, and project quality assessments.

4.1 Carbon Offsets

Voluntary carbon offsets do not have centralized transaction-level purchase data; offset use becomes observable only when credits are formally retired. We therefore assemble project-level retirement data from the four largest voluntary carbon registries: Verra, Gold Standard, the American Carbon Registry, and the Climate Action Reserve, using data compiled by the Berkeley Carbon Trading Project ([Berkeley Carbon Trading Project, 2025](#)). The resulting dataset covers 3,565 distinct offset projects and more than 350,000 retirement transactions between 2009 and 2024, corresponding to a total of 759 million tons of CO₂ retired over the period (Table 1, Panel A).

To identify corporate users of offsets, we extract and standardize information from unstructured text fields describing the retirement beneficiary. We use large-language-model-assisted text processing to harmonize beneficiary names, resolve spelling and formatting inconsistencies, and map retirements to ultimate parent companies. Retirements attributed to subsidiaries are assigned to the parent firm, and all retirements are aggregated at the headquarters (parent-company) level. We do not attempt to distinguish between direct subsidiary purchases and centralized corporate procurement. This procedure identifies 1,582 publicly listed firms that retire offsets from 1,889 distinct projects across 88 countries. Projects used by listed firms account for approximately 87% of total retired volume in the full sample, or 659 million tons of CO₂ (Table 1, Panel A).

Offset quality ratings. To measure offset quality, we match projects to MSCI’s Carbon Project Ratings ([MSCI, 2024](#)), which provide ex-post assessments of project additionality, permanence, leakage risk, and monitoring credibility. MSCI covers 2,549 projects in our sample and assigns each project both a

letter grade — ranging from *AAA* (highest integrity) to *CCC* (lowest integrity) — and an associated numerical integrity score, where higher values denote higher environmental quality. As shown in Table 1, the average numerical rating is 2.99 in the full sample and 2.98 among projects used by listed firms. The distribution reflects a meaningful quality spread; the interquartile range spans from a numerical score of 2.69 (corresponding to a letter grade of *B*) at the 25th percentile to 3.24 (*BBB*) at the 75th percentile.

We complement MSCI ratings with project assessments from BeZero Carbon and Sylvera, two independent analytics providers that specialize in evaluating the environmental integrity and risk profile of carbon offset projects, which together cover 375 and 358 projects in our sample, respectively. These ratings provide additional standardized evaluations of project credibility and permanence. A key feature of all three rating datasets is that, over our sample period, they are largely ex post. MSCI released its ratings in bulk only in late 2024 for projects that had already issued and retired credits, while BeZero Carbon and Sylvera began assigning ratings only after 2021, often after projects had already begun issuing credits and retirements had taken place. As a result, these ratings generally do not reflect the information available to firms at the time of offset purchase or retirement.

Carbon offset prices. Information on carbon offset prices is drawn from Sylvera, which compiles transaction-level data and broker-quoted information from voluntary carbon markets and reports average prices at the project-vintage level. Sylvera aggregates observed trades and indicative quotes into weekly average prices for each project and vintage, reflecting prevailing market conditions rather than contract-specific terms. Coverage is most comprehensive from 2020 onward, coinciding with the rapid expansion of trading activity and price reporting in voluntary carbon markets. We aggregate these weekly data to the annual level by project and vintage to align with the timing of observed retirements. Price information is available for 1,049 projects in the full sample and 750 projects used by listed firms (Table 1, Panel A).

These prices should be interpreted as indicative measures of the market value of a credit from a given project and vintage around the time of retirement, rather than as the exact prices paid by the retiring firm. We do not observe the purchase date, so realized acquisition costs may differ from the price observed at retirement if firms acquired credits earlier and held them in inventory. At the same time, each retirement is linked to the issuance date of the retired credits, which allows us to observe the gap between issuance and retirement. In our sample, that gap is relatively short: the median is about 0.9 years, and 80 percent of retired CO₂ tonnes are retired within two years of issuance. Since purchase must occur no earlier than issuance and no later than retirement, these patterns imply that retirement-year prices are likely to be reasonably close to the market conditions prevailing when firms acquired the credits. In the analysis, prices are therefore used as comparative measures of the opportunity cost and market value of retiring credits across projects and vintages, rather than as exact transaction prices faced by firms.

As shown in Panel B, average offset prices are similar across samples, with a mean price of \$6.35 per ton in the full sample and \$6.34 among projects used by listed firms. Price dispersion is substantial, with standard deviations exceeding the mean and prices ranging from less than a dollar to over \$115 per ton. This wide variation reflects heterogeneity in project type, location, vintage, and perceived quality, as well as thin trading and episodic transactions in voluntary carbon markets.

4.2 Firms' Geographical Footprint, Emissions, and Financial Data

To characterize firms' global operations, we merge our carbon offset retirement data with *FactSet GeoRevenue*, a database that provides a granular breakdown of publicly listed firms' revenues by country. GeoRevenue uses firms' segment disclosures, regulatory filings, and proprietary text analysis to assign, for each firm-year, the percentage of consolidated revenue generated in each country of operation. This allows us to construct a consistent and time-varying measure of firms' geographical footprint across markets.

Our final matched sample includes 1,450 publicly listed firms domiciled in 54 countries that retired at least one metric ton of CO₂ to offset emissions at some point between 2009 and 2024. The largest headquarters locations are the United States (325 firms), the United Kingdom (128), Japan (94), and Australia (81). Publicly listed firms account for approximately 40% of all carbon offset retirements observed in the combined registry dataset over the sample period. For each firm-year, we observe the full distribution of revenues across countries, which we use to identify whether a firm has an operational presence — defined as positive revenue — in a given country at the time offsets are retired.

Table 2 highlights a striking asymmetry between the geographic origin of firms retiring offsets and the location of the underlying carbon offset projects. While the majority of listed firms using carbon offsets are domiciled in high-income countries, most carbon offset projects are located in low- and middle-income countries, including India, China, Indonesia, Brazil, and several countries in Sub-Saharan Africa. This pattern underscores the role of voluntary carbon markets as a channel through which firms headquartered in advanced economies direct resources toward mitigation projects in developing countries. As such, voluntary carbon markets are often viewed as a potential source of private climate finance for emissions-reduction projects in low- and middle-income countries, where abatement opportunities may be relatively inexpensive and financing constraints more binding (World Bank, 2025).

Firm emissions and financial characteristics. Firm carbon emissions and financial information are obtained from ISS STOXX. We use Scope 1 and Scope 2 emissions, total sales, and balance-sheet characteristics, all reported at the annual frequency. Emission intensity is constructed as total emissions divided by revenue, providing a standardized measure of firms' baseline carbon exposure that is comparable across industries and over time.

Climate targets. Information on firms' climate commitments is obtained from the Carbon Disclosure Project (CDP). CDP is a global non-profit platform through which firms voluntarily report standardized information on greenhouse gas emissions, climate risks, and mitigation strategies to investors and other stakeholders. Using these data, we construct an indicator for whether a firm has an approved Science Based Targets initiative (SBTi) emissions-reduction target in a given year. This variable captures the formal adoption of externally validated decarbonization commitments and allows us to distinguish firms with explicit climate targets from those without.

Table 3 reports summary statistics for firm characteristics and carbon offset retirements at the firm-year level. Panel A summarizes overall firm attributes and aggregate offset activity. The average firm in our sample reports annual sales of \$20.35 billion and emits approximately 3.36 million metric tons of CO₂ per year. Roughly 30% of firm-year observations occur after a firm has obtained SBTi target approval. On average, firms retire offsets from 4.7 distinct projects across 2.6 countries, and the share of firm revenue generated in these retirement countries averages 7.5%.

Offset usage is economically meaningful but modest relative to firms' scale. The average firm retires about 153 thousand tons of CO₂ annually through carbon credits, with an average credit price of \$6.73 and an average MSCI project rating of 2.96 (both measures value-weighted across projects).

Panels B and C split offset retirements based on whether credits are sourced from countries where the firm has an operational presence, defined as positive revenue, or from countries where it does not. Firms retire a larger share of offsets in countries where they operate. These operational-country retirements are also, on average, lower rated by MSCI and lower priced than retirements in other countries. The next subsection examines these geographic and quality patterns more directly, before turning to the regression analysis.

4.3 Aggregate patterns

Figures 1 and 2 characterize the geographic concentration of carbon offset retirements. As shown in Figure 1, firms allocate a significantly larger share of their portfolio to countries within their operational footprint throughout the sample period. While the share of offsets sourced from countries without a revenue footprint has increased over time, retirements remain predominantly concentrated in countries where firms maintain a physical presence.

Figure 2 explores the intensity of this preference by relating retirement location to the scale of foreign operations. For each firm, we compute the average revenue share across the foreign countries where it has any operational presence (excluding the firm's home country to ensure the results are not driven solely by domestic bias), and sort firms into four quartiles based on this firm-level measure. Within each quartile, the figure plots the mean share of the firm's total carbon-offset retirements that are allocated to foreign countries where it operates, with 95% confidence intervals. We find a monotonic relationship:

the mean operational retirement share rises from 68% in the lowest revenue quartile to 83% in the highest quartile. This evidence suggests that the tilt toward local offsets is not merely binary but scales with the intensity of a firm's local economic activity. Together, these aggregate patterns reject the benchmark of geographic fungibility and provide initial support for the local bias predicted by Hypothesis 1.

We next examine whether the quality of retired offsets varies systematically with firms' geographic presence, using MSCI project ratings. Figure 3 reports the distribution of offset retirements across rating quartiles, disaggregated by firms' operational presence in the country where the offset project is located.

Panel A contrasts firms with and without an operational presence in the host country. For each MSCI rating quartile, the figure reports the share of corporate offset retirements accounted for by the two groups; by construction, these shares sum to one within each quartile, so differences in levels partly reflect overall participation in voluntary carbon markets. The salient pattern lies in the quality gradient. Firms with an operational presence account for roughly two-thirds of retirements in the lowest quality quartile (Q1) and less than one-third in the highest quality quartile (Q4). Firms without local operations exhibit the opposite pattern: they account for about one-third of retirements in Q1 but nearly two-thirds in Q4, indicating a pronounced tilt toward higher-rated projects.

Panel B zooms in on firms with an operational presence and further disaggregates them by the intensity of their local footprint and by firm domicile. Foreign firms with a high revenue footprint in the host country (above the sample median revenue share) display the strongest quality gradient, with retirements concentrated in intermediate and lower-rated projects: over 40 percent of their retirements are in Q2 projects, while less than 15 percent are allocated to the highest-rated offsets. In contrast, foreign firms with a low revenue footprint exhibit a relatively balanced allocation across the quality spectrum, with shares in each quartile ranging between roughly 14 and 18 percent of total retired volume. Domestic firms — those headquartered in the same country as the offset project — retire substantially fewer offsets overall, and their allocations are also tilted toward lower-rated projects; the relatively flat profile across rating categories reflects the smaller scale of domestic retirements rather than the absence of a quality gradient.

Taken together, these descriptive patterns indicate that the geographic bias in offset retirements is accompanied by systematic differences in the quality composition of firms' offset portfolios. While carbon offsets are, in principle, fungible across locations, firms with a local operational footprint — particularly those with substantial revenue exposure — tend to concentrate their retirements in lower-rated projects in the countries where they operate. By contrast, firms without local operations disproportionately retire higher-quality credits.

5 Empirical Methodology

To examine how firms’ operational presence relates to the location and quality of carbon offset retirements, we estimate regressions that compare offset choices across projects and locations within the same firm over time.

Firms retiring offsets in countries where they operate may differ systematically from those that do not in size, sector, emissions profile, or sustainability strategy. Similarly, countries attracting more foreign retirements may differ in project availability or baseline quality. To isolate firms’ location-specific preferences, our empirical approach conditions on observable firm and project characteristics and exploits within-firm variation in offset choices across destinations.

Our baseline specification is:

$$S_{kiy} = \delta_1 \text{Operation}_{ij} + \delta_2 \text{Rating}_k + \delta_3 (\text{Operation}_{ij} \times \text{Rating}_k) + \mathbf{X}'_{iy} \boldsymbol{\gamma} + \text{Fixed Effects} + \varepsilon_{kiy}, \quad (17)$$

where S_{kiy} denotes the share of carbon offsets retired by firm i from project k in year y , defined as the volume of credits retired from project k divided by the firm’s total retirements in that year. The indicator Operation_{ij} equals one if firm i has an operational presence in country j , where project k is located. Rating_k captures the ex-post quality rating assigned to project k .

Equation (17) tests whether operational presence changes both the level and the quality composition of firms’ offset allocations. We also distinguish among different forms of exposure to the country of operations, separating domestic firms, foreign firms with more salient host-country operations, and foreign firms with more limited exposure.

The vector \mathbf{X}_{iy} includes firm-level controls such as emissions, revenue, and SBTi target status. All specifications include project classification, vintage, and country-of-retirement-by-year fixed effects to account for supply-side variation. This absorbs time-varying factors common to projects available in the same country-year, including changes in project entry, exit, registry composition, and the local distribution of ratings, and therefore addresses the concerns that firms with operations may simply retire offsets in country-years where the available project menu is systematically different. We report specifications with firm fixed effects as well as our most restrictive specification with firm-by-year fixed effects, which absorb all time-varying firm characteristics and compare offset quality choices within the same firm and year across locations.

Standard errors are clustered at the firm level to account for serial correlation in retirement choices.

Hypothesis 3, which concerns the relationship between prices and project quality, is tested separately using project-level price regressions described in Section 8.

6 Results

Table 4-Panel A reports baseline estimates using MSCI project ratings. The table compares three forms of operational exposure. Columns (1)–(2) define *Operation* broadly, equal to one for domestic firms and foreign firms with positive revenues in the project country. Columns (3)–(4) focus on foreign firms with above-median revenue exposure in the project country; in these columns, $Operation_{ij}$ equals zero for domestic firms, foreign firms with below-median exposure, and firms with no operations. Columns (5)–(6) instead define $Operation_{ij}$ as one for foreign firms with below-median exposure, excluding foreign firm-country pairs with above-median exposure from the sample.

Because each specification interacts *Operation* with *Rating*, the coefficient on *Rating* captures the quality slope for the omitted comparison group in each column. Across all six specifications, this coefficient is positive and statistically significant. Thus, within the omitted comparison group in each column, retirement shares increase with project quality. When location-specific incentives are weaker or absent, project integrity appears to be an important determinant of corporate demand.

The key result is the interaction between *Operation* and *Rating*. In columns (3)–(4), which focus on foreign firms with above-median revenue exposure, the interaction is negative and statistically significant. Thus, while firms in the comparison group allocate larger shares to higher-rated projects, this quality sensitivity is substantially weaker for foreign firms in countries where they have a salient operational footprint. Column (4) makes the pattern clear. The coefficient on *Rating* is 0.048, while the coefficient on $Operation \times Rating$ is -0.079 . The quality slope is therefore positive for the comparison group but negative for foreign firms with above-median revenue exposure in the project country ($0.048 - 0.079 = -0.031$). In other words, within the same firm-year and conditional on country-year supply conditions, operational exposure reverses the usual positive relationship between project quality and retirement shares.

The implied marginal effects show that this is a composition shift across project quality. In column (4), operational presence raises the retirement share by about 4.3 percentage points for projects rated one standard deviation below the mean, has essentially no effect at the mean rating, and lowers the share by about 4.6 percentage points for projects rated one standard deviation above the mean.⁶ Thus, within a firm-year, operational exposure is associated with higher retirement shares for lower-rated local projects and lower shares for higher-rated local projects.

Columns (5)–(6) show weaker evidence for foreign firms with below-median revenue exposure. With firm fixed effects, the interaction between *Operation* and *Rating* is negative and statistically significant, indicating a weaker quality slope for firms operating in these lower-exposure countries than for the

⁶The marginal effect of operational presence at rating r is $\hat{\delta}_1 + \hat{\delta}_3 r$. Using the column (4) estimates, $\hat{\delta}_1 = 0.234$ and $\hat{\delta}_3 = -0.079$, and the project-level MSCI rating mean and standard deviation among projects used by listed firms ($\bar{R} = 2.98$, $\sigma_R = 0.56$; Table 1, Panel B), the implied effects are $0.234 - 0.079(2.42) = 0.043$ at one standard deviation below the mean, $0.234 - 0.079(2.98) = -0.001$ at the mean, and $0.234 - 0.079(3.54) = -0.046$ at one standard deviation above the mean.

comparison group. However, once firm-year fixed effects are included, the interaction falls sharply in magnitude and is no longer statistically significant. Thus, for firms with limited host-country revenue exposure, the association between operational presence and offset quality is not robust when identification comes from comparing the same firm’s portfolio allocation across locations within the same year.⁷

This contrast helps sharpen the interpretation of the results. Although the conceptual framework is written in terms of operational footprint more generally, the negative quality gradient in the data is concentrated among foreign firms with more salient host-country exposure, rather than among domestic firms or foreign firms with limited exposure. We confirm the pattern in an alternative specification reported in Appendix Table A2, where offset use by the three groups — domestic firms, foreign firms with below-median revenue exposure, and foreign firms with above-median revenue exposure — is compared against a fixed control group: retirements by the same firms in countries where they have no operational presence. The negative quality gradient is absent among domestic firms; negative but not robust for foreign firms with below-median revenue exposure; and largest and robust for foreign firms with above-median revenue exposure. Overall, consistent with the goodwill channel, the strongest effects appear among firms that must build and maintain local stakeholder relationships outside their home country, where the reputational return to visibly sourcing offsets locally is likely to be highest.

6.1 Project menus, ex-post ratings, and firm-level estimates

The negative *Operation* \times *Rating* coefficient in Table 4-Panel A admits several alternative interpretations. We address three in turn: first, that the gradient could reflect supply-side differences in the available project menu rather than demand-side preferences; second, that ex-post ratings released in 2024 were not observable to firms at the time of retirement; and third, that the project-level specification could be sensitive to portfolio fragmentation. Each test confirms the baseline pattern.

First, a concern with the baseline interpretation is that the country-of-retirement-by-year fixed effects, while absorbing changes in the cross-country distribution of available credits, cannot fully control for within-country-year differences in the quality mix of available carbon offset projects. The negative coefficient for the interaction, *Operation* \times *Rating*, in Table 4-Panel A could therefore reflect that firms tend to operate in countries with weaker high-quality menus, rather than a reduced sensitivity to quality conditional on the available set of projects. To address this, we construct an external measure of menu quality from supply-side data. We use project-year issuance and retirement volumes for the universe of voluntary-market projects across our registries. For each project we compute the net supply of credits available in each year as cumulative issuances minus cumulative retirements through that year. We then

⁷Appendix Table A1 shows a similar cross-group pattern using BeZero/Sylvera ratings. The negative *Operation* \times *Rating* interaction is concentrated among foreign firms with above-median revenue exposure: it is significant with firm fixed effects (column 3) and remains negative, though imprecise, with firm-by-year fixed effects (column 4). The interaction is small and insignificant for below-median exposure (columns 5–6) and absent in the broader domestic-plus-foreign specification (columns 1–2). The weaker precision is expected given the smaller BeZero/Sylvera coverage.

merge our project-level MSCI ratings and, for each (c,y) pair, compute the volume-weighted average MSCI rating across projects with positive net supply. $HighMenu_{cy}$ equals one when this country-year average exceeds the cross-country median in year y .

Re-estimating the baseline specification with the additional interactions $Operation \times HighMenu$ and $Operation \times Rating \times HighMenu$ (Table 4-Panel B) leaves the headline gradient essentially unchanged: the negative $Operation \times Rating$ coefficient survives in both magnitude and significance, and the triple interaction with $HighMenu$ is small and statistically insignificant across specifications. The negative quality gradient is therefore not explained by cross-country differences in the supply of high-quality projects.

Second, project rating was not directly observed at the time firms retire offsets — MSCI released its ratings only in late 2024, and the BeZero and Sylvera coverage is also largely recent. The negative interaction in Table 4-Panel A should therefore be read first as a *realized* integrity gradient; projects retired in foreign firms’ operational footprints later receive lower ex-post quality ratings. To assess whether this pattern is also present in information observable at the time of choice, we construct an ex-ante proxy based only on harmonized project type (for example, renewable energy, cookstoves, or REDD+), developer status (that is, whether the developer operates only locally or across multiple countries), and registry (such as Verra, Gold Standard, CAR, or ACR). For each (p,d,r) bucket, we compute the retired-volume-weighted share of projects classified as high quality under the MSCI measure and assign that score back to each project.

Re-estimating Table 4-Panel A with this ex-ante observable proxy yields a similar pattern (Appendix Table A3). For foreign firms with an above-median revenue footprint in the project country, the interaction between operational presence and ex-ante quality remains negative and statistically significant. The negative quality gradient is therefore not solely an artifact of ex-post ratings released in 2024; it is also visible across project categories defined by characteristics firms could plausibly observe at the time of retirement.

Third, a concern with the project-level specification in Equation (17) is that the dependent variable is defined at the firm-project level rather than at the level of aggregate firm allocations across exposure and quality groups. In principle, the project-level estimates could therefore be sensitive to portfolio fragmentation. The firm-year-cell specification also provides a more direct test of the paper’s core allocation margin: how a firm divides its annual offset portfolio across locations with different operational relevance and across high- versus low-quality credits. In our setting fragmentation concerns appear limited. As shown in Table 3, the median firm-year involves retirements in only one country and two projects; even at the 75th percentile, firms retire in only three countries and five projects. There is little scope for the main results to be driven mechanically by firms splitting a given country allocation across many small projects.

To verify directly, we collapse retirements into firm-year cells defined jointly by the three-way exposure category $c \in \{Foreign-no-ops, Domestic, Foreign-with-Ops\}$ and a binary quality indicator $q \in \{low, high\}$, and regress the within-firm-year retirement share on cell indicators and their interactions with quality. The aggregated firm-year-cell results, reported in Table 5, confirm the main pattern. Relative to foreign countries with no operational presence, the overall share allocated to foreign countries where the firm operates (summing across the low- and high-quality bins $2\hat{\delta}_2 + \hat{\delta}_5$) is 11.5 percentage points higher in the firm and year fixed-effects specification and 8.9 percentage points higher in the within-firm-year specification. This excess allocation is concentrated almost entirely in below-median-quality projects: the low-quality differential is $\hat{\delta}_2 = 12.0$ (11.6) percentage points in column 1 (column 2), while the high-quality differential, $\hat{\delta}_2 + \hat{\delta}_5$, is near zero; -0.5 percentage points in column 1 and -2.7 percentage points in column 2. The negative coefficients on *Foreign with Ops* \times *High Quality* imply a 12.5- to 14.3-percentage-point shift away from above-median-quality projects within foreign operational markets.⁸ These results are consistent with Hypotheses 1 and 2 and indicate that the project-level coefficients reflect genuine allocation shifts rather than portfolio fragmentation.

Taken together, these three checks rule out the leading alternative interpretations of the baseline gradient. The negative *Operation* \times *Rating* coefficient is not explained by cross-country differences in the supply of high-quality projects, not driven by ex-post information firms could not have observed at the time of retirement, and not an artifact of fragmentation across many small projects. We retain the project-level specification as our main design because it allows us to condition directly on project characteristics and absorb supply-side heterogeneity more finely; the geographic concentration and quality gradient survive in both the project-level and firm-year-cell estimates.

7 Mechanism Heterogeneity and Comparative Statics

The baseline results document a negative quality gradient associated with firms' operational presence. We now examine how the strength of this gradient varies across environments where the marginal benefits of integrity or local signaling plausibly differ.

7.1 Experience in Carbon Markets

We first examine whether the negative relationship between operational presence and offset quality attenuates when informational frictions are likely to be lower. Although project transparency is not directly observable, we proxy for information acquisition using firms' experience in carbon offset markets. Experience is defined in two ways: overall participation in offset markets and repeated activity within a specific country.

⁸The firm \times year fixed effects in column (2) absorb firm-years that retire in a single (q, c) cell, which trims the sample by roughly half. The interaction estimate is robust to retaining these firm-years by zero-filling cells with no retirement and re-estimating on the resulting balanced panel of 17,784 firm-year-cells, where the coefficient on *Foreign with Ops* \times *High Quality* is -0.133 (s.e. 0.021).

Figure 4 illustrates the role of country-level experience. The figure separates retirements by foreign firms with an operational footprint into those occurring during the first two years of activity in a country and those occurring after at least two years. In early years, retirement shares decline steeply with project quality, indicating a pronounced negative gradient. After firms accumulate country-level experience, this gradient becomes noticeably flatter.

Table 6 confirms this attenuation in regression form. Across both definitions of experience, the interaction between *Operation* and *Rating* remains negative but declines meaningfully in magnitude with experience. For example, comparing columns (2) and (4), which differ only in firms' overall offset experience, the interaction coefficient is substantially smaller once firms have accumulated several years of participation. A similar pattern emerges when experience is defined at the firm-country level. These results indicate that greater familiarity with offset markets mitigates, though does not eliminate, the negative quality gradient.

7.2 Retirement Intensity

The conceptual framework implies that the strength of the quality gradient depends on the marginal value firms attach to local offset sourcing relative to environmental integrity. When offsetting represents a larger share of a firm's emissions management strategy, the non-pecuniary benefits associated with the geographic allocation of retirements may carry greater weight in portfolio choice. In such cases, operational presence should exert a stronger influence on the composition of offsets.

To examine this idea, we measure firms' reliance on offsets using *retirement intensity*, defined as the ratio of total offset retirements to total Scope 1 and 2 emissions. The sample is split at the median level of retirement intensity (Table 7, Columns 1-4).

Among firms with high retirement intensity, the *Operation* \times *Rating* interaction is large, negative, and statistically significant: -0.130 in column (3) and -0.128 in column (4). These estimates are more than 60 percent larger in absolute value than the baseline interaction in Table 4, column (4). The composition shift documented in the baseline therefore intensifies for firms that rely more heavily on offsets. By contrast, among firms with low retirement intensity, the interaction is small and not statistically significant (-0.012 and -0.039). These results suggest that the local quality gradient is concentrated among firms for which offsets represent a more meaningful part of emissions management, consistent with geographic visibility becoming more relevant when offsetting is more central to the firm's climate strategy.

7.3 Emission Intensity

Finally, we consider heterogeneity by emission intensity. Firms with higher emissions per unit of revenue plausibly face a higher marginal benefit from credible abatement. When splitting firms at the

median of emission intensity (Table 7, Columns 5-8), the negative quality gradient is strongest among low-emission-intensity firms and substantially attenuated among high-emission-intensity firms.

For firms with greater emissions exposure, portfolio quality is more uniform across locations, whereas firms with lower emissions intensity exhibit a sharper tilt toward lower-rated projects within their operational footprints. This heterogeneity suggests that the relative valuation of environmental integrity versus local signaling varies with the firm's underlying emissions profile.

Overall, these comparative statics indicate that the magnitude of the negative quality gradient varies systematically with proxies for informational frictions and with the relative importance of offsetting within the firm's emissions strategy.

7.4 Supply-Side Matching and Developer Access

A remaining concern is that the negative quality gradient reflects mechanical matching between firms and project developers rather than strategic allocation across projects. In the voluntary carbon market, some international developers operate across multiple jurisdictions and tend to implement larger projects with higher average integrity. If firms without local operations disproportionately source from these developers, observed quality differences could arise from supply-side access rather than location-specific demand.

To address this concern, we re-estimate the baseline specification including an indicator for international developer status and its interaction with operational presence (Table 7, Columns 9-10). The coefficient on international developer status is statistically insignificant, and its interaction with *Operation* is economically small. Crucially, the negative interaction between *Operation* and *Rating* remains quantitatively similar to the baseline and highly significant.

These results indicate that the quality gradient is not driven by differential access to international versus domestic developers. Instead, firms reallocate retirement shares across projects conditional on geography, consistent with demand-side portfolio choices rather than mechanical supply matching.

8 Price Implications and the Decoupling of Prices from Quality

We next examine whether the allocation patterns documented above extend to price formation in the voluntary carbon market. Table 8 reports project-level regressions of average offset prices on project quality (ratings), the share of retirements by firms with a local operational footprint, and their interaction. This analysis directly tests Hypothesis 3, which predicts that the relationship between prices and environmental integrity weakens in markets dominated by firms with local operational exposure.

Across all specifications, project ratings are positively and significantly associated with prices. Higher-rated projects command a premium, indicating that environmental integrity remains an important determinant of valuation.

We then examine whether this price-quality relationship varies with the composition of demand. Columns (2), (5), and (8) show that projects with a larger share of retirements by firms with a local operational footprint tend to have lower average prices. This pattern is consistent with the allocation results, which show that local demand is tilted toward lower-rated projects.

The central test lies in the interaction between project ratings and the share of retirements by firms with local operations. In Columns (4), (7), and (10), this interaction is negative and statistically significant. While higher ratings are associated with higher prices overall, the sensitivity of prices to ratings is weaker for projects whose demand is dominated by firms with an operational presence in the host country. In other words, the sensitivity of prices to project ratings declines in markets where local operational demand is more prominent.

This attenuation is strongest when local demand is driven by foreign firms with high revenue exposure in the project country and is negligible for firms with low exposure (Columns 11-13). The heterogeneity mirrors the allocation results and indicates that the strength of the price-quality relationship depends on the intensity of operational presence.

Taken together, these findings indicate that the association between prices and environmental integrity varies systematically with the composition of buyers. In markets characterized by stronger local operational demand, prices are less tightly aligned with project ratings.

9 Local visibility and the strength of the goodwill channel

The preceding results show that firms disproportionately retire offsets in countries where they operate, but that these local retirements are tilted toward lower-quality projects. This pattern is consistent with the goodwill channel: operational presence raises the private value of supporting local projects, even when they offer weaker environmental integrity. This section provides three additional pieces of evidence on the strength of this channel. First, the local allocation effect and the negative quality gradient are concentrated in countries with higher climate ambition, where climate action is likely to be more visible to stakeholders. Second, the quality gradient is also stronger in weaker-governance environments, where firms may place greater value on local engagement. Third, using subsidiary filings for U.S.-listed firms, we show that offset retirements increase after firms establish subsidiaries in a country, and that these entry-country retirements are lower quality than the same firms' retirements elsewhere.

9.1 Climate ambition and the quality gradient

We measure country-level climate ambition using the Climate Change Performance Index (CCPI), an annually published index produced by Germanwatch. The CCPI evaluates countries' climate mitigation performance on a 0–100 scale, combining indicators on greenhouse gas emissions, renewable energy, energy use, and climate policy.⁹ We merge the CCPI overall score to the retirement-level data at the country-year level and classify project countries into high- and low-climate-ambition groups based on the year-specific median CCPI score. The CCPI covers 63 countries, which account for about 90% of global emissions. In our sample, 29 of the 88 project countries where firms retire offsets have CCPI coverage, and these countries account for roughly two-thirds of retirement observations. The 59 countries without CCPI scores are predominantly developing countries that supply offsets but where listed firms have limited operational presence; nearly half are in Sub-Saharan Africa. This coverage limitation makes the CCPI exercise a more demanding test, since it excludes many arm's-length offset-supplying countries that would otherwise sharpen the contrast between local and non-local sourcing.

Table 9 reports the results. Two patterns reinforce the goodwill interpretation. First, in high-CCPI countries, the omitted comparison group tilts retirement shares toward higher-rated projects: the coefficient on *Rating* is positive under both firm fixed effects and firm-by-year fixed effects (0.056 and 0.080, respectively). In low-CCPI countries, by contrast, the corresponding coefficients are small and statistically insignificant. Thus, in countries where climate action is more salient, project quality appears to matter more for buyers without salient local operational exposure.

Second, this is precisely where firms with operational presence exhibit the strongest negative quality gradient. In high-CCPI countries, the *Operation* \times *Rating* interaction is negative and statistically significant: -0.097 with firm fixed effects and -0.081 with firm-by-year fixed effects. These magnitudes are comparable to the baseline interaction in Table 4, columns (3) and (4). In low-CCPI countries, the interaction is small and statistically insignificant. The negative quality gradient is therefore concentrated in the countries where local climate visibility is likely to be most valuable, consistent with the goodwill channel.¹⁰

Taken together, the CCPI results support the goodwill interpretation. Project integrity appears to matter more in high-climate-ambition countries for firms without salient local operational exposure. But precisely in these countries, operational presence generates the strongest negative quality gradient: exposed firms shift retirements toward lower-rated local projects. This pattern is difficult to reconcile with a pure information channel. If local operations primarily improved project screening, operational

⁹Recent economics and finance research uses the CCPI to capture differences in national climate performance and policy ambition; see, for example, Bolton and Kacperczyk (2023) and Demirgüç-Kunt et al. (2024).

¹⁰Columns (5)–(6) estimate a pooled specification that interacts operational presence and project rating with the standardized CCPI score. The triple interaction, *Operation* \times *Rating* \times *CCPI*, is negative though not statistically significant. While imprecise, the sign is consistent with the split-sample evidence that the negative quality gradient is stronger in higher-CCPI countries.

exposure should strengthen the quality slope in high-CCPI countries. Instead, the evidence suggests that local visibility raises the private value of sourcing offsets within the firm’s operational footprint, reducing the marginal role of environmental integrity in portfolio choice.

9.2 Governance and the goodwill channel

We next examine whether the quality gradient also varies with broader institutional quality. The goodwill channel should be stronger in environments where firms may rely more heavily on visibility, local engagement, and relationship management to maintain stakeholder support. In such settings, the reputational value of supporting local projects may be especially high.

To examine this implication, we measure country-year governance using the World Bank Worldwide Governance Indicators and classify project countries into “low governance” and “high governance” groups based on the sample median in each year. Figure 5 reports the distribution of retirement shares across MSCI rating quartiles, separately for firms with and without operational presence in the project country.

The contrast is sharp. In low-governance project countries (Panel A), firms with an operational presence allocate a substantially larger share of retirements to lower-rated projects and a smaller share to higher-rated projects relative to firms without local operations. The negative quality gradient is pronounced. In high-governance countries (Panel B), by contrast, these differences are markedly attenuated, and the allocation of retirements across rating quartiles is more similar across firms.

This heterogeneity is consistent with the goodwill channel. The link between operational presence and lower-quality offset sourcing is strongest in environments where stakeholder engagement plausibly carries greater value. Combined with the CCPI evidence, the governance patterns suggest that firms’ operational footprints shape not only where offsets are retired, but also the integrity of the projects they choose.

9.3 Subsidiary entry and local offset sourcing

The main analysis measures operational presence using the geographic distribution of firm revenue. This measure has broad coverage, but it may not fully capture the strength of a firm’s local engagement. Firms can generate revenue in a country through exports, distributors, or intermediaries without having a meaningful physical presence. To address this concern, we complement the baseline results with an alternative measure of operational footprint based on subsidiary locations.

We identify subsidiary entry events using Exhibit 21 filings from U.S.-listed firms. Exhibit 21, filed with annual 10-K reports, lists a firm’s significant subsidiaries and their jurisdictions of incorporation. By tracking the first year in which a country appears in a firm’s subsidiary filings, we identify new country-entry events. We collect Exhibit 21 filings for 322 U.S.-listed firms in our offset sample over

2010–2024 and obtain subsidiary information for 308 firms. After standardizing jurisdiction names and excluding tax havens, we identify 2,248 country-entry events across 219 firms and 130 countries. Restricting the sample to countries with voluntary carbon offset projects yields 1,299 firm-country entry pairs across 206 firms.

This setting provides a useful validation exercise for the goodwill channel. Establishing a subsidiary represents a discrete increase in the firm’s legal and physical presence in a country. If local offset sourcing reflects visibility and stakeholder engagement, retirement behavior should shift toward newly entered countries after subsidiary entry. Moreover, if this motive weakens the link between offset demand and environmental integrity, entry-country retirements should be lower quality than the same firm’s retirements elsewhere.

Figure 6 examines the timing of retirement behavior around subsidiary entry. The unit of observation is the firm-country pair, and the dependent variable is the absorbing-state indicator that the firm has ever retired offsets in that country by the year of observation. The event study includes firm-country and year fixed effects, with the year before entry as the omitted category; control units are firm-country pairs in the same set of entry countries where the firm has no operations and never enters. Pre-entry coefficients are close to zero, providing little evidence of differential pre-trends. After entry, the cumulative probability of having retired in the entry country rises monotonically, reaching 4.2 percentage points three years after entry — roughly a 70% increase relative to the pre-entry mean of 6.0% for treated firm-country pairs. Consistent with the country-level evidence in Section 9.1, this response is concentrated in high-CCPI destinations: estimating the same specification separately by climate-ambition bin (Appendix Figure A1) yields a cumulative probability of 9.6 percentage points in high-CCPI countries by the third year after entry, against essentially zero in low-CCPI countries.

We next examine whether offsets retired in entry countries are of lower quality. Because relatively few firms retire offsets in the same country both before and after subsidiary entry, this exercise is not designed as a difference-in-differences test. Instead, we compare the quality of offsets retired in countries where the firm has recently established a subsidiary with the quality of other offsets retired by the same firm in other countries. This within-firm comparison holds fixed the firm’s overall offsetting strategy and asks whether the local component of its portfolio is tilted toward lower-quality projects.

Table 10 reports the quality results. Offsets retired in entry countries are of lower quality than other offsets retired by the same firm elsewhere. This pattern appears both when quality is measured by the share of retirements going to above-median MSCI projects and when it is measured by the average MSCI rating. The gap is largest in the first two years after entry and becomes smaller over time, suggesting that the tilt toward lower-quality local projects is strongest when the firm first establishes a local presence.

Panels B and C address a potential composition concern. The negative quality gradient could arise if firms enter countries where the available offset supply is generally of lower quality. To examine this possibility, we split project countries by their average MSCI quality. The quality gap is concentrated in above-median-quality countries: even in countries where the local offset supply is relatively strong, firms entering through subsidiaries retire lower-quality projects than their own retirements elsewhere. In below-median-quality countries, the estimates are smaller and statistically insignificant. This pattern suggests that the result is not simply driven by firms entering low-quality offset markets.

These results reinforce the interpretation of the baseline findings. The local offset bias and the negative quality gradient are not specific to a revenue-based measure of operational presence; they also emerge when operational footprint is measured by subsidiary entry. After entry, firms are more likely to source offsets from the entry country, but those offsets are lower quality than the same firm's retirements elsewhere. Consistent with the goodwill channel, local presence increases the private value of visible local sourcing, especially when firms first enter a country, without necessarily directing demand toward higher-integrity projects.

10 Conclusion

This paper studies how firms allocate carbon-offset retirements across countries and projects when the climate benefit of abatement is geographically fungible but the private benefits of visible local action are not. In a frictionless offset market, firms seeking credible climate mitigation should allocate demand toward high-integrity projects regardless of location. We show instead that corporate offset portfolios are shaped by firms' operational footprints: foreign firms source more offsets from countries where they operate, and these local retirements are systematically tilted toward lower-quality projects.

Our analysis yields three main findings. First, foreign firms with substantial host-country exposure tilt their operational-country retirements toward lower-rated projects and away from higher-rated ones. Second, this negative local quality gradient is concentrated where visibility incentives are likely to be strongest: in high-climate-ambition countries and in weaker-governance environments. The pattern is attenuated with firm experience, suggesting that information frictions matter, but the dominant heterogeneity points to goodwill rather than screening. Subsidiary-entry events provide complementary within-firm evidence: after entry, firms become more likely to retire offsets in the entry country, and those retirements are lower quality than the same firms' other retirements. Third, the demand distortion has price implications. In project segments dominated by foreign firms with local operations, prices are less informative about project quality than in segments dominated by other foreign buyers.

The results identify a demand-side distortion in markets for global public goods. When buyers receive location-specific reputational benefits, demand can shift toward projects that are privately valuable because they are visible locally, even when they are weaker in environmental integrity. This mechanism

helps explain why improving project certification alone may be insufficient: even conditional on the set of issued credits, buyer incentives can weaken the link between demand, prices, and environmental integrity.

The findings also have implications beyond voluntary carbon markets. Many climate-finance instruments bundle global environmental claims with locally visible corporate action. Green bonds, sustainability-linked finance, supply-chain programs, and offset credits used in emerging compliance frameworks all face the same basic tension: local reputational benefits may be socially valuable, but they should not be conflated with verified global abatement. Separating local engagement claims from climate-mitigation claims—through disclosure of project location, buyer operational exposure, and integrity-adjusted credit use—would preserve incentives for local environmental investment while more closely aligning corporate climate accounting with actual abatement.

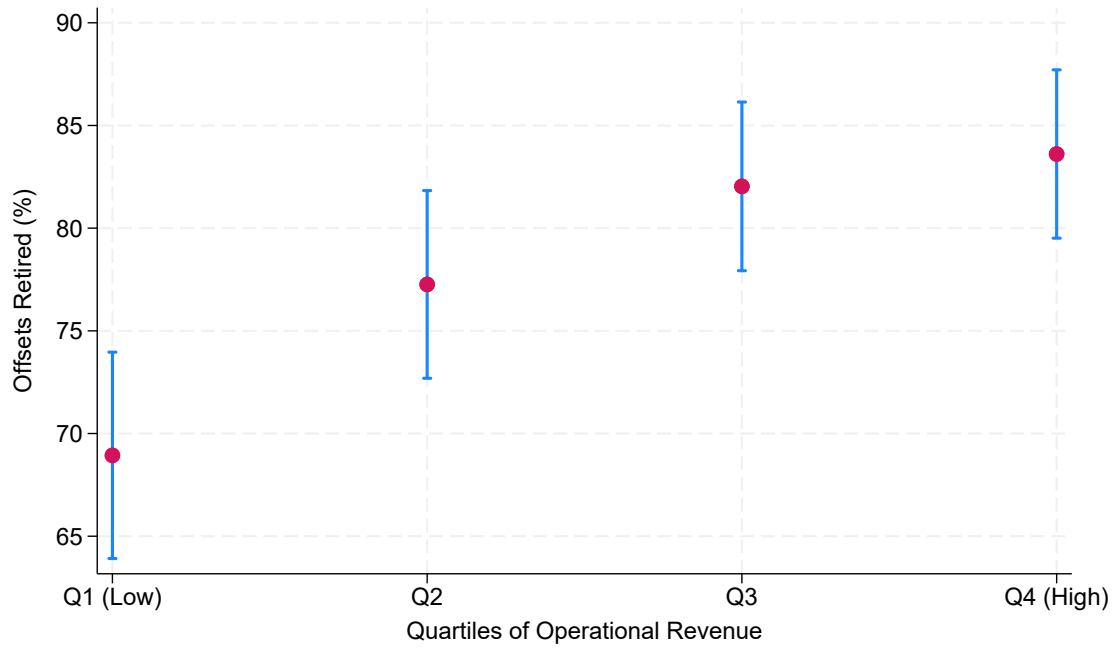
References

- Albuquerque, R., Koskinen, Y., and Zhang, C. (2019). Corporate social responsibility and firm risk: Theory and empirical evidence. *Management Science*, 65(10):4451–4469.
- Beck, A. and Pedraza, A. (2026). Deep and shallow decarbonization in supply chains. *Policy Research Working Paper World Bank*.
- Ben-David, I., Jang, Y., Kleimeier, S., and Viehs, M. (2021). Exporting pollution: Where do multinational firms emit CO₂? *Economic Policy*, 36(107):377–437.
- Berg, F., Ceccarelli, M., Heeb, F., Ivashchenko, A., Rigobón, R., and Zwinkels, R. C. (2025). The market for voluntary carbon offsets. Technical report, SAFE Working Paper.
- Berkeley Carbon Trading Project (2025). Voluntary registry offsets database. <https://gspp.berkeley.edu/berkeley-carbon-trading-project/offsets-database>. Accessed January 3, 2024.
- Bertrand, M., Bombardini, M., Fisman, R., Hackinen, B., and Trebbi, F. (2021). Hall of mirrors: Corporate philanthropy and strategic advocacy. *Quarterly Journal of Economics*, 136(4):2413–2465.
- Bertrand, M., Bombardini, M., Fisman, R., and Trebbi, F. (2020). Tax-exempt lobbying: Corporate philanthropy as a tool for political influence. *American Economic Review*, 110(7):2065–2102.
- Bolton, P. and Kacperczyk, M. (2023). Global pricing of carbon-transition risk. *The Journal of Finance*, 78(6):3677–3754.
- Broccardo, E., Hart, O., and Zingales, L. (2022). Exit versus voice. *Journal of Political Economy*, 130(12):3101–3145.
- Calel, R., Colmer, J., Dechezleprêtre, A., and Glachant, M. (2025). Do carbon offsets offset carbon? *American Economic Journal: Applied Economics*, 17(1):1–40.

- Colmer, J., Martin, R., Muûls, M., and Wagner, U. J. (2025). Does pricing carbon mitigate climate change? firm-level evidence from the european union emissions trading system. *Review of Economic Studies*, 92(3):1625–1660.
- Coval, J. D. and Moskowitz, T. J. (1999). Home bias at home: Local equity preference in domestic portfolios. *The Journal of Finance*, 54(6):2045–2073.
- Coval, J. D. and Moskowitz, T. J. (2001). The geography of investment: Informed trading and asset prices. *Journal of political Economy*, 109(4):811–841.
- Demirgüç-Kunt, A., Pedraza, A., Pulga, F., and Ruiz-Ortega, C. (2024). Global bank lending under climate policy. *IMF Economic Review*, 72(2):858–901.
- Fabra, N. and Reguant, M. (2014). Pass-through of emissions costs in electricity markets. *American Economic Review*, 104(9):2872–2899.
- Fowlie, M. and Reguant, M. (2022). Mitigating emissions leakage in incomplete carbon markets. *Journal of the Association of Environmental and Resource Economists*, 9(2):307–343.
- Gill-Wiehl, A., Kammen, D. M., and Haya, B. K. (2024). Pervasive over-crediting from cookstove offset methodologies. *Nature Sustainability*, 7(2):191–202.
- Godfrey, P. C., Merrill, C. B., and Hansen, J. M. (2009). The relationship between corporate social responsibility and shareholder value: An empirical test of the risk management hypothesis. *Strategic Management Journal*, 30(4):425–445.
- Gunningham, N., Kagan, R. A., and Thornton, D. (2004). Social license and environmental protection: why businesses go beyond compliance. *Law & Social Inquiry*, 29(2):307–341.
- Hamilton, K., Sjardin, M., Marcello, G., and Xu, J. (2012). State of the voluntary carbon markets 2012.
- Hamilton, K., Sjardin, M., Shapiro, A., and Marcello, G. (2010). State of the voluntary carbon markets 2010.
- Hartzmark, S. M. and Shue, K. (2025). Counterproductive sustainable investing: The impact elasticity of brown and green firms. *Journal of Finance*. Forthcoming.
- Heeb, F., Kölbl, J. F., Paetzold, F., and Zeisberger, S. (2023). Do investors care about impact? *Review of Financial Studies*, 36(5):1737–1787.
- Henisz, W. J., Dorobantu, S., and Nartey, L. J. (2014). Spinning gold: The financial returns to stakeholder engagement. *Strategic Management Journal*, 35(12):1727–1748.
- Kim, S., Li, T., and Wu, Y. (2023). Carbon offsets: Decarbonization or transition-washing? *Review of Financial Studies*, 36(11):4585–4630.
- Lins, K. V., Servaes, H., and Tamayo, A. (2017). Social capital, trust, and firm performance: The value of corporate social responsibility during the financial crisis. *Journal of Finance*, 72(4):1785–1824.

- Manshadi, V., Monachou, F., and Morgenstern, I. (2025). Offsetting carbon with lemons: Adverse selection and certification in the voluntary carbon market. *Available at SSRN 5792202*.
- MSCI (2024). Carbon project ratings. <https://www.msci.com/data-and-analytics/carbon-markets/carbon-project-ratings>. Accessed December 20, 2024.
- Oehmke, M. and Opp, M. M. (2025). A theory of socially responsible investment. *Review of Economic Studies*, 92(2):1193–1225.
- Pástor, L., Stambaugh, R. F., and Taylor, L. A. (2025). Sustainable investing. *Annual Review of Financial Economics*, 17.
- Schneider, L., La Hoz Theuer, S., Broekhoff, D., Füssler, J., Kizzier, K., Korberg, M., Lazarus, M., and Howard, A. (2019). Environmental integrity of international carbon market mechanisms under the paris agreement. *Climate Policy*, 19(3):386–400.
- Shapiro, J. S. and Walker, R. (2025). Is air pollution regulation too lenient? Evidence from US offset markets. *American Economic Review*, 115(9):3058–3080.
- West, T. A., Wunder, S., Sills, E. O., Börner, J., Rifai, S. W., Neidermeier, A. N., Frey, G. P., and Kontoleon, A. (2023). Action needed to make carbon offsets from forest conservation work for climate change mitigation. *Science*, 381(6660):873–877.
- World Bank (2025). *State and Trends of Carbon Pricing 2025*. World Bank, Washington, DC. License: Creative Commons Attribution CC BY 3.0 IGO.

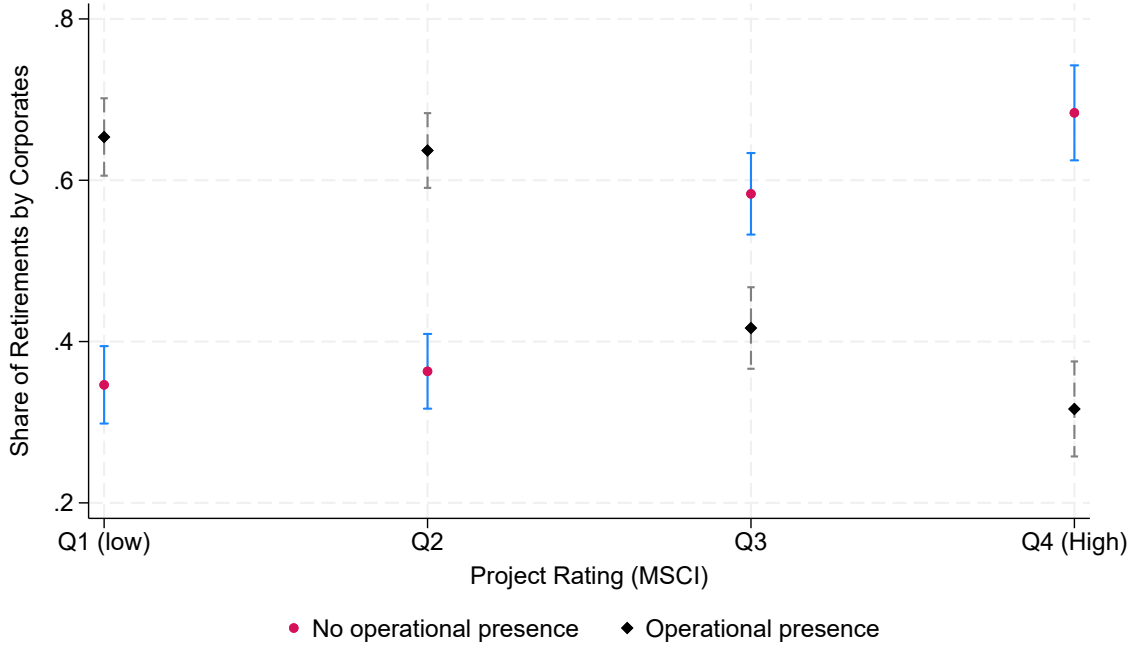
Figure 2: Revenue Exposure and Concentration of Offset Retirements



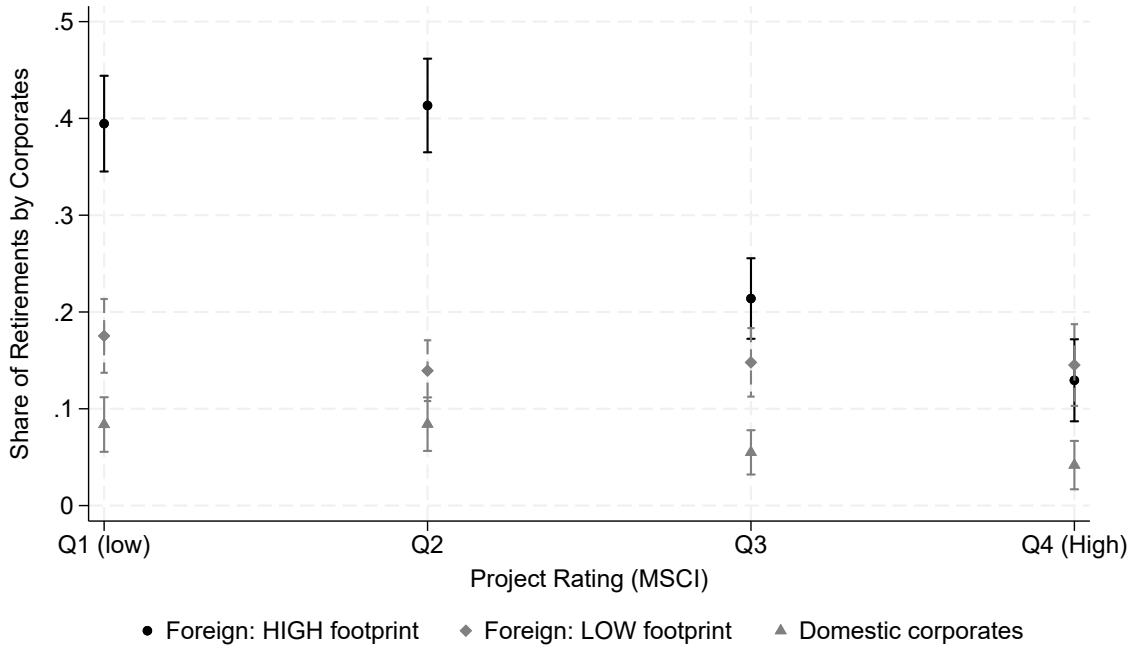
Note: This figure plots the share of carbon offsets that firms retire in foreign countries where they have an operational presence, by quartiles of firm-level foreign revenue exposure. The unit of observation is the firm. For each firm, we compute the average revenue share across foreign countries in which the firm reports positive revenue, excluding the firm's home country. Firms are then sorted into four equal-sized groups based on this measure. Within each quartile, points show the cross-firm mean share of total offset retirements allocated to foreign operational countries. Vertical bars indicate 95 percent confidence intervals based on the cross-firm standard error.

Figure 3: Carbon offset retirements by project quality and firm presence

Panel A: Operational presence vs. no operational presence

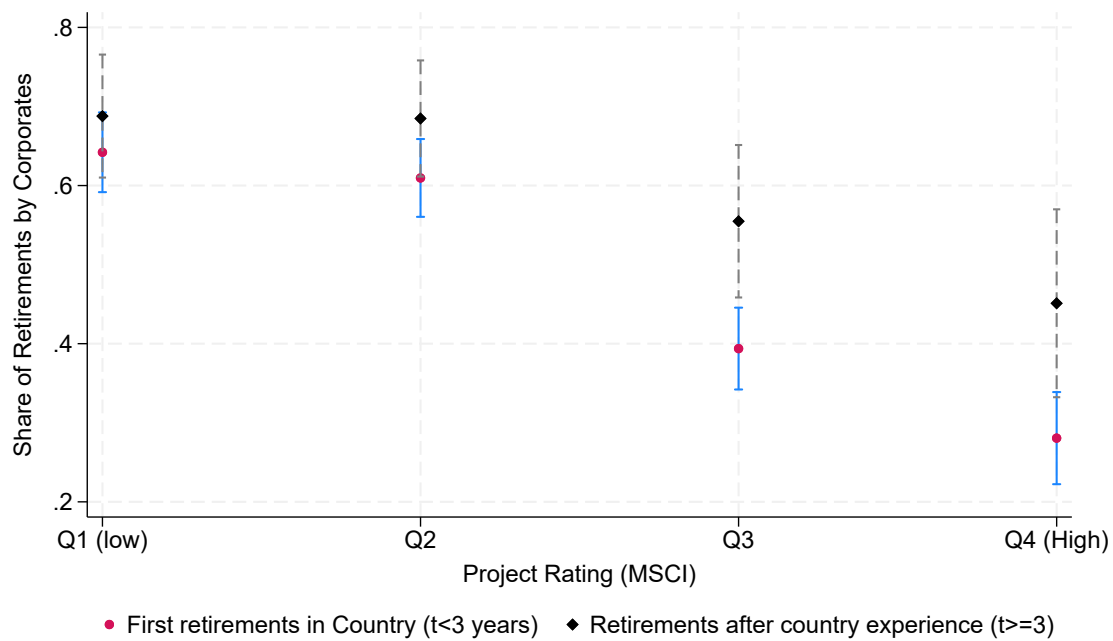


Panel B: Disaggregation among firms with operational presence



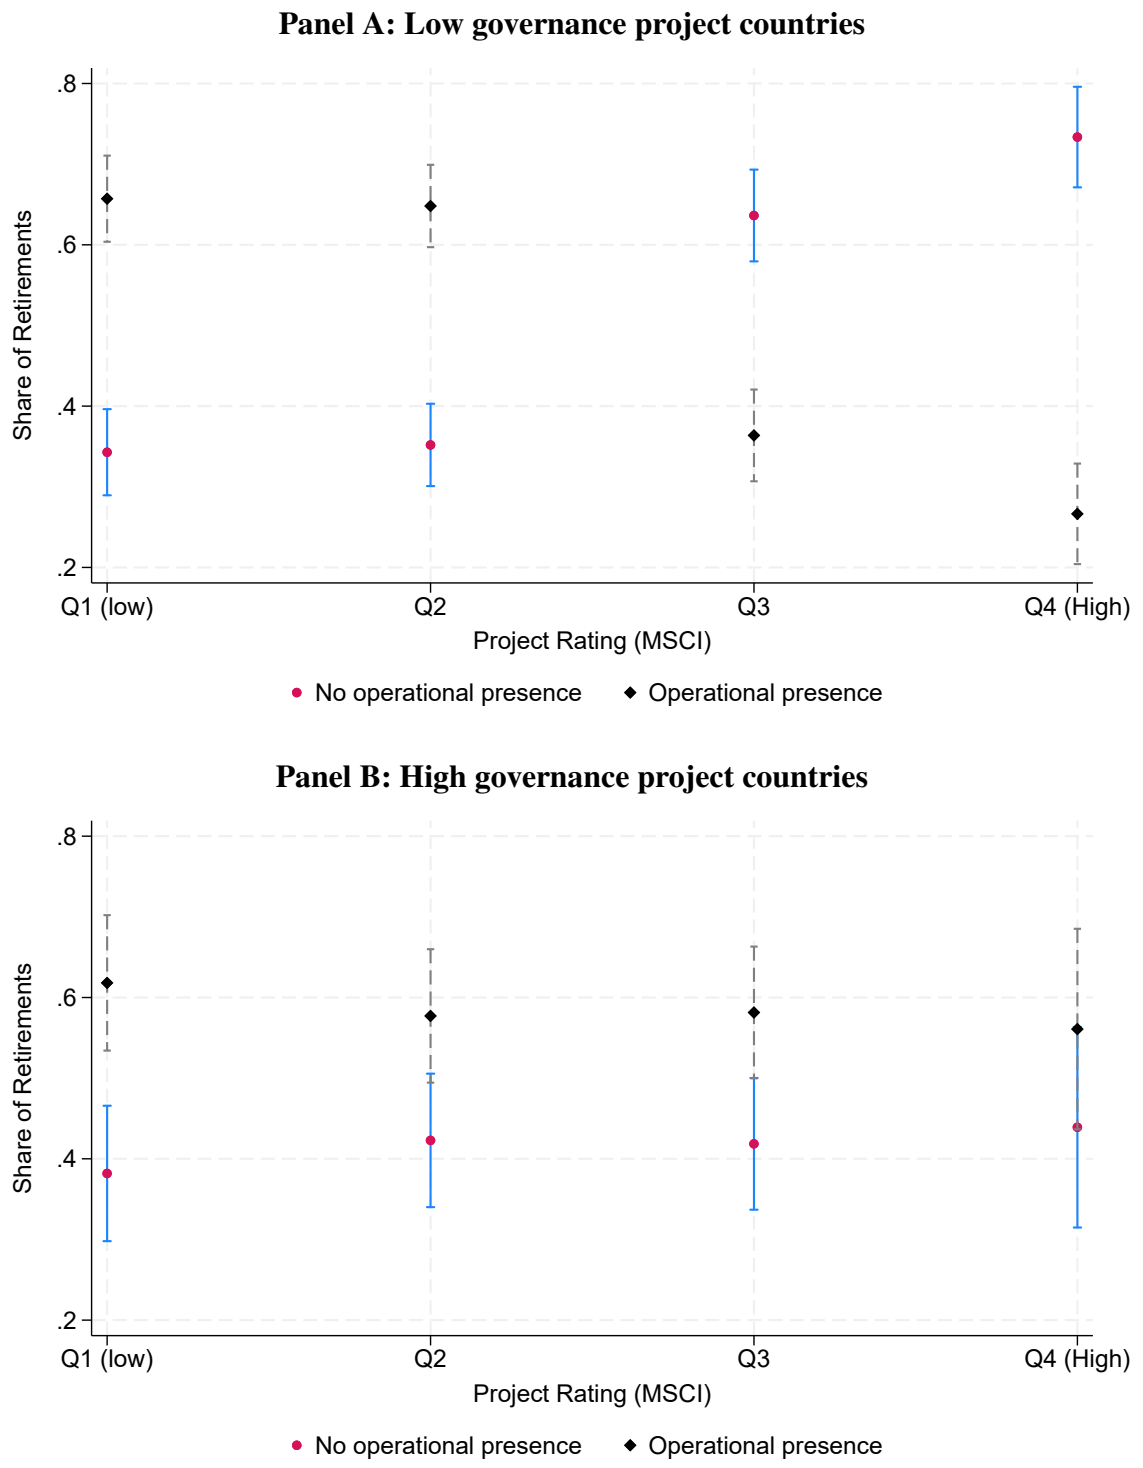
Notes: Panel A reports the share of carbon offset retirements by publicly listed firms across quartiles of project quality, as measured by MSCI project ratings. Firms are classified based on whether they have an operational presence in the country where the offset project is located, defined as positive firm revenues in that country at the time of retirement. Panel B focuses on the subset of firms with an operational presence and further disaggregates them into foreign firms with a high revenue footprint, foreign firms with a low revenue footprint (relative to the sample median), and domestic firms domiciled in the project country. In both panels, markers represent mean shares within each rating quartile, and vertical bars denote 95% confidence intervals.

Figure 4: Offset Quality and Firm Experience within Project Countries



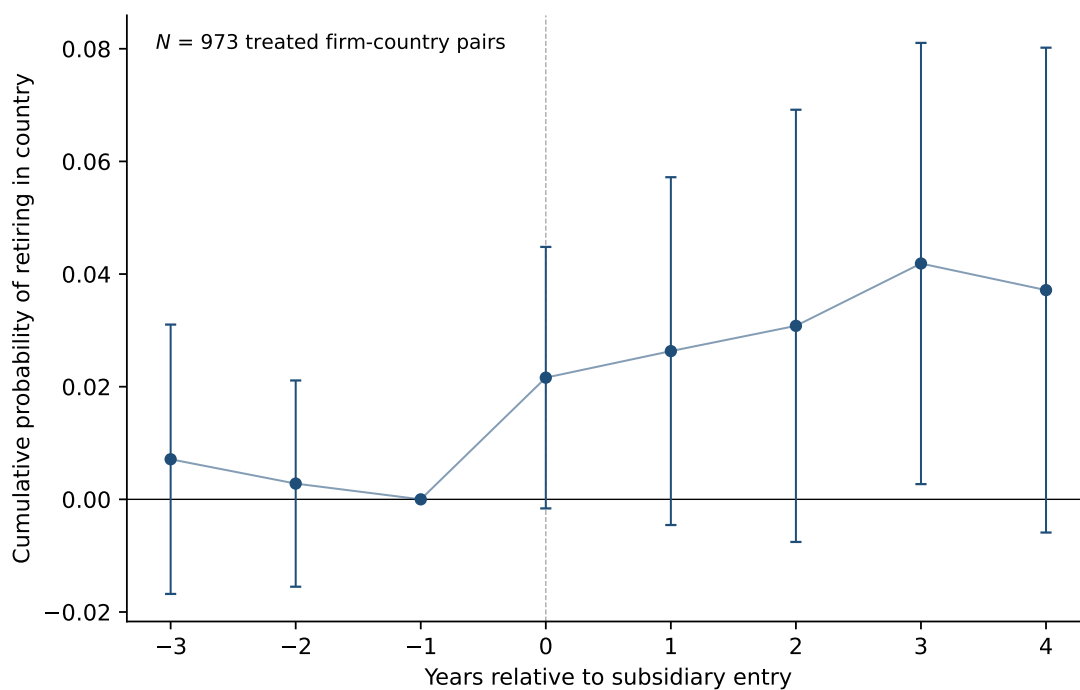
Note: The figure reports the share of carbon offset retirements by foreign firms with an operational presence in a country, disaggregated by project quality (MSCI rating quartiles) and by firm experience in that country. “First retirements” correspond to offsets retired within the first two years in which a firm retires any project in a given country. “Later retirements” correspond to offsets retired after at least two years of country-level experience. Points represent mean shares within each rating quartile, and vertical bars indicate 95% confidence intervals.

Figure 5: Carbon offset retirements by project quality and firm presence: High vs. Low governance countries



Notes: Share of carbon offset retirements by publicly listed firms across quartiles of project quality, as measured by MSCI project ratings. Firms are classified based on whether they have an operational presence in the country where the offset project is located, defined as positive firm revenues in that country at the time of retirement. Panel A is the sample of countries with below median WGI indicator and Panel B above median. In both panels, markers represent mean shares within each rating quartile, and vertical bars denote 95% confidence intervals.

Figure 6: Subsidiary entry and offset retirement in entry country



Notes: Event-study estimates from the firm-country-year regression $y_{ict} = \sum_{k \neq -1} \delta_k \mathbf{1}\{t - g_{ic} = k\} + \alpha_{ic} + \gamma_t + \varepsilon_{ict}$, where $y_{ict} = \mathbf{1}\{\exists s \leq t : q_{ics} > 0\}$ is the absorbing-state indicator that firm i has retired offsets in country c by year t . Treated units are firm-country pairs (i, c) with an SEC Exhibit 21 subsidiary entry in c between 2010 and 2022, with g_{ic} the entry year; control units are firm-country pairs in the same set of entry countries where the firm has no operations and never enters. α_{ic} are firm-country fixed effects; γ_t are year fixed effects. The omitted period is $k = -1$. Standard errors are clustered at the firm level. Error bars show 95% confidence intervals.

Table 1: Carbon Offset Projects: Summary Statistics

Data are from the Voluntary Registry Offsets Database developed by the Berkeley Carbon Trading Project. The sample includes projects with retirements between 2009 and 2024 from four major voluntary carbon registries: American Carbon Registry, Climate Action Reserve, Gold Standard, and Verra. The full sample comprises the universe of retirement transactions reported in the database over this period. "Used by listed firms" refers to projects from which at least one ton of CO₂ was retired by a publicly listed company. Panel A reports the number of unique projects and countries, as well as data coverage across rating agencies and prices. Panel B reports project-level means, standard deviations, minimums, and maximums.

	Full Sample				Used by Listed Firms			
Panel A. Number of projects and sample size								
Projects	3,565				1,889			
Countries	103				88			
Projects with MSCI ratings	2,549				1,395			
Projects with BeZero ratings	375				303			
Projects with Sylvera ratings	358				295			
Projects with price data	1,049				750			
Total credits retired (million tons)	759				659			
	Mean	Std.	Min	Max	Mean	Std.	Min	Max
Panel B. Summary statistics by project								
MSCI rating	2.99	0.58	1.4	4.4	2.98	0.56	1.4	4.4
Price (USD)	6.35	6.68	0.11	115.10	6.34	7.20	0.11	115.10
Credits retired (1,000 tons)	213	847	0	24,454	349	1,136	0	24,454

Table 2: Top Countries by Corporate Origin and Carbon Project Location

The table reports the top 15 countries by total carbon offset retirements, measured in million tons of CO₂ (MtCO₂). The left panel ranks countries by the headquarters location of parent firms retiring offsets, while the right panel ranks countries by the location of carbon offset projects. Firm counts refer to the number of listed firms headquartered in each country, and project counts refer to the number of distinct offset projects located in each country.

Rank	Parent Company HQ			Carbon Project		
	Country	Firms	Retirements (MtCO ₂)	Country	Projects	Retirements (MtCO ₂)
1	United States	325	74.55	India	306	34.64
2	Germany	74	26.97	United States	331	32.24
3	United Kingdom	128	22.64	China	282	32.14
4	Japan	94	22.25	Indonesia	15	20.72
5	Australia	81	15.52	Brazil	105	13.94
6	Italy	34	13.77	Peru	18	13.23
7	France	56	8.78	Colombia	25	10.30
8	Switzerland	36	6.59	Congo	2	8.27
9	China	48	5.86	Turkey	130	6.70
10	Brazil	61	5.86	Zambia	6	6.18
11	Colombia	8	3.78	Kenya	39	5.17
12	Canada	43	2.99	South Korea	5	4.51
13	Hong Kong	24	2.53	Zimbabwe	2	3.17
14	Sweden	51	1.87	Cambodia	6	2.77
15	Spain	39	1.27	Tanzania	5	2.57

Table 3: Listed companies retiring carbon credits: Summary statistics

This table reports summary statistics for firm-level characteristics and carbon offset retirements. Panel A summarizes firm characteristics and overall offset use. Panels B and C report statistics separately for credits retired in countries where firms operate and in other countries. All variables are reported at the firm–year level.

	Mean	Std.	Min	P25	Median	P75	Max
Panel A. Firms' summary statistics							
Sales (billion USD)	20.35	44.41	0.00	1.48	5.36	18.58	588.37
Emissions (million tons)	3.36	12.29	0.00	0.01	0.08	0.78	120.70
Emission intensity	198.75	519.80	0.07	3.45	17.93	102.61	3188.17
SBTi target approved	0.30	0.44	0.00	0.00	0.00	0.88	1.00
Number of countries	2.62	2.84	1.00	1.00	1.00	3.00	30.00
Number of projects	4.71	8.50	1.00	1.00	2.00	5.00	118.00
Share of revenue (%)	7.51	14.63	0.10	0.78	2.21	7.42	100.00
Retired offsets (1,000 tons)	152.80	835.86	0.00	0.20	3.00	29.16	12,982.65
MSCI rating	2.96	0.46	1.60	2.69	2.92	3.24	4.30
Price (USD)	6.73	5.12	0.32	3.50	6.03	8.04	115.10
Panel B. Carbon credits retired in countries with operations							
Share of retirements (%)	61.85	44.67	0.00	0.00	96.29	100.00	100.00
MSCI rating	2.93	0.48	1.60	2.63	2.90	3.24	4.30
Price (USD)	6.34	6.86	0.35	2.61	5.17	7.26	115.10
Panel C. Carbon credits retired in other countries							
Share of retirements (%)	38.15	44.67	0.00	0.00	3.71	100.00	100.00
MSCI rating	3.03	0.46	1.60	2.71	3.00	3.30	4.30
Price (USD)	7.45	5.97	0.32	3.42	6.60	9.08	46.00

Table 4: Operational Presence, Revenue Footprint, and Offset Quality

This table reports estimates from the following specification:

$$S_{kiy} = \delta_1 \text{Operation}_{ij} + \delta_2 \text{Rating}_k + \delta_3 (\text{Operation}_{ij} \times \text{Rating}_k) + \mathbf{X}'_{iy} \boldsymbol{\gamma} + \text{Fixed Effects} + \varepsilon_{kiy}.$$

The dependent variable S_{kiy} is the share of firm i 's annual carbon offset retirements allocated to project k in year y . Operation_{ij} is a dummy equal to one if firm i has operations in the country j where project k is located. In columns (1)-(2), Operation_{ij} equals one for firms with operations in the project country, including both foreign firms with positive revenues and domestic firms headquartered in that country. In columns (3)-(4), Operation_{ij} equals one for foreign firms with a revenue footprint above the sample median in the project country. In columns (5)-(6), Operation_{ij} equals one for foreign firms with a below-median revenue footprint, and the sample excludes foreign firm-country pairs with above-median revenue footprint. Rating_k denotes the project's MSCI rating. Panel A reports the baseline specification above. Panel B extends the baseline with the additional terms $\text{Rating} \times \text{HighMenu}$, $\text{Operation} \times \text{HighMenu}$, and $\text{Operation} \times \text{Rating} \times \text{HighMenu}$, where HighMenu_{cy} equals one if country c 's year- y volume-weighted average MSCI rating — computed across rated projects with positive net supply (cumulative issuances minus cumulative retirements) — exceeds the cross-country median in year y . All specifications include project classification, vintage, and country-of-retirement-by-year fixed effects, as well as firm-level controls $\text{Log}(\text{Emissions})$, $\text{Log}(\text{Revenue})$, and Target Approved . Standard errors are clustered at the firm level and reported in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	Operational Presence Domestic + Foreign		Foreign Above median revenue		Foreign Below median revenue	
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Baseline</i>						
OPERATION	0.143** (0.072)	0.048 (0.072)	0.244*** (0.075)	0.234*** (0.082)	0.199** (0.085)	0.016 (0.089)
RATING	0.070*** (0.021)	0.042** (0.021)	0.061*** (0.017)	0.048*** (0.017)	0.080*** (0.021)	0.052** (0.024)
OPERATION \times RATING	-0.051** (0.024)	-0.026 (0.024)	-0.085*** (0.025)	-0.079*** (0.029)	-0.074*** (0.028)	-0.024 (0.030)
R-squared	0.550	0.477	0.550	0.478	0.574	0.501
<i>Panel B: With project menu controls</i>						
OPERATION	0.163** (0.083)	0.078 (0.082)	0.249*** (0.081)	0.249*** (0.084)	0.250*** (0.095)	0.089 (0.106)
RATING	0.068*** (0.025)	0.050** (0.025)	0.054*** (0.020)	0.050** (0.020)	0.082*** (0.024)	0.065** (0.028)
OPERATION \times RATING	-0.059** (0.029)	-0.043 (0.029)	-0.086*** (0.028)	-0.085*** (0.030)	-0.097*** (0.033)	-0.058 (0.038)
RATING \times HighMenu	0.013 (0.038)	-0.005 (0.038)	0.023 (0.030)	-0.005 (0.032)	0.010 (0.038)	-0.030 (0.040)
OPERATION \times HighMenu	-0.202 (0.155)	-0.038 (0.169)	-0.444* (0.269)	-0.134 (0.329)	-0.087 (0.178)	-0.103 (0.196)
OPERATION \times RATING \times HighMenu	0.063 (0.050)	0.029 (0.055)	0.127 (0.086)	0.047 (0.105)	0.044 (0.057)	0.057 (0.063)
R-squared	0.550	0.478	0.551	0.478	0.575	0.503
Observations	6,268	5,206	6,268	5,206	4,902	4,030
Firm FE	Yes	No	Yes	No	Yes	No
Firm \times Year FE	No	Yes	No	Yes	No	Yes

Table 5: Aggregated Firm-Year-Cell Results: Three-Category Exposure Buckets

This table reports estimates from the following firm-year-cell specification:

$$S_{iqcy} = \delta_1 \mathbb{1}(c = 1) + \delta_2 \mathbb{1}(c = 2) + \delta_3 q + \delta_4 (\mathbb{1}(c = 1) \times q) + \delta_5 (\mathbb{1}(c = 2) \times q) + \text{Fixed Effects} + \varepsilon_{iqcy}.$$

The dependent variable S_{iqcy} is the share of firm i 's annual carbon offset retirements in year y allocated to cell (q, c) . The quality indicator q equals one if the project is above the median MSCI rating and zero otherwise. The exposure variable c is defined using three mutually exclusive categories: $c = 0$ for foreign countries where the firm has no operations, $c = 1$ for the firm's domestic country, and $c = 2$ for foreign countries where the firm has operations. Each firm-year contributes up to six cells: $(q, c) \in \{0, 1\} \times \{0, 1, 2\}$. The omitted category is $c = 0$. Standard errors are clustered at the firm level and reported in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1)	(2)
Domestic (c=1)	0.003 (0.066)	-0.024 (0.105)
Foreign with Ops (c=2: any operational presence)	0.12*** (0.029)	0.116*** (0.039)
High Quality (q=1: above-median MSCI rating)	0.009 (0.026)	-0.006 (0.032)
Domestic x High Quality (c=1 x q)	0.056 (0.068)	0.1 (0.121)
Foreign with Ops x High Quality (c=2 x q)	-0.125*** (0.035)	-0.143*** (0.044)
Observations	3486	1826
R-squared	0.38	0.107
Firm FE	Yes	No
Year FE	Yes	No
Firm \times Year FE	No	Yes

Table 6: Operational Presence, Experience, and Offset Quality

This table reports estimates from the following specification:

$$S_{kiy} = \delta_1 \text{Operation}_{ij} + \delta_2 \text{Rating}_k + \delta_3 (\text{Operation}_{ij} \times \text{Rating}_k) + \mathbf{X}'_i \gamma + \text{Fixed Effects} + \varepsilon_{kiy}.$$

The dependent variable S_{kiy} is the share of firm i 's annual carbon offset retirements allocated to project k in year y . Operation_{ij} is a dummy equal to one if firm i has operations in the country j where project k is located. Columns (1)-(4) split the sample based on firms' overall experience using carbon offsets, regardless of location. Firms are classified as inexperienced if they have retired offsets for fewer than three years and as experienced if they have retired offsets for three or more years. Columns (5)-(8) define experience at the firm-country level, based on whether firm i has retired offsets in country j for fewer than three years or for three or more years. Rating_k denotes the project MSCI rating. All specifications include firm log sales, log emissions, an indicator if the firm has SBTi target approved as controls, and project classification, vintage, and country-of-retirement-by-year fixed effects. Standard errors are clustered at the firm level and reported in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	Firm experience using offsets				Firm-country experience using offsets			
	$t < 3$ years		$t \geq 3$ years		$t < 3$ years		$t \geq 3$ years	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
OPERATION	0.273*** (0.063)	0.307*** (0.090)	0.207*** (0.039)	0.152** (0.065)	0.255*** (0.064)	0.225*** (0.054)	0.126 (0.121)	0.115 (0.177)
RATING	0.073*** (0.018)	0.076*** (0.027)	0.050*** (0.013)	0.022 (0.013)	0.057*** (0.014)	0.060*** (0.018)	0.055 (0.036)	0.024 (0.026)
OPERATION \times RATING	-0.103*** (0.020)	-0.115*** (0.032)	-0.061*** (0.014)	-0.049* (0.024)	-0.090*** (0.019)	-0.084*** (0.020)	-0.033 (0.047)	-0.036 (0.062)
Constant	0.675* (0.357)	0.028 (0.080)	0.189 (0.233)	0.112*** (0.039)	0.516** (0.203)	0.037 (0.055)	0.800* (0.418)	0.105 (0.076)
Observations	2,936	2,399	3,300	2,908	5,048	4,123	1,248	1,003
R-squared	0.620	0.453	0.482	0.446	0.546	0.471	0.555	0.427
Year FE	Yes	No	Yes	No	Yes	No	Yes	No
Firm FE	Yes	No	Yes	No	Yes	No	Yes	No
Firm \times Year FE	No	Yes	No	Yes	No	Yes	No	Yes

Table 7: Retirement Intensity, Operational Presence, and the Quality Gradient

This table reports estimates from variants of the following specification:

$$S_{kiy} = \delta_1 \text{Operation}_{ij} + \delta_2 \text{Rating}_k + \delta_3 (\text{Operation}_{ij} \times \text{Rating}_k) + \mathbf{X}'_i \gamma + \text{FE} + \varepsilon_{kiy}.$$

The dependent variable S_{kiy} is the share of firm i 's carbon offset retirements in year y allocated to project k . Operation_{ij} is an indicator equal to one if firm i has an operational presence in country j , where project k is located. Rating_k denotes the environmental quality rating of project k . Control variables include indicators for international project developers, firm sales, emissions, and target commitments. All specifications include classification fixed effects, project vintage fixed effects, and country-of-retirement fixed effects. Columns alternate between year and firm fixed effects and firm-by-year fixed effects, as reported. Standard errors are clustered at the country-of-retirement level. In columns (1)-(4), firms are split by *retirement intensity*, defined as total carbon offset retirements divided by total emissions at the firm level. In columns (5)-(8), firms are split by *emission intensity*, defined as yearly emissions divided by firm revenue. Columns labeled "Low" ("High") report estimates for firms below (above) the median of each measure. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Retirement Intensity				Emission Intensity				All Firms	
	Low		High		Low		High		(9)	(10)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)		
OPERATION	0.043 (0.061)	0.153** (0.059)	0.370*** (0.047)	0.326*** (0.098)	0.369*** (0.109)	0.347*** (0.122)	0.070 (0.048)	0.129*** (0.044)	0.222*** (0.042)	0.235*** (0.040)
RATING	0.005 (0.017)	0.024 (0.020)	0.082*** (0.009)	0.063*** (0.018)	0.052*** (0.019)	0.055*** (0.019)	0.054*** (0.014)	0.040* (0.020)	0.051*** (0.010)	0.046*** (0.014)
OPERATION × RATING	-0.012 (0.020)	-0.039 (0.024)	-0.130*** (0.016)	-0.128*** (0.036)	-0.125*** (0.035)	-0.134*** (0.040)	-0.025 (0.022)	-0.037* (0.019)	-0.076*** (0.013)	-0.084*** (0.015)
International Developer									0.001 (0.010)	0.010 (0.013)
OPERATION × Int. Developer									-0.015 (0.022)	0.001 (0.029)
Observations	2,871	2,327	3,491	2,976	3,080	2,555	3,277	2,752	6,371	5,316
R^2	0.559	0.453	0.503	0.457	0.492	0.441	0.563	0.467	0.519	0.443
Year FE	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No
Firm FE	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No
Firm × Year FE	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes

Table 8: Prices, Buyer Composition, and Project Quality

This table reports estimates from the following specification:

$$P_{ky} = \alpha + \delta_1 \text{Rating}_k + \delta_2 \text{Share}_{ky} + \delta_3 (\text{Rating}_k \times \text{Share}_{ky}) + \mu_{c(k)} + \lambda_y + \varepsilon_{ky}.$$

The dependent variable P_{ky} is the average transaction price of carbon offsets for project k in year y . Rating_k denotes the environmental quality rating of project k . The variable Share_{ky} measures the share of total offset retirements in project k and year y accounted for by different groups of firms. In columns (2)–(4), Share_{ky} is computed using all firms with operations in the project country, including both domestic firms headquartered in that country and foreign firms with positive revenues there. In columns (5)–(7), Share_{ky} is computed using only domestic firms. In columns (8)–(10), Share_{ky} is computed using only foreign firms whose revenue footprint in the project country is above the sample median. In columns (11)–(13), Share_{ky} is computed using only foreign firms whose revenue footprint in the project country is below the sample median. All specifications include classification fixed effects $\mu_{c(k)}$ and year fixed effects λ_y . Standard errors are clustered at the project level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Operation (Domestic + Foreign)			Domestic Firms			Foreign: High Revenue			Foreign: Low Revenue			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
Rating	1.962*** (0.722)		2.032*** (0.722)	3.167*** (0.894)		1.962*** (0.721)	2.165*** (0.764)		2.034*** (0.722)	2.728*** (0.861)		1.962*** (0.717)	1.501*** (0.576)
Share		-0.736* (0.398)	-0.846** (0.374)	4.795* (2.667)	-1.298* (0.684)	-1.830*** (0.561)	5.352 (3.411)	-0.790*** (0.297)	-0.875*** (0.320)	7.375*** (2.580)	0.301 (0.423)	0.322 (0.431)	-4.811 (4.066)
Share × Rating				-1.946** (0.986)			-2.465** (1.238)			-2.852*** (0.943)			1.789 (1.512)
Constant	-0.208 (2.051)	6.067*** (0.317)	0.037 (2.075)	-3.264 (2.491)	5.749*** (0.201)	-0.088 (2.043)	-0.674 (2.160)	5.827*** (0.218)	-0.243 (2.033)	-2.251 (2.422)	5.573*** (0.200)	-0.292 (2.084)	1.045 (1.647)
Observations	713	845	713	713	845	713	713	845	713	713	845	713	713
R ²	0.436	0.373	0.442	0.448	0.372	0.443	0.447	0.372	0.440	0.451	0.369	0.437	0.441

Table 9: Country Climate Ambition and Firms' Retirements

This table reports estimates from variants of the following specification:

$$S_{kijy} = \delta_1 \text{Operation}_{ij} + \delta_2 \text{Rating}_k + \delta_3 (\text{Operation}_{ij} \times \text{Rating}_k) + \mathbf{X}'_{iy} \gamma + \text{FE} + \varepsilon_{kijy}.$$

The dependent variable S_{kijy} is the share of firm i 's carbon offset retirements in year y allocated to project k located in country j . Operation_{ij} is an indicator equal to one if firm i has above-median revenue in country j . Rating_k denotes the MSCI numerical quality rating of project k . Columns (1)–(2) restrict the sample to project countries above the year-specific median of the Climate Change Performance Index (CCPI), while columns (3)–(4) restrict the sample to project countries below the year-specific median. Columns (5)–(6) use the pooled sample and interact operational presence and project quality with the standardized CCPI score of the project country. Control variables include log emissions, log revenue, and an indicator for whether the firm has a net-zero target. All specifications include country-of-retirement-by-year fixed effects. Columns alternate between firm fixed effects and firm-by-year fixed effects, as reported. Standard errors are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	High CCPI		Low CCPI		CCPI Interaction	
	(1) Firm	(2) Firm × Year	(3) Firm	(4) Firm × Year	(5) Firm	(6) Firm × Year
Operation	0.327*** (0.124)	0.208 (0.139)	0.117 (0.113)	−0.062 (0.097)	0.233*** (0.082)	0.167* (0.087)
Rating	0.056* (0.031)	0.080** (0.036)	0.028 (0.037)	−0.047 (0.031)	0.050** (0.023)	0.037* (0.022)
Operation × Rating	−0.097** (0.041)	−0.081* (0.048)	−0.037 (0.039)	0.022 (0.037)	−0.078*** (0.028)	−0.058* (0.032)
Operation × CCPI					0.190** (0.093)	0.090 (0.102)
Rating × CCPI					0.019 (0.023)	0.011 (0.021)
Operation × Rating × CCPI					−0.051 (0.034)	−0.029 (0.036)
Observations	2,253	1,691	1,625	1,222	4,066	3,215
R-squared	0.562	0.460	0.632	0.538	0.541	0.439
Firm FE	Yes	No	Yes	No	Yes	No
Firm × Year FE	No	Yes	No	Yes	No	Yes

Table 10: Subsidiary Entry and Offset Quality

This table examines the quality of offsets retired in countries where firms establish new subsidiaries. Panel A reports estimates from the following specification:

$$Q_{ict} = \delta_1 \mathbf{1}\{\text{Entry}_{ic}, 0\text{--}2 \text{ years}\} + \delta_2 \mathbf{1}\{\text{Entry}_{ic}, 3\text{--}5 \text{ years}\} + \delta_3 \mathbf{1}\{\text{Entry}_{ic}, 6+ \text{ years}\} + \mu_{it} + \varepsilon_{ict}.$$

Panels B and C report estimates from the following specification:

$$Q_{ict} = \beta \text{NewEntry}_{ic} + \mu_{it} + \varepsilon_{ict}.$$

The dependent variable Q_{ict} denotes the quality of offsets retired by firm i in country c in year t , measured either as an indicator for above-median MSCI quality, *High Quality*, or as the continuous MSCI rating, *Avg MSCI*. NewEntry_{ic} is an indicator equal to one if country c is a subsidiary-entry country for firm i . The omitted category is countries in which firm i retires offsets but has no operations; countries with established operations are excluded from the comparison. μ_{it} denotes firm-by-year fixed effects. Panels B and C estimate the specification separately for above- and below-median-quality countries, where country quality is measured using the country-level average MSCI rating and the sample median. Standard errors are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	High Quality (1)	Avg MSCI (2)
<i>Panel A: By years since entry</i>		
Entry, 0–2 years	–0.340** (0.135)	–0.474** (0.217)
Entry, 3–5 years	–0.239** (0.104)	–0.187 (0.127)
Entry, 6+ years	–0.069 (0.103)	–0.092 (0.129)
<i>N</i>	2,238	2,238
<i>Panel B: Above-median quality countries</i>		
New entry	–0.283* (0.147)	–0.387** (0.175)
<i>N</i>	784	784
<i>Panel C: Below-median quality countries</i>		
New entry	–0.134 (0.093)	–0.091 (0.186)
<i>N</i>	1,030	1,030

A Appendix Tables

Table A1: Operational Presence, Revenue Footprint, and Offset Quality (Sylvera and BeZero ratings)

This table reports estimates from the following specification:

$$S_{kiy} = \delta_1 \text{Operation}_{ij} + \delta_2 \text{Rating}_k + \delta_3 (\text{Operation}_{ij} \times \text{Rating}_k) + \mathbf{X}'_i \boldsymbol{\gamma} + \text{Fixed Effects} + \varepsilon_{kiy}.$$

The dependent variable S_{kiy} is the share of firm i 's annual carbon offset retirements allocated to project k in year y . Operation_{ij} is a dummy equal to one if firm i has operations in the country j where project k is located. In columns (1)-(2), Operation_{ij} equals one for firms with operations in the project country, including both foreign firms with positive revenues and domestic firms headquartered in that country. In columns (3)-(4), Operation_{ij} equals one for foreign firms with a revenue footprint above the sample median in the project country. In columns (5)-(6), Operation_{ij} equals one for foreign firms with a below-median revenue footprint, and the sample excludes foreign firms-country pairs with above-median revenue footprint. Rating_k is a binary indicator equal to one if the project is rated BB or higher based on the combined letter ratings from BeZero and Sylvera, and zero otherwise. The sample is restricted to projects rated at least one of these agencies. All specifications include project classification, vintage, and country-of-retirement-by-year fixed effects. Standard errors are clustered at the firm level and reported in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	Operational Presence Domestic + Foreign		Foreign Above median revenue		Foreign Below median revenue	
	(1)	(2)	(3)	(4)	(5)	(6)
OPERATION	0.005 (0.023)	-0.040* (0.023)	0.052** (0.026)	0.035 (0.028)	-0.017 (0.025)	-0.046* (0.027)
RATING	0.059* (0.032)	0.052 (0.037)	0.058** (0.025)	0.074*** (0.026)	0.044 (0.036)	0.085** (0.042)
OPERATION \times RATING	-0.036 (0.036)	0.020 (0.040)	-0.112** (0.044)	-0.040 (0.056)	-0.002 (0.037)	-0.030 (0.042)
Log Emissions	-0.013 (0.015)		-0.011 (0.015)		-0.013 (0.015)	
Log Sales	-0.033 (0.022)		-0.037* (0.022)		-0.036 (0.023)	
Target Approved	0.029 (0.041)		0.029 (0.041)		-0.021 (0.045)	
Observations	4,089	3,171	4,089	3,171	3,145	2,406
R-squared	0.540	0.456	0.541	0.456	0.575	0.486
Firm FE	Yes	No	Yes	No	Yes	No
Firm \times Year FE	No	Yes	No	Yes	No	Yes

Table A2: Operational Presence and Offset Quality (fixed baseline)

This table reports estimates from the following specification:

$$S_{kiy} = \delta_1 \text{Operation}_{ij} + \delta_2 \text{Rating}_k + \delta_3 (\text{Operation}_{ij} \times \text{Rating}_k) + \mathbf{X}'_i \boldsymbol{\gamma} + \text{Fixed Effects} + \varepsilon_{kiy}.$$

The dependent variable S_{kiy} is the share of firm i 's annual carbon offset retirements allocated to project k in year y . Operation_{ij} is a dummy equal to one if firm i has operations in the country j where project k is located. Columns (1)–(2) compare domestic firms against no-ops (sample drops foreign-with-ops). Columns (3)–(4) compare foreign above-median revenue firms against no-ops only (sample drops domestic and below-median). Columns (5)–(6) compare foreign below-median firms against no-ops only (sample drops domestic and above-median). Rating_k is a binary indicator equal to one if the project is rated BB or higher based on the combined letter ratings from BeZero and Sylvera, and zero otherwise. The sample is restricted to projects rated at least one of these agencies. All specifications include classification, vintage, and country-of-retirement-by-year fixed effects, as well as the firm-level controls $\text{Log}(\text{Emissions})$, $\text{Log}(\text{Revenue})$, and Target Approved . Standard errors are clustered at the firm level and reported in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	Domestic vs no-ops		Foreign Above median vs no-ops		Foreign Below median vs no-ops	
	(1)	(2)	(3)	(4)	(5)	(6)
OPERATION	-0.153 (0.174)	-0.121 (0.250)	0.272*** (0.093)	0.130 (0.095)	0.178** (0.086)	-0.023 (0.085)
RATING	0.075*** (0.025)	0.044* (0.025)	0.081*** (0.023)	0.053** (0.022)	0.064*** (0.021)	0.041* (0.022)
OPERATION \times RATING	0.075 (0.056)	0.063 (0.081)	-0.105*** (0.030)	-0.063* (0.033)	-0.065** (0.028)	-0.007 (0.028)
Observations	3,268	2,598	4,153	3,377	4,441	3,716
R-squared	0.620	0.544	0.586	0.545	0.573	0.526
Firm-level controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	No	Yes	No	Yes	No
Firm \times Year FE	No	Yes	No	Yes	No	Yes

Table A3: Operational Presence, Revenue Footprint, and Ex-Ante Quality

This table reports estimates from the following specification:

$$S_{kiy} = \delta_1 \text{Operation}_{ij} + \delta_2 Q_{pdr} + \delta_3 (\text{Operation}_{ij} \times Q_{pdr}) + \mathbf{X}'_i \gamma + \text{Fixed Effects} + \varepsilon_{kiy}.$$

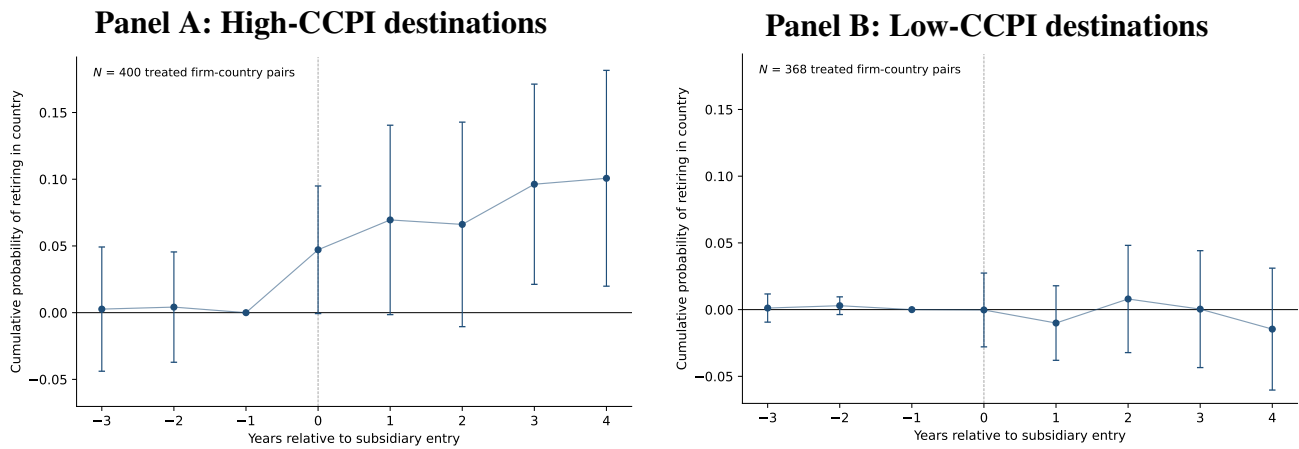
The dependent variable S_{kiy} is the share of firm i 's annual carbon offset retirements allocated to project k in year y . Operation_{ij} is a dummy equal to one if firm i has operations in the country j where project k is located. In columns (1)–(2), Operation_{ij} equals one for firms with operations in the project country, including both foreign firms with positive revenues and domestic firms headquartered in that country. In columns (3)–(4), Operation_{ij} equals one for foreign firms with a revenue footprint above the sample median in the project country. In columns (5)–(6), Operation_{ij} equals one for foreign firms with a below-median revenue footprint, and the sample excludes foreign firm-country pairs with above-median revenue footprint. Q_{pdr} is an ex-ante observable quality proxy constructed using observable characteristics at the time of retirement. For each bucket (p, d, r) ,

$$Q_{pdr} = \frac{\sum_{k \in (p,d,r)} \text{Volume}_k \cdot \mathbf{1}[\text{HighQuality}_k]}{\sum_{k \in (p,d,r)} \text{Volume}_k},$$

where p denotes harmonized project type, d developer status (local-only versus operating in multiple countries), and r registry (VCS, Gold Standard, ACR, or CAR). HighQuality_k equals one if project k 's MSCI rating is above the sample median, and zero otherwise. Weights are total retired volume. The classification variable p contains 9 categories, yielding 72 (p, d, r) buckets in total. All specifications include project classification, vintage, and country-of-retirement-by-year fixed effects, as well as the firm-level controls $\text{Log}(\text{Emissions})$, $\text{Log}(\text{Revenue})$, and Target Approved . Standard errors are clustered at the firm level and reported in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	Operational Presence		Foreign		Foreign	
	Domestic + Foreign		Above median revenue		Below median revenue	
	(1)	(2)	(3)	(4)	(5)	(6)
OPERATION	0.029 (0.021)	-0.014 (0.020)	0.034 (0.022)	0.036 (0.022)	0.023 (0.023)	-0.022 (0.023)
QUALITY	0.042 (0.029)	-0.004 (0.026)	0.010 (0.025)	0.001 (0.022)	0.025 (0.030)	-0.026 (0.028)
OPERATION \times QUALITY	-0.064** (0.030)	-0.026 (0.030)	-0.065* (0.035)	-0.071* (0.038)	-0.070** (0.035)	-0.052 (0.037)
Observations	7,808	6,653	7,808	6,653	6,069	5,088
R-squared	0.543	0.471	0.545	0.489	0.568	0.492
Firm-level controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	No	Yes	No	Yes	No
Firm \times Year FE	No	Yes	No	Yes	No	Yes

Figure A1: Subsidiary entry and offset retirement in entry country, by country-level climate ambition (CCPI)



Notes: Event-study estimates from the firm-country-year regression $y_{ict} = \sum_{k \neq -1} \delta_k \mathbf{1}\{t - g_{ic} = k\} + \alpha_{ic} + \gamma_t + \varepsilon_{ict}$, where $y_{ict} = \mathbf{1}\{\exists s \leq t : q_{ics} > 0\}$ is the absorbing-state indicator that firm i has retired offsets in country c by year t . Treated units are firm-country pairs (i, c) with an SEC Exhibit 21 subsidiary entry in c between 2010 and 2022, with g_{ic} the entry year; control units are firm-country pairs in the same set of entry countries where the firm has no operations and never enters. Panels A and B split the country set by Climate Change Performance Index: a country is classified as high-CCPI (low-CCPI) if its sample-mean overall CCPI score over 2009–2024 is above (below) the median across panel countries with CCPI coverage; the split bins the 25 covered countries into 13 high-CCPI and 12 low-CCPI. Coefficients are estimated jointly via a fully-interacted specification, with implied low-CCPI coefficients computed as the sum of the main effect and the interaction. α_{ic} are firm-country fixed effects; γ_t are year fixed effects. The omitted period is $k = -1$. Standard errors are clustered at the firm level. Error bars show 95% confidence intervals.

B Model Appendix

Proposition 1 (proof). We want to determine the conditions under which the share of offsets purchased locally is larger than that from the foreign country. This amounts to verify

$$\frac{q_{H,l}^* + q_{L,l}^*}{q_{H,l}^* + q_{L,l}^* + q_{H,f}^* + q_{L,f}^*} > \frac{1}{2}, \quad (18)$$

with equilibrium average quantities $\{q_{H,j}^*, q_{L,j}^*\}_{j=l,f}$ from (10). This is equivalent to verifying that

$$q_{H,l}^* + q_{L,l}^* > q_{H,f}^* + q_{L,f}^*. \quad (19)$$

Substituting the explicit expressions from the first order conditions, we get

$$q_{H,j}^* + q_{L,j}^* = \frac{1}{\phi + \theta^2 \gamma \rho_j (1 - \rho_j)} \left(\rho_j B_{H,j}(H) + (1 - \rho_j) B_{H,j}(L) + \rho_j B_{L,j}(L) + (1 - \rho_j) B_{L,j}(H) \right) \quad (20)$$

Because the firm's posterior expectation of quality depends exclusively on the realized signal rather than the unobservable true project state, the valuation of any project is identical conditional on receiving the same signal. Combined with symmetric production costs ($\phi_L = \phi_H = \phi$), we have that $B_{L,j}(L) = B_{H,j}(L) = B_{L,j}$ and $B_{H,j}(H) = B_{L,j}(H) = B_{H,j}$, thus the expression in (19) simplifies to

$$\left(B_{H,l} + B_{L,l} \right) > \left(\frac{\gamma \rho_l (1 - \rho_l) \theta^2 + \phi}{\gamma \rho_f (1 - \rho_f) \theta^2 + \phi} \right) \left(B_{H,f} + B_{L,f} \right) \quad (21)$$

where substituting the benefits terms and rearranging, we get

$$\left([(\beta)^\sigma + (\rho_l \theta)^\sigma]^{\frac{1}{\sigma}} + [(\beta)^\sigma + ((1 - \rho_l) \theta)^\sigma]^{\frac{1}{\sigma}} \right) > \left(\frac{\gamma \rho_l (1 - \rho_l) \theta^2 + \phi}{\gamma \rho_f (1 - \rho_f) \theta^2 + \phi} \right) \theta. \quad (22)$$

When $\sigma = 1$, condition (22) can be rearranged to obtain a lower bound on the local benefit β :

$$\beta > \kappa(\rho_f, \rho_l) \theta, \quad (23)$$

where

$$\kappa(\rho_f, \rho_l) = \frac{1}{2} \left(\frac{\gamma \rho_l (1 - \rho_l) \theta^2 + \phi}{\gamma \rho_f (1 - \rho_f) \theta^2 + \phi} - 1 \right).$$

If $\rho_l > \rho_f$, then $\kappa(\rho_f, \rho_l) < 0$, so the lower bound is slack: even $\beta = 0$ is sufficient to generate a local bias. Conversely, if $\rho_f > \rho_l$, then $\kappa(\rho_f, \rho_l) > 0$, and local benefits must be sufficiently large to overturn the information disadvantage.

When $\sigma < 1$, we use the property of the CES aggregator that for $0 < \sigma < 1$ and $a, b \geq 0$,

$$(a^\sigma + b^\sigma)^{1/\sigma} \geq a + b.$$

This implies that a sufficient condition for (19) to hold is again

$$\beta > \kappa(\rho_l, \rho_f) \theta. \quad (24)$$

Therefore, when $\sigma < 1$, the same threshold as in the linear case $\sigma = 1$ is sufficient. Moreover, since for $\sigma < 1$ the CES aggregator is strictly convex, the true lower bound on β is weakly smaller than in the linear benchmark.

Proposition 2 (proof). We seek conditions under which the firm holds a larger share of low-quality offsets locally, that is

$$\frac{q_{H,l}^*}{q_{H,l}^* + q_{L,l}^*} < \frac{q_{H,f}^*}{q_{H,f}^* + q_{L,f}^*}. \quad (25)$$

with equilibrium average quantities $\{q_{H,j}^*, q_{L,j}^*\}_{j=l,f}$ from (10). From the previous argument, the local high-quality share, $\frac{q_{H,l}^*}{q_{H,l}^* + q_{L,l}^*}$, can be written as

$$f(\beta, \sigma) = \frac{\rho_l A(\beta, \sigma) + (1 - \rho_l) B(\beta, \sigma)}{A(\beta, \sigma) + B(\beta, \sigma)}, \quad (26)$$

where

$$A(\beta, \sigma) = [\beta^\sigma + (\rho_l \theta)^\sigma]^{1/\sigma}, \quad B(\beta, \sigma) = [\beta^\sigma + ((1 - \rho_l) \theta)^\sigma]^{1/\sigma}. \quad (27)$$

The foreign high-quality share simplifies to

$$f_f = \rho_f^2 + (1 - \rho_f)^2. \quad (28)$$

As $\beta \rightarrow \infty$, both $A(\beta, \sigma)$ and $B(\beta, \sigma)$ become asymptotically linear in β , implying

$$\lim_{\beta \rightarrow \infty} f(\beta, \sigma) = \frac{1}{2},$$

independently of ρ_l . Since $f_f > \frac{1}{2}$ whenever $\rho_f > \frac{1}{2}$, the inequality is eventually satisfied for sufficiently large β . The problem therefore reduces to characterizing the threshold level of β above which $f(\beta, \sigma) < f_f$.

Consider first the benchmark case $\sigma = 1$. The CES aggregator becomes linear and the inequality reduces to

$$\frac{\beta/\theta + \rho_l^2 + (1 - \rho_l)^2}{2\beta/\theta + 1} < \rho_f^2 + (1 - \rho_f)^2. \quad (29)$$

Let $s(\rho) = \rho^2 + (1 - \rho)^2 > \frac{1}{2}$. Rearranging yields

$$\beta > v(\rho_l, \rho_f) \theta, \quad (30)$$

where

$$v(\rho_l, \rho_f) = \frac{s(\rho_l) - s(\rho_f)}{2s(\rho_f) - 1}, \quad (31)$$

and $s(\rho_j) = \rho_j^2 + (1 - \rho_j)^2 > \frac{1}{2}$. If $\rho_l < \rho_f$, then $s(\rho_l) < s(\rho_f)$ and the lower bound is negative, implying that even $\beta = 0$ generates a negative quality gradient. If instead $\rho_l > \rho_f$, then $v(\rho_l, \rho_f) > 0$ and local goodwill must be sufficiently large to overturn the informational advantage.

For the general case $\sigma < 1$, the function $f(\beta, \sigma)$ is continuous and strictly decreasing in β . Moreover,

$$f(0, \sigma) = \rho_l^2 + (1 - \rho_l)^2, \quad \lim_{\beta \rightarrow \infty} f(\beta, \sigma) = \frac{1}{2}.$$

Hence, whenever $\rho_l > \rho_f$, there exists a unique cutoff $\underline{\beta}(\sigma)$ such that $f(\beta, \sigma) < f_f$ for all $\beta > \underline{\beta}(\sigma)$. Furthermore, since the CES aggregator is convex for $\sigma < 1$, one has

$$A(\beta, \sigma) \geq \beta + \rho_l \theta, \quad B(\beta, \sigma) \geq \beta + (1 - \rho_l) \theta, \quad (32)$$

which implies

$$f(\beta, \sigma) \leq f(\beta, 1), \quad \forall \beta > 0. \quad (33)$$

Therefore, the linear benchmark in (30) provides a sufficient condition for the inequality to hold, namely

$$\beta > v(\rho_l, \rho_f) \theta \implies f(\beta, \sigma) < f_f. \quad (34)$$

Proposition 3 (proof). We show that the presence of the local benefit β reduces price sensitivity to quality. Formally, we want to establish

$$\frac{\partial p_{k,l}^*}{\partial \theta} < \frac{\partial p_{k,f}^*}{\partial \theta}, \quad k \in \{H, L\}. \quad (35)$$

with $p_{k,j}^*$ average equilibrium prices from (9). Fix for simplicity $k = H$. Using the equilibrium price formula, local and foreign prices can be written as

$$p_l^* = \left(\rho_l A(\beta, \sigma) + (1 - \rho_l) B(\beta, \sigma) \right) \frac{\phi}{\gamma \rho_l (1 - \rho_l) \theta^2 + \phi}, \quad (36)$$

where

$$A(\beta, \sigma) = [\beta^\sigma + (\rho_l \theta)^\sigma]^{1/\sigma}, \quad B(\beta, \sigma) = [\beta^\sigma + ((1 - \rho_l) \theta)^\sigma]^{1/\sigma},$$

while foreign prices simplify to

$$p_f^* = \theta (\rho_f^2 + (1 - \rho_f)^2) \frac{\phi}{\gamma \rho_f (1 - \rho_f) \theta^2 + \phi}. \quad (37)$$

Consider first the benchmark case $\sigma = 1$. The CES aggregator becomes linear, so that

$$\rho_l A + (1 - \rho_l) B = \beta + \theta (\rho_l^2 + (1 - \rho_l)^2).$$

Hence,

$$p_l^* = \left(\beta + \theta \mu(\rho_l) \right) \frac{\phi}{\gamma \rho_l (1 - \rho_l) \theta^2 + \phi}, \quad \mu(\rho) = \rho^2 + (1 - \rho)^2. \quad (38)$$

The marginal derivative of the local price with respect to θ thus becomes

$$\frac{\partial p_l^*}{\partial \theta} = \frac{\partial}{\partial \theta} \left[\frac{\phi \theta \mu(\rho_l)}{\gamma \rho_l (1 - \rho_l) \theta^2 + \phi} \right] - 2\beta \phi \left[\frac{\gamma \rho_l (1 - \rho_l) \theta}{(\gamma \rho_l (1 - \rho_l) \theta^2 + \phi)^2} \right] \quad (39)$$

whereas for the foreign price,

$$\frac{\partial p_f^*}{\partial \theta} = \frac{\partial}{\partial \theta} \left[\frac{\phi \theta \mu(\rho_f)}{\gamma \rho_f (1 - \rho_f) \theta^2 + \phi} \right] \quad (40)$$

Note that the term in first brackets in (39) is strictly increasing in ρ , whereas the second term is decreasing in β . Thus if $\rho_f > \rho_l$, the inequality in (35) holds necessarily for any β . However, for any given $\beta > 0$, local price sensitivity to quality may dampened even if $\rho_l > \rho_f$. Specifically, solving the derivatives and rearranging, one has that the inequality in (35) holds if and only if

$$\beta > \frac{1}{2a_l \theta} \left[\mu(\rho_l)(\phi - a_l \theta^2) - \mu(\rho_f)(\phi - a_f \theta^2) \frac{(a_l \theta^2 + \phi)^2}{(a_f \theta^2 + \phi)^2} \right]. \quad (41)$$

with $a_j = \gamma \rho_j (1 - \rho_j)$. When $\sigma < 1$, local benefits partially substitute quality thereby dampening price sensitivity further. As before, we establish a sufficient lower bound on β such that the inequality in (35) holds even when $\rho_l > \rho_f$. Equilibrium prices can be written in the form

$$p_j^*(\theta) = M_j(\theta, \beta) D_j(\theta), \quad (42)$$

$$D_j(\theta) = \frac{\phi}{\gamma \rho_j (1 - \rho_j) \theta^2 + \phi}, \quad (43)$$

$$M_l(\theta, \beta) = \rho_l A(\beta, \sigma) + (1 - \rho_l) B(\beta, \sigma), \quad (44)$$

$$M_f(\theta) = \theta \mu(\rho_f), \quad \mu(\rho) = \rho^2 + (1 - \rho)^2, \quad (45)$$

with

$$A(\beta, \sigma) = [\beta^\sigma + (\rho_l \theta)^\sigma]^{1/\sigma}, \quad B(\beta, \sigma) = [\beta^\sigma + ((1 - \rho_l) \theta)^\sigma]^{1/\sigma}. \quad (46)$$

The derivative of the local price with respect to quality is

$$\frac{\partial p_l^*}{\partial \theta} = \frac{\partial M_l}{\partial \theta} D_l + M_l \frac{\partial D_l}{\partial \theta}. \quad (47)$$

Note that $\frac{\partial D_l}{\partial \theta} < 0$, so the second term in (47) is negative, whereas the first one is positive since $\frac{\partial M_l}{\partial \theta} > 0$. Moreover, the marginal effect of M_l scales at rate $\beta^{1-\sigma}$, while the contribution of D_l scales at rate β . Hence, for β sufficiently large, the second (negative) term dominates and the derivative becomes negative. It follows that, for large enough β , one has necessarily

$$\frac{\partial p_l^*}{\partial \theta} < \frac{\partial p_f^*}{\partial \theta},$$

since $\frac{\partial p_f^*}{\partial \theta} = c_f > 0$. We can derive a closed-form upper bound beyond which this inequality holds. Expanding (47), we obtain

$$\frac{\partial p_l^*}{\partial \theta} = [\beta^\sigma + (\rho_l \theta)^\sigma]^{1-\frac{1}{\sigma}} \left[\rho (\rho_l \theta)^{\sigma-1} D_l + (\beta^\sigma + (\rho_l \theta)^\sigma) D_l' \right]. \quad (48)$$

Since $D'_l < 0$, a sufficient condition for $\frac{\partial p_l^*}{\partial \theta} < 0$ is that the term in brackets in (48) be negative. Rearranging yields

$$\beta > \left[\theta^{\sigma-1} \frac{D_l}{|D'_l|} - \theta^\sigma \right]^{1/\sigma} = \rho_l \left[\frac{1}{2} \theta^{\sigma-2} \left(\frac{\phi}{\gamma \rho_l (1 - \rho_l)} - \theta^2 \right) \right]^{1/\sigma}. \quad (49)$$

Notably, this threshold is strictly increasing in the accuracy of the local signal, ρ_l . As local information improves, expected quality becomes less noisy and demand becomes more sensitive to project fundamentals. If, instead of zero, we compare with the constant $c_f = \frac{\partial p_f^*}{\partial \theta} > 0$, which does not depend on β and is increasing in ρ_f , the same reasoning implies that the threshold remains increasing in ρ_l and decreasing in ρ_f .