Integrity and Unraveling in Carbon Offset Markets

Eduardo Azevedo[1] Susanna Berkouwer[1,4] Prakash Mishra[2]

Anna Russo[3] Arthur van Benthem[1,4] *

Preliminary and incomplete – please do not cite without permission.

Abstract

Carbon offset markets aim to lower abatement costs through trade, but when adverse selection drives a wedge between credits and true carbon reductions, enforcing unbiased trade—counting credits only for expected reductions—may unravel the market. We investigate this tension in tropical forest offsets. Credits exceed reductions by an order of magnitude. A natural experiment provides direct evidence of unraveling: when a standardized contract pooled heterogeneous credits at a single price, prices fell, adverse selection worsened, and trade collapsed. Yet buyers value nominal credits beyond their underlying carbon reductions. An equilibrium model of land use, crediting, and demand shows that this demand acts as an implicit subsidy, sustaining trade that adverse selection would otherwise unravel, but with small, vastly overstated abatement. Under unbiased trade, the market unravels even at high carbon prices, despite socially valuable abatement opportunities. Preserving both trade and environmental integrity would require a profound change in market design, substantial subsidies, and major improvements in baseline measurement. Without this, gains from offset trade may not be realizable.

*We thank Yuyang Wang and Xiaomeng Wu for their excellent research assistance. We are grateful to Nick Ryan, [...], and seminar and conference participants at Texas A&M, Arizona State, University of Arizona, Harvard Salata Institute, Coase, UCSB, Dartmouth, PERC, Emory, Yale School of the Environment, and the 2026 Yale Climate, Energy, and Growth Conference for insightful comments and suggestions. Van Benthem thanks Penn Global, the Kleinman Center for Energy Policy, Wharton AI & Analytics Initiative, Wharton Global Initiatives, Jacobs Levy Equity Management Center for Quantitative Financial Research, and the Wharton Dean's Research Fund for generous support.

[1] University of Pennsylvania. [2] Caltech. [3] Harvard University. [4] National Bureau of Economic Research.

1 Introduction

Carbon offset markets aim to direct emissions reductions toward low-cost opportunities in sectors and geographies not covered by carbon pricing, such as land use and forestry. By allowing suppliers to voluntarily sell CO_2 reductions or removals to higher-cost emitters, offset markets can realize these opportunities through mutually beneficial trade without imposing a globally binding emissions cap. To date, direct carbon pricing covers only 29% of global emissions (World Bank, 2026), motivating substantial policy and market interest in offset mechanisms, which operate in a voluntary carbon market serving 40% of Fortune 500 companies, in the 40% of compliance markets that permit offsetting (Macintosh et al., 2025), and as longstanding and central components of international climate agreements.

Carbon offset markets are nevertheless controversial. Their central challenge is that emissions reductions are difficult to measure and participation is voluntary, creating scope for adverse selection that can drive a wedge between issued credits and true reductions (Calel et al., 2025; Chen et al., 2025). In this paper, we argue that addressing this challenge by enforcing unbiased trade — so that credits represent the emissions reductions they claim — can unravel the market, creating a tension between environmental integrity and trade. Can an opt-in market be designed to simultaneously maintain unbiased trade and realize the low-cost abatement opportunities these markets aim to capture? We develop a conceptual framework and apply it to new data on tropical-forest offsets to answer this question, assessing whether and when adverse selection leads to biased trade, inefficiently limited trade, or complete market unraveling.

We begin with a stylized framework of an offset market in which projects and buyers each choose whether to trade, with participating projects receiving credits based on estimated emissions reductions. On the supply side, imperfect measurement introduces adverse selection: projects for which credits overstate true reductions are more likely to find entry profitable (Aspelund and Russo, 2026). On the demand side, buyers value expected emissions reductions, nominal credits, or both (Kotchen and Vogt, 2026). Together, adverse selection and buyer preferences can produce two undesirable equilibria. If buyers — or a regulatory environment enforcing trading rules — condition on the expected effectiveness of credits in the market, environmental integrity is preserved, but the market may support little trade or unravel completely (Akerlof, 1970), leaving low-cost abatement opportunities untapped. If buyers — or a regulatory environment that accepts credits at face value — instead demand nominal credits regardless of the reductions they represent, trade can be sustained, but credits may be issued and projects may incur costs even when true emissions reductions are limited or

absent. Thus, the same willingness to tolerate biased trade that sustains the market can undermine its environmental integrity.

To assess the empirical relevance of this tension, we assemble new data on both sides of the tropical-forest offset market. On the supply side, we compile geographic boundaries and annual credit issuance for 450 tropical-forest preservation and restoration projects from registration documents—over an order of magnitude more than prior influential studies (West et al., 2020; Guizar-Coutiño et al., 2022; West et al., 2023). We link these projects to satellite-derived panels of tropical forest cover, which allow us to estimate true emissions reductions, and to measures of agricultural productivity and transport costs, which capture projects' opportunity costs. On the demand side, we obtain new transaction records covering prices and quantities for project-specific over-the-counter trades and the universe of trades and deliveries for a standardized, pooled exchange-traded offset contract. Together, these data allow us to document three empirical facts consistent with our framework.

First, issued credits substantially exceed estimated emissions reductions among projects that enter the market. We estimate changes in forest cover using a matched event-study design that compares project areas with geographically and economically similar control areas, leveraging the timing of each project's first credit issuance.[1] By ten years after entry, forest cover increases by approximately 0.9 percentage points on average, with larger effects for restoration projects than for avoided-deforestation projects. Converting these changes into CO_2 equivalents using standard carbon-accounting methods, we find that cumulative issued credits exceed our preferred estimates of cumulative CO_2 reductions by approximately an order of magnitude, implying average credit effectiveness below ten percent in both sectors. Effectiveness also varies substantially across projects, and most of this variation is not explained by observable project characteristics. These findings are consistent with adverse selection into the market.

Second, a natural experiment provides direct evidence that adverse selection can unravel trade in offset markets. The introduction of a standardized, pooled exchange-traded contract gave holders of eligible credits a choice: sell bilaterally in the over-the-counter (OTC) market at project-specific prices, or deliver into the pooled contract at a single common price, with buyers unable to choose which credit they received. The pooled contract initially traded at the average OTC price of eligible credits. Lower-quality credits quickly selected into the pool: credits with lower OTC prices and from projects later shown to have low ex post effectiveness

[1] Relative to prior studies (West et al., 2020; Guizar-Coutiño et al., 2022; West et al., 2023), we use a substantially larger sample of projects, a systematic procedure for constructing matched controls, and an Empirical Bayes correction for project-level estimation error and deviations from parallel trends.

were disproportionately delivered into the pooled contract. As the composition of the pool deteriorated, its equilibrium price fell, inducing further adverse selection until trade in the pooled contract ceased at zero prices. This episode provides an unusually clean illustration of Akerlof (1970) unraveling, complementing the few existing direct (Cutler and Reber, 1998) or indirect (Hendren, 2013; Herbst and Hendren, 2024) empirical demonstrations. Relative to a control market that differs only in pooling, we observe the entire process of unraveling: composition and price deteriorate together until prices reach zero and trade stops entirely.

Third, we estimate demand under the *status quo* regulatory environment, separating buyers' valuation of true emissions reductions from their valuation of nominal credits. We exploit widely publicized information shocks about the effectiveness of a subset of projects in a hedonic triple-difference design. Before these shocks, prices were largely unrelated to estimated effectiveness. We find that buyers value a one-ton increase in expected CO_2 reductions at approximately \$3–\$6, a positive but modest amount, and place an additional value of approximately \$6 on nominal credits. This demand for nominal credits acts as an implicit subsidy, sustaining trade at low prices even among projects with low effectiveness. Beyond informing our analysis, these revealed-preference estimates point to limits of voluntary environmental action and contribute broadly to literatures on corporate social responsibility (Bénabou and Tirole, 2010; Kitzmueller and Shimshack, 2012), private provision of public goods (Bergstrom et al., 1986; Andreoni, 1988), warm glow (Andreoni, 1989, 1990), and demand for green goods (Kotchen, 2006; Kotchen and Moore, 2007).

In the second part of the paper, we build on our stylized framework to develop a simple empirical model that embeds these facts and quantifies the tension between biased trade and unraveling. Potential entrants choose whether to participate in the offset market and, conditional on participation, allocate land between agriculture and forest. Entry depends on expected credit revenue, agricultural opportunity costs, and transaction costs. An intermediary issues credits based on a noisy estimate of each project's true emissions reduction and may scale issuance upward or downward. Buyers choose among credits that differ in price, sector, and expected effectiveness. We estimate the supply side using observed land-use decisions (following Souza-Rodrigues (2019)), offset prices, credit issuance, and the distribution of event-study estimates of forest and CO_2 impacts. We estimate demand using the information shocks and empirical strategy described above.

Our parameter estimates indicate the data are best rationalized by crediting that is unbiased on average across potential projects but strongly biased among participants due to selection. In other words, credits can be issued using scientifically defensible methods and still substantially

overstate reductions among the projects for which trade is profitable. At low prices, projects enter only when they receive particularly generous credit issuance or face unusually low opportunity costs; as prices rise, this selection weakens. Demand estimates align with the hedonic estimates and further reveal high price sensitivity, consistent with most purchases occurring in the unregulated voluntary carbon market (VCM). Although trade achieves limited CO_2 reductions under the *status quo*, our estimates imply that forest cover is highly elastic to conservation payments: a perfectly and costlessly targeted payment for *net* forest change could abate approximately 50 $GtCO_2$ at less than \$30 per ton (see also Jayachandran et al., 2017; Souza-Rodrigues, 2019; Araujo et al., 2026; Assunção et al., 2026; Mishra, 2026).

We use these estimates to explore the tension between biased trade and unraveling by comparing the *status quo* with a counterfactual opt-in compliance market that enforces unbiased trade. This counterfactual mirrors active proposals to admit offsets into compliance markets with safeguards for environmental integrity (see, e.g., Pande et al. (2025), among others). In the counterfactual, credits offset compliance obligations at the regulated carbon price, and trading ratios ensure that traded credits are unbiased in expectation.

We find that this tension is empirically relevant and has direct implications for current policy. In the *status quo*, willingness to pay for nominal credits acts as an implicit subsidy, sustaining trade at low prices even though entrants are highly selected on low true reductions, and the resources devoted to ineffective projects can generate a net welfare loss. Imposing unbiased trade instead causes the market to unravel, even at carbon prices as high as \$80.[2] Because credits must substitute one-for-one for emissions in expected tons, their value is determined by the *average* effectiveness of the pool, which we estimate falls below cost at every quantity. Yet the *marginal* credit is more effective than the average, and its social value exceeds its cost. The EU, for instance, plans to meet part of its 2040 climate target with international offsets, a measure motivated by cost containment (European Commission, 2025). Our estimates imply that meeting this target through unbiased trade would require an explicit subsidy of approximately \$15 per ton of true abatement beyond the ETS price, replacing an implicit subsidy that would sustain cheap trade but deliver only a small fraction of the intended reductions.

Finally, we ask why such severe adverse selection persists despite its social costs. In the *status quo*, selection is precisely what makes offset credits cheap, and market participants benefit even as it undermines environmental integrity. Adverse selection is thus privately valuable but socially costly, leaving little incentive to improve information. Enforcing unbiased

[2]This result parallels evidence from other markets that fixing one friction can exacerbate the effects of adverse selection (Handel, 2013; Polyakova, 2016).

trade reverses these incentives. Reforming market design is therefore necessary but, given current information, not sufficient. Tropical forests indeed offer large, low-cost abatement opportunities, but opt-in trade alone cannot realize them.

2 Institutional Details

An *offset market* is any trading mechanism that features voluntary participation of suppliers (opt-in). The objective of carbon offset markets is to leverage trade to incentivize low-cost CO_2 reductions without imposing a global binding cap. Three classes of actors — suppliers, buyers, and intermediaries — make up the market, which trades *credits* denominated in tons of CO_2. Suppliers are concentrated in low- and middle-income countries and buyers in high-income countries (Figure 1). Appendix A provides additional institutional details.

Market Structure. Suppliers (*Projects*) undertake activities that reduce emissions or remove CO_2 from the atmosphere, relative to a counterfactual baseline in the absence of the project. Acceptable activities are broad. They include forest preservation and restoration (our focus, which comprises approximately 20% of the market), the expansion of renewable energy, and the adoption of clean household devices, among others (see Appendix Table A.1). Projects can occur anywhere in the world and can be conducted by non-governmental organizations, for-profit entities, and local governments or indigenous community associations. The broad scope is intended to cover low-abatement-cost activities in geographies or sectors without a carbon price.

Intermediaries (*Registries*) are non-profit entities that quantify and record CO_2 reductions to facilitate trade. Specifically, they develop rules to estimate CO_2 reduced by project activities (*methodologies*), validate and verify these activities (typically delegating audit to accredited third parties), issue credits based on estimated reductions, and maintain a ledger of credits. Registries typically charge both a fixed fee and a fee per credit. Credits are typically traded over-the-counter in bilateral transactions.

Buyers demand credits. We distinguish between two types. First, compliance buyers purchase credits to satisfy legally-binding regulatory requirements for emissions reductions in jurisdictions where carbon is priced. Approximately 40% of carbon pricing schemes allow offset credits, usually for a limited share of total required emissions reductions (4-10% of the total compliance obligation; see International Carbon Action Partnership (2023); Macintosh et al. (2025) and Appendix Figure A.2).[3] Compliance markets tie demand to regulated

[3]Examples include California's Cap-and-Trade, China's National Emissions Trading System (ETS), South

allowance prices and scheme-specific eligibility rules. Second, voluntary buyers purchase credits to support net-zero or carbon-neutrality commitments (referred to as the VCM, or Voluntary Carbon Market). Approximately 44% of Fortune Global 500 companies purchase offset credits (Climate Impact Partners, 2026). Cumulatively, 2.5 billion credited tons have been issued for voluntary buyers (Forest Trends' Ecosystem Marketplace, 2024).

Tropical Forest Conservation. We focus on the market for tropical forest conservation or restoration activities. The two dominant activities (or *sectors*) of this type are REDD+ (Reducing Emissions from Deforestation and forest Degradations) and ARR (Afforestation, Reforestation, and Revegetation). REDD+ projects are issued credits for avoided CO_2 emissions, or prevented (counterfactual) deforestation that would have otherwise occurred. ARR projects are issued credits for CO_2 removal from planting trees or restoring vegetation on otherwise bare or abandoned agricultural land. Carbon offset markets are a primary market-based mechanism for monetizing the CO_2 impacts of tropical forest conservation or restoration.

Credit Issuance and Retirement. Credits must be issued for trade to occur, but *measuring* reductions presents two challenges: (1) defining a counterfactual in terms of forest cover and (2) measuring CO_2 reductions associated with the project's changes in forest biomass relative to this counterfactual. In practice, a Registry publishes a menu of methodologies that define acceptable baselines (challenge (1)) and carbon accounting (challenge (2)). A baseline defines a counterfactual rate of forest loss (REDD+) or regeneration (ARR). The baseline is unobservable and must be modeled, e.g. via a representative reference region, historical average rates of forest loss, or ecological models. Carbon accounting—the measurement and calculation of changes in forest carbon stocks as CO_2-equivalent emissions reductions or removals—varies across regions and is subject to substantial uncertainty. Credits are issued over many years as activities are monitored and CO_2 reductions accrue.

Once issued, buyers can purchase and trade credits. When a final buyer uses the credit to support a climate claim (e.g., toward its voluntary emissions target) or surrenders the credit in a compliance market, the credit is *retired*: the registry permanently removes it from circulation by recording its retirement in the project's registry account, ensuring that it cannot be traded or claimed again. If (and only if) a carbon offset removes one ton of CO_2 from the atmosphere in expectation, it compensates one ton emitted elsewhere. If this condition is met, the offset satisfies *perfect substitutability*.

Korea's ETS, and CORSIA for international aviation; the European Union Emissions Trading System (ETS) is currently reviewing limited use of offset credits.

Offset Market Policy. Offset credit trading mechanisms have been implemented at scale across a range of market designs, including the Kyoto Protocol's Clean Development Mechanism, the Chicago Climate Exchange, numerous compliance markets, and a decentralized voluntary carbon market.[4] A common critique across implementations is an *ex post* divergence between issued credits and CO_2 reductions, as documented in West et al. (2020); Guizar-Coutiño et al. (2022); West et al. (2023). Despite this record, there exists substantial policy interest in maturing and growing this market: Article 6 of the Paris Agreement, finalized in November 2024, establishes two new mechanisms for opt-in trade (Articles 6.2 and 6.4), many net-zero plans of countries and of firms rely on offset credits, and numerous policy proposals argue for scaling trade through improved market design.[5] Many of these focus on frameworks to achieve perfect substitutability (see Appendix Table A.2). The prevailing view is that the historic challenges represent a measurement problem that better methodologies can solve.

3 Theoretical Framework

We develop a stylized model of offset credit trade to illustrate sources of inefficiency, highlight implications for market design, and guide subsequent empirical analysis.

Supply. There exists a continuum of entities, e.g., projects or jurisdictions, indexed by $j \in \mathcal{J}$, distributed according to measure μ. Each project can choose to trade a net reduction in CO_2 relative to baseline. Define Δ_j as the true net reduction in CO_2 should they choose to trade and c_j as the private cost of this reduction. Δ_j is not observable or contractible and must be estimated for trade to occur. Define *estimated* reductions as:

$$q_j = \Delta_j + \varepsilon_j. \tag{1}$$

We refer to the estimate q_j as offset *credits*. In this Section, we assume that q_j is an unbiased estimate of Δ_j in the population, so $\mathbb{E}[\varepsilon_j] = 0$.[6]

Project j chooses to trade when it is profitable to do so. Given a price p per offset credit, the

[4] Globally, the voluntary market is currently three orders of magnitude smaller than the compliance markets, but ambitious growth projections see the carbon offset market grow to \$7-35 billion by 2030 and \$45-250 billion by 2050 (https://www.msci.com/research-and-insights/blog-post/frozen-carbon-credit-market-may-thaw-as-2030-gets-closer).

[5] See, as one prominent example, U.S. Department of the Treasury et al. (2024).

[6] In Section 6, we extend this model to allow for biased estimates on average.

set of trading projects is:

$$\Theta(p) := \{j \in \mathcal{J} : p\, q_j - c_j \geq 0\}. \tag{2}$$

While $\mathbb{E}[q_j] = \mathbb{E}[\Delta_j]$, it is likely that $\mathbb{E}[q_j \,|\Theta(p)] \geq \mathbb{E}[\Delta_j \,|\, \Theta(p)]$.[7] Projects that opt in to trade will be positively selected on q_j, in which a high draw of ε_j increases offset credit revenue. They will also be negatively selected on costs c_j, which may be correlated with Δ_j.

Define the *credit supply curve* $Q(p) := \int_{\Theta(p)} q_j \, d\mu(j)$ as the quantity of credits supplied at price p, which we assume is continuous and strictly increasing, yielding inverse credit supply curve $P_S(Q)$. The quantity of *true CO_2 reductions* at any quantity of *credits* Q is $R(Q) := \int_{\Theta(P_S(Q))} \Delta_j \, d\mu(j)$. It will be useful to define the *average effectiveness* of credits in the market as $AE(Q) := R(Q)/Q$ and the *marginal effectiveness*, $ME(Q) := R'(Q)$, as the CO_2 reduction generated by a marginal expansion in credit supply. The market exhibits adverse selection when $ME(Q) > AE(Q)$.

Regulatory Environment. The regulatory environment determines offset credit demand. First, a regulator could enforce *unbiased trade*, in which demand is a function of *only* expected true reductions, $AE(Q)$. When credits are used to offset buyer emissions ton-for-ton, enforcing *unbiased trade* ensures perfect substitutability across emissions sources despite the potential for selection.[8] Alternatively, the regulatory environment could tolerate *biased trade*, in which buyers can value credits separately from the expected emissions reduction. *Biased trade* violates perfect substitutability and introduces the possibility that the existence of offset credit trade *increases* aggregate CO_2 emissions.[9] Both opt-in compliance markets and the voluntary carbon market could feature *biased* or *unbiased* trade.

Demand. Offset credits are purchased by a large number of price-taking buyers. The demand curve is determined by the regulatory environment:

$$P_D(Q) = \phi AE(Q) + \psi. \tag{3}$$

[7] This is always the case if ε_j is conditionally independent of (c_j, Δ_j) (see Appendix B).

[8] This can be enforced either by requiring that credited quantities condition on selection (rather than crediting the naive estimate q_j), or equivalently, that transacted q_j be discounted such that that buyers buy $\frac{1}{AE(Q)}$ credits to offset one ton.

[9] In a compliance offset market, in which firms face a regulated emissions cap, biased trade mechanically increases CO_2 emissions.

The first term, $\phi\, AE(Q)$, conditions on the selection equation in (2), where ϕ denotes the value of the expected reduction in CO_2 to the buyer.[10] The second term, ψ, is the buyer's value of a nominal offset credit. Unbiased trade requires $\psi = 0$. Biased trade allows $\psi > 0$.[11]

Equilibrium. An equilibrium is a price-quantity pair $(p^{eq}, Q(p^{eq}))$ such that either $p^{eq} = \phi\, AE(Q(p^{eq})) + \psi$ or, if no $p > 0$ satisfies $p \leq \phi\, AE(Q(p)) + \psi$ with $Q(p) > 0$, the market unravels and $(p^{eq}, Q^{eq}) = (0, 0)$.

Socially-Optimal vs. Equilibrium Outcomes. The social value of a marginal increase in credits is $\Phi ME(Q)$, where Φ denotes the social value of one ton of CO_2 removed from the atmosphere. Buyers value it at $P_D(Q)$. The wedge between social and private valuations is:

$$\underbrace{\Phi(ME(Q) - AE(Q))}_{\text{selection effect}} \quad - \quad \underbrace{\psi}_{\text{nominal credit demand}} \quad + \quad \underbrace{(\Phi - \phi)\, AE(Q)}_{CO_2 \text{ valuation difference}} . \tag{4}$$

There are three sources of inefficiency, which may vary across regulatory environments. The first term captures selection: unbiased offset trade implies that demand must depend on $AE(Q)$ to ensure perfect substitutability, but the social planner values the reductions in CO_2 generated by the *marginal* expansion of supply, $ME(Q)$. This source of inefficiency arises in any opt-in trading mechanism enforcing perfect substitutability. The second term captures demand for nominal credits in regulatory environments that tolerate *biased trade*. A positive ψ raises the private valuation above the expected CO_2 reductions. The third term captures any difference in the valuation of CO_2 between the offset buyer and the social planner.

Graphical Analysis. Figure 2 illustrates implications in two example (polar) cases.[12] Both sub-figures plot an (identical) inverse supply curve of credits $P_S(Q)$, which denotes the total quantity of credits Q (x-axis) supplied at each price p (y-axis). Next, both sub-figures plot the social value of the *marginally-supplied* credit, $\Phi ME(Q)$. It is socially-efficient for an entity to opt in and implement CO_2 reductions whenever the $\Phi ME(Q)$ curve lies above the inverse supply curve $P_S(Q)$ ($Q^* > 0$, $p^* > 0$).

Figure 2a plots demand under a trading mechanism that enforces *unbiased trade* $P_D(Q) = \Phi\, AE(Q)$. This demand curve is upwards-sloping because it conditions on the selection

[10] In a compliance market, the regulator controls ϕ, the allowance price, through the regulated market design. In voluntary markets, ϕ represents the value to firms of voluntary reductions in CO_2 emissions due to broader regulatory, investor, or consumer pressure.

[11] One could interpret $\psi > 0$ as warm glow, private branding/marketing benefits, political pressure to tolerate cheap credits, or ignorance of the selection equation (2).

[12] Appendix B provides additional examples.

equation (2) and therefore reflects the potential for adverse selection in the market.[13] In the example of Figure 2a, the requirement of *unbiased trade* and the existence of adverse selection together *unravel trade* (Akerlof, 1970). The demand curve ($P_D(Q)$) lies below the inverse supply curve ($P_S(Q)$) everywhere. Perfect substitutability makes it unprofitable to trade at every price, even at the planner's valuation of Φ.[14]

Figure 2b plots an alternative polar case of *biased trade* with buyers who value *only* nominal credits $P_D(Q) = \psi$. In this case, *trade occurs* at quantity Q^{eq} because the marginal entrant can be profitable by selling *all* credits q_j at ψ. The revenue from credit sales in excess of the valuation of expected CO_2 reductions acts as an implicit subsidy that supports trade, but by definition violates the substitutability of emissions that underlies the premise of trade.

In the two cases in Figure 2, *biased trade* is necessary for trade to exist at all. But these examples are only illustrative: the existence and implications of selection, the limits to trade, and the nature of demand for true vs. nominal CO_2 reductions are all empirical questions.

4 Data

We assemble a dataset combining information on credited quantities q_j, data that facilitates estimation of CO_2 reductions Δ_j, and transaction prices p. Appendix C provides details.

Project Information from Registry Ledgers. Registries record detailed information about all credited offset projects. We scrape and digitize information from these registry ledgers. We collect each project's unique identifier, sector (e.g., REDD+ or ARR), entry year, and credits issued q_j (by year). Our main analysis focuses on tropical forest conservation and reforestation, where we are able to obtain geo-spatial boundaries of each project from maps in the recorded project documentation. Our primary analysis sample includes 450 projects with usable forest-cover outcomes that comprise approximately two-thirds of all credits issued from tropical forest projects. Panel A of Table 1 presents summary statistics of project information.

Forest Cover and Other Geo-Spatial Information. We measure tropical forest cover with the European Commission Joint Research Centre Tropical Moist Forest (TMF) product (Vancutsem et al., 2021), derived from 30-meter resolution Landsat imagery. This dataset classifies the tropical band into undisturbed, degraded, regrowing, and deforested forest

[13] Appendix B provides details on when this curve would be upwards-sloping.

[14] This is a polar case: adverse selection could exist while still only modestly limiting trade. See Appendix Figure A.3.

annually from 1990 to 2023. We also compile several other datasets describing the biome, market access (transport costs), and agricultural profitability (yields, prices) of each location (see Appendix C for more details). Finally, we measure geographically varying forest carbon with IPCC measures of carbon stock in standing forest (Intergovernmental Panel on Climate Change, 2003) and carbon sequestration in secondary forest regrowth (Robinson et al., 2025).

Panel B of Table 1 compares geo-spatial characteristics of offset projects (column (1)) to other tropical land (column (2)). Projects that select into trade have higher forest cover, lower deforestation, lower market access, and higher carbon stocks (existing and regrowth) than the average tropical forest location. This comparison suggests that projects select into trade based on low opportunity costs (low payoffs outside the offset market) and high relative carbon stocks (high payoffs inside the offset market).

Credit Transaction Prices. We collect information about credit transactions — including prices — from Xpansiv/CBL, a centralized digital marketplace for trading offset credits and "the world's largest and most active spot exchange for environmental commodities." The digital marketplace provides infrastructure to settle bilateral, project-specific over-the-counter (OTC) trades and records the unique project identifier, the transacted quantity of credits, the per-credit transaction price, and the transaction date. During our sample period (2019-2025), Xpansiv/CBL also introduced a standardized, exchange-traded contract (the Standard Instruments Program, or SIP) that represents a bundle of credits that meet eligibility standards trading at a single price (without project identifiers).

Panel C of Table 1 summarizes the transaction data (column (3)). The average credit price is \$7.02, and the mean transaction contains 20,957 credits.

5 Evidence

In this section, we document three empirical facts consistent with the model in Section 3. First, we estimate a stark divergence between issued credits (q) and associated CO_2 reductions (Δ) among entrants. Second, we present direct evidence of market unraveling using a natural experiment exploiting variation across offset contracts. Third, we estimate demand for credits that rationalizes *status quo* trade, separating willingness-to-pay for CO_2 reductions (Δ) and nominal credits (q).

5.1 Supply-Side Evidence: Credits vs. CO_2

We compare issued credits to estimated CO_2 reductions in three steps. First, we estimate changes in tropical forest cover in project areas relative to a counterfactual baseline in a matched difference-in-differences design. Second, we convert changes in forest cover to reductions in CO_2 (Δ) using best practices in carbon accounting. Third, we compare these estimates, both market-wide and project-by-project, to issued credits (q).

We divide the tropical forest region into a square kilometer grid. For each project $j \in \mathcal{J}$ with entry year E_j, let $\mathcal{T}_j$ and $\mathcal{C}_j$ denote the sets of treated and control grid cells. We match each treated cell to three never-treated control cells in the same country-biome, without replacement, on the straight-line distance to the nearest city, port, main road, and local road, and on levels and changes in forest cover in the five years before E_j.[15] We do not match on pre-period coefficients prior to the five years before treatment in order to assess the validity of the design. The entry year E_j is defined as the calendar year in which the project begins issuing credits. Identification comes from variance in the timing of project entry. There is substantial and varied delay (often exceeding five years) between the year a project begins the issuance process and the ultimate issuance year, E_j. Appendix Figure A.4 plots a histogram of entry years E_j and the issuance delay.

Pooled Effects. We estimate average treatment effects over time, pooling across projects, using a stacked event study design:

$$y_{gt} = \sum_{k \neq -1} \beta_k \, D_g \cdot \mathbf{1}[t - E_j = k] + D_g \times \kappa_j + \delta_t \times \kappa_j + \varepsilon_{gt} \tag{5}$$

where y_{gt} is the forest fraction in grid cell g in year t; $D_g \equiv \mathbf{1}[g \in \mathcal{T}_j]$ indicates treatment; κ_j is a fixed effect for project j's matched set $\mathcal{T}_j \cup \mathcal{C}_j$; and δ_t is a year fixed effect. The interaction $D_g \times \kappa_j$ absorbs level differences between treated and control cells within each project's matched set, and $\delta_t \times \kappa_j$ absorbs common time shocks within each set.[16] Because both unit and time effects are interacted with project fixed effects, each project's treated and control cells are compared only to each other, avoiding the negative weighting problem identified by Sun and Abraham (2021), following an approach similar to Cengiz et al. (2019). To account for spatial correlation, we report Conley standard errors with a uniform kernel and a 50km radius cutoff.

[15] The no-replacement restriction applies within a project-specific match; the same donor cell may be used in the matched samples of different projects.

[16] Appendix Figure A.6 plots robustness to alternative fixed effects.

Project-Specific Effects. For each project j, we estimate:

$$y_{gt} = \sum_{k \neq -1} \beta_j^k \, D_g \cdot \mathbf{1}[t - E_j = k] + D_g + \delta_t + \varepsilon_{gt} \tag{6}$$

where δ_t are year fixed effects. The project-level treatment effect is $\beta_j^{\bar{k}}$, the coefficient at the last issuing year $\bar{k}$ (or the last observed year for projects still issuing at the end of the sample). We convert each $\beta_j^{\bar{k}}$ to CO_2 equivalents using the carbon accounting methods described in detail in Appendix D. For REDD+, we use a forest stock accounting method following IPCC methodology (Intergovernmental Panel on Climate Change, 2003) that values avoided deforestation by converting standing biomass into CO_2 that would have been lost. For ARR, we use a location- and age-dependent Chapman-Richards growth curve (Robinson et al., 2025) to capture gradual sequestration in young stands.

We face two challenges in interpreting the distribution of project-specific estimates. First, each effect is estimated with error. Second, while the assumption of parallel trends may hold on average, it may not hold exactly *for every project* with a single uniform matching strategy. We address these challenges with an Empirical Bayes procedure that extends Rambachan and Roth (2023); Kwon and Roth (2024); Walters (2024) to adjust our estimates for both estimation error and project-specific (potential) deviations from parallel trends. Appendix E.2 provides details.

Taken together, our estimates build on earlier, influential studies of overcrediting among select REDD+ projects (West et al., 2020; Guizar-Coutiño et al., 2022; West et al., 2023) and expand and improve upon them in three ways. First, we present estimates at a market-wide scale (compared to 12, 40, and 27 projects, respectively).[17] Second, we use a standard systematic method to construct matched controls.[18] Finally, our Empirical Bayes approach estimates the distribution of project-specific effects while accounting for project-level estimation error and potential deviations from parallel trends.

Results. Figure 3 and Table 2 present results. Figure 3a plots the pooled event study coefficients β_k from equation (5), with estimates reported in Panel A of Table 2).[19] Held-out

[17] Our sample of 450 offset projects exceeds prior analysis by more than an order of magnitude and contains both REDD+ and ARR projects. Market-scale analysis is necessary to test for a *market-wide* divergence between Δ and q (as in Section 3) vs. divergences for an individual or a possibly selected subset of projects.

[18] Our donor pool consists of every untreated grid cell in the same country and biome with a common matching algorithm for each project; earlier work has used project-specific synthetic controls from donor pools that are restricted ex ante according to project-specific criteria or adjusted ex post only among projects where the synthetic control fails to achieve balance.

[19] Appendix Figure A.7 presents results with a uniform weight for each project.

pre-period coefficients are close to zero, supporting the validity of the design. By event year ten, forest cover increases by 0.9 percentage points. Table 2 Panel A and Appendix Figure A.5 report estimates separately by sector and document intuitive patterns of heterogeneity: by event year 10, treatment effects for ARR projects are approximately 2.2 percentage points, compared with 0.8 percentage points for REDD+ projects. The change in forest cover due to a reduction in the latent risk of deforestation (REDD+) is lower than that of projects that restore tree cover on degraded lands (ARR).

Figure 3b and Panel B of Table 2 compare aggregate estimates of CO_2 reductions (Δ) to issued credits (q). Figure 3b sums the cumulative project-specific CO_2 reductions, obtained by multiplying estimates from equation (6) by CO_2 accounting factors, and compares this sum to the cumulative issued credits. Issued credits exceed estimated reductions by approximately an order of magnitude—our preferred credit-level average effectiveness is below ten percent. Table 2 and Appendix Figure A.5 document similar results across sectors. The equally-limited effectiveness for REDD+ vs. ARR may be surprising in light of the much-higher transaction prices for tree-planting offsets, suggesting that market participants perceive them as higher quality (World Bank, 2026).

Figure 3c plots the distribution of estimated project effectiveness (Δ_j/q_j) (Panel C of Table 2 reports means). Figure 3c documents substantial heterogeneity and that most of the distribution lies far below one. Appendix Figure A.8 documents that observable project characteristics explain less than 10% of this variation.[20]

These facts together are consistent with q_j as an imperfect measure of Δ_j. The selection equation (2) implies that even unbiased credit issuance *on average* would imply the over-stated issuance among entrants (the analysis sample) observed in Figure 3 and Table 2.

Appendix Table A.6 uses detailed project documentation to investigate this result and the relevance of strategic site selection. Projects choose to enter in areas where favorable baselines can be justified by crediting rules, thereby allowing large volumes of issued credits, but where realized changes in forest (relative to matched controls) are modest. Comparisons across crediting rules illustrate: when rules use historic deforestation to justify baselines, projects have higher historic deforestation. When nearby reference regions are used, nearby deforestation is substantially higher than when it is irrelevant to crediting. Reforestation (ARR) projects are located in areas with naturally occurring regeneration.[21]

[20] Appendix Figure A.9 also documents no trends across issuance years.

[21] In addition to the clear scope for site selection, we also observe that only 27% of projects that begin the crediting process ultimately complete it and issue credits.

While these results are consistent with adverse selection, selection need not be the only explanation for the documented wedge between q and Δ. One role of the empirical model in Section 6 is to assess the role of selection versus crediting that is biased on average.

Taken together, the empirical evidence in this subsection establishes that offset markets have important ingredients of an adversely selected market: significant over-crediting and strategic selection into the market. We next examine evidence on the implications for trade.

5.2 Equilibrium Evidence: Unraveling

We exploit a natural experiment to provide direct evidence on the equilibrium consequences of adverse selection for trade. We study the introduction of a new contract that pools heterogeneous offset credits into a standardized contract. We trace the unraveling — in equilibrium prices and composition — of this pooled contract, relative to a separating comparison group. It is important to note, however, that while this natural experiment provides unusually direct evidence on the relevance of unraveling in our context, it examines one particular trading instrument. It alone therefore cannot answer the question of whether the market itself would unravel under counterfactual regulations that impose *unbiased trade*.

The Trading Experiment. Trade in offset markets typically occurs over-the-counter (OTC), in which credits from different projects and years trade as different assets, each with different observable attributes. In late 2021, a standardized, exchange-traded contract, the Standard Instrument Program (SIP), was introduced with an objective to increase standardization and market thickness while imposing minimum quality standards. A SIP contract trades at a single price, observed to the buyer. Upon purchase, credits — differentiated across projects — are delivered to the buyer among the pool of sellers of eligible credits willing to settle trades through the SIP at equilibrium prices. The buyer learns the identity of the individual project(s) only after purchase of the pooled contract. We study the Nature-based Global Emissions Offset, or N-GEO contract, which pools eligible Agriculture, Forestry, and Other Land Use (AFOLU) credits.

These two parallel trading structures create a sharp test of the classic unraveling logic of Akerlof (1970) by comparing pooled vs. individual offset contracts. Sellers and buyers *choose* whether to transact OTC (observables revealed) or through the SIP (pooled). If buyers value credits heterogeneously, sellers of less-valued credits earning low OTC prices should disproportionately select into the pooled contract. Buyers should rationally expect an adversely selected pool of credits, causing equilibrium prices to adjust, further inducing adverse selection.

Results. Figure 4 presents the prices and composition of SIP-eligible credits that trade OTC (separated) vs. through the SIP (pooled) over time.[22] Figure 4a plots the time series of prices for these two groups. When the SIP contract was first introduced, the two series tracked each other closely. But for selection, they represent the same underlying asset. However, relatively quickly, the SIP price series begins to diverge and ultimately collapse. This divergence, and ultimate unraveling, may be driven by adverse selection (Akerlof, 1970).

Figures 4b and 4c test this hypothesis by examining the composition of the SIP (pooled) trades, relative to the OTC (separated) trades, over time. In 2022, the composition of OTC and SIP credits is similar. As prices diverge, lower-quality credits increasingly select into the SIP. Figure 4b shows that credits with lower OTC prices disproportionately select into the SIP, and increasingly so as SIP prices decline. Figure 4c documents that projects that were publicly exposed as having low *ex-post* effectiveness (Δ_j/q_j) are increasingly likely to transact through the SIP rather than OTC.[23] Appendix Figure A.12 documents adverse selection into the SIP on additional dimensions and impacts on traded volumes.

Together, Figure 4 establishes that there is scope for adverse selection in the market that can limit or unravel trade. However, further interpretation is subject to two important caveats. First, we examine selection into one particular trading instrument (the SIP). Second, the SIP conceals *all* project attributes. Buyers may value underlying CO_2 reductions (ϕ), heterogeneous observable project characteristics unrelated to CO_2 reductions (ψ), or both. Estimating these parameters is essential for quantifying the scope for inefficiency in the offset credit market and evaluating alternative market designs.

5.3 Demand-Side Evidence: Willingness to Pay for Credits vs. CO_2

We separately estimate willingness to pay (WTP) for expected CO_2 reductions (ϕ) versus nominal credits (ψ) under the *status quo* regulatory environment. We leverage a market-wide information shock, which provides quasi-random variation in information about offset effectiveness.

The Information Experiment. In January 2023, a series of media outlets widely publicized academic estimates (West et al., 2020; Guizar-Coutiño et al., 2022; West et al., 2023; The Guardian, 2023) of *ex-post* evaluations of the ratio Δ_j/q_j for a subset of tropical forest projects. At approximately the same time (late 2022), third-party rating firms began publicizing estimates of Δ_j/q_j. This event combines three useful features. First, it revealed

[22] 93% of credits transacted through the SIP are also traded OTC in a given year.
[23] We explore this source of variation in more detail in Section 5.3.

project-specific estimates of the effectiveness ratio Δ_j/q_j.[24] Second, it was high-profile: information about Δ_j/q_j was readily available and sent a shockwave through the market.[25] Third, only a subset of projects received these *ex-post* evaluations, due to the time-intensive process and data requirements for conducting them.

Empirical Strategy. We use this variation in a hedonic triple-differences design to estimate the magnitude of ϕ. We compare prices across three dimensions: (i) projects with higher vs. lower published effectiveness ratios ($\tilde{\Delta}_j^{pub}$), (ii) projects with vs. without a published estimate, and (iii) before vs. after the information shock:[26]

$$p_{ljt} = \phi Post_t Pub_j \tilde{\Delta}_j^{pub} + \nu^1 Pub_j \tilde{\Delta}_j^{pub} + \nu^2 Pub_j Post_t + \nu^3 Pub_j + x_j'\gamma + \delta_t + \varepsilon_{ljt} \quad (7)$$

where p_{ljt} is the transaction price of transaction l for a credit of project j in year t, x_j are project controls, and δ_t are year fixed effects. The coefficient of interest ϕ measures the differential post-shock price response to revealed project effectiveness. Pub_j absorbs selection into which projects had published estimates; $Pub_j\tilde{\Delta}_j^{pub}$ absorbs fixed differences among assessed projects with published estimates correlated with $\tilde{\Delta}_j^{pub}$; and $Pub_j Post_t$ absorbs any aggregate shocks to projects with published estimates during and following 2023. We weight each transaction by the quantity of credits transacted.

Using equation (7) to estimate and interpret ϕ (as in Section 3) requires three further assumptions (see also Appendix E.4.2). First, we assume that buyers have no information about effectiveness prior to the information provision. This assumption is testable. Second, we assume that all buyers update their expectations about effectiveness to the published estimates. This is not directly testable, but may be justified given the nature of the shock. Finally, in the presence of buyer heterogeneity, we note that the hedonic approach in equation (7) recovers the willingness to pay of the marginal buyer.

[24] All but one source provided numeric estimates of the ratio; one rating firm published letter grades with an explicit numerical mapping to an effectiveness ratio.

[25] The coverage was sufficiently prominent to have market-wide effects. The CEO of the largest registry (Verra) resigned, and industry groups (e.g., Forest Trends' Ecosystem Marketplace (2024)) attribute a subsequent market-wide contraction and price collapse to this coverage. Firms purchasing credits were individually named in media reports.

[26] In our baseline estimates we assume the shock occurred uniformly in 2023. In Appendix Table A.8, we present results assigning a project-specific year of the shock equal to the first year in which any published estimate of $\tilde{\Delta}_j^{pub}$ was available (e.g., from a research paper). Appendix Table A.8 also presents results with additional time-varying controls. Results are similar.

Results. Table 3 reports results.[27] Column (1) first documents no correlation between price and effectiveness without the revealed public information, supporting the first assumption above. Columns (2)-(7) report estimates of ϕ from different specifications of equation (7); magnitudes can be interpreted as the change in price moving from a credit representing zero to one ton of CO_2 reductions. Columns (2) and (3) pool all estimates of $\tilde{\Delta}_j^{pub}$, separately parameterizing $\tilde{\Delta}_j^{pub}$ as a continuous variable or a re-scaled indicator (above vs. below median). Columns (4) and (5) restrict to publicized estimates from academic research (Probst et al., 2024). Columns (6) and (7) restrict to estimates that are reported directly as a numeric effectiveness ratio (dropping those reported as letter grades with associated numeric interpretations).

Across specifications, estimates of ϕ range from approximately \$3 to \$6. Although positive and typically statistically distinguishable from zero, these values are modest. These estimates provide quantitative evidence on the limits to voluntary action and reflect the fact that compliance markets accepting offsets tend to also have low carbon prices (see Appendix Figure A.2).

Given estimates of ϕ, we recover WTP for nominal credit attributes as a residual (reported in column (8)). Heterogeneous credit attributes cause ψ_j to vary across credits. We calculate ψ_j as the average value of $p_{ljt} - \hat{\phi}\tilde{\Delta}_j^{pub}$ which amounts to to approximately \$6.50. Given these magnitudes, we decompose the revealed value of an offset credit that reduces one ton of CO_2 into: (1) the value of the reduction itself ($\sim$40%) and (2) the value of the nominal credit ($\sim$60%). This highlights the relevance of both terms in equation (3) and the potential for *biased trade* to be sustained by a positive WTP for credits independent of expected CO_2 reductions.

Discussion. These results reflect the *status quo* regulatory environment, not innate preferences. Any alternative market design would affect demand and equilibrium market outcomes. Indeed, these results suggest potential scope for improvement via enforcement of *unbiased trade* and higher values of ϕ, which can occur in a compliance-market setting. These are the aims of current reforms. However, it remains ambiguous whether such reforms would improve or unravel the market.

[27] Appendix Figure A.13 plots dynamic event study coefficients.

6 Empirical Model

We develop an empirical model that extends the stylized model in Section 3 and embeds the results in Section 5 to quantitatively evaluate the tension between biased trade and unraveling.

6.1 Model

Land Use. Each potential entrant, indexed by j, consists of N_j plots of land, indexed by g, and may either enter and trade offset credits ($x_j = 1$) or not ($x_j = 0$). Potential entrants observe plot-level profitability in agriculture (π^a_{jg}) and forest (π^f_{jg,x_j}). When $x_j = 0$ (baseline), the payoff to forest is normalized to zero ($\pi^f_{jg,0} = 0$); when $x_j = 1$, the payoff to forest is inclusive of offset credit revenue. Plots are allocated to agriculture when more profitable than forest ($\pi^a_{jg} \geq \pi^f_{jg,x_j}$).

Define θ_{j0} and θ_{j1} as the fraction of j's plots used for forest at baseline ($x_j = 0$) and when trading offset credits ($x_j = 1$). The net reduction in CO_2 emissions is therefore:

$$\Delta_j = m_j \left(\theta_{j1} - \theta_{j0}\right), \tag{8}$$

where m_j converts land use change to CO_2 emissions. We assume that conditional on trade, Δ_j is unobserved and non-contractible, but that the forest share θ_{j1} is observed and contractible.

Credit Issuance. An intermediary estimates reductions and issues q_j offset credits,

$$q_j = \alpha \left(\Delta_j + \varepsilon_j\right), \tag{9}$$

where ε_j is the estimation error ($\mathbb{E}\left[\varepsilon_j\right] = 0$) and α is a shading parameter that scales estimates into offset credits (q_j) to be traded. While the stylized model (Section 3) imposed that credit issuance is unbiased on average ($\alpha = 1$), we introduce the possibility that the intermediary either shades credit awards downward (to account for selection, $\alpha < 1$) or upward (to issue excess credits, $\alpha > 1$). The intermediary earns a per-credit fee τ.

Entry. A potential entrant enters and trades ($x_j = 1$) if expected revenue exceeds costs,

$$\left(\mathbb{E}\left[p_j\right] - \tau\right) q_j - \omega_j - \chi_j \geq 0, \tag{10}$$

where $\mathbb{E}\left[p_j\right]$ is the expected price per offset credit, ω_j is the opportunity cost (agricultural profits π^a_{jg} forgone on plots switched from agriculture to forest from θ_{j0} vs. θ_{j1}), and χ_j is an

entry or transaction cost. Entry costs χ_j are real resource costs associated with preparing a project for issuance and providing necessary documentation, but τq_j is a transfer.

Demand. Offset credits are purchased by price-taking buyers, indexed by i. The payoff to buyer i from purchasing one credit from project j is:

$$u_{ij} = \phi\, AE(\mathbf{z}_j) + \psi_{ij} - p_j, \tag{11}$$

where the first term denotes the value of the expected reduction in CO_2 per credit (tons, units of Δ) valued at ϕ (dollars per ton, as in equation (3)). $\mathbf{z}_j$ denotes buyers' information about project j, and average effectiveness,

$$AE(\mathbf{z}) := \frac{\mathbb{E}\left[\Delta_k \mid x_k = 1,\ \mathbf{z}_k = \mathbf{z}\right]}{\mathbb{E}\left[q_k \mid x_k = 1,\ \mathbf{z}_k = \mathbf{z}\right]},$$

is the credit-weighted average effectiveness of traded credits that buyers cannot distinguish from j's. When $\mathbf{z}_j$ contains only an indicator or a coarse attribute (e.g., sector), $AE(\mathbf{z}_j)$ pools across entrants (as in Section 3). When $\mathbf{z}_j$ includes a project-specific estimate (as for projects with published estimates following the information shock in Section 5.3), $AE(\mathbf{z}_j)$ is that project's own effectiveness, Δ_j/q_j. The second term ψ_{ij} denotes the willingness-to-pay (in dollars per credit) for nominal credit attributes unrelated to reductions in CO_2, which may differ across credits j and buyers i. Buyers make a discrete choice among credits j and the outside option of no purchase, normalized to zero. Define $D_j(\mathbf{p})$ as the demand for credits j, given a vector of prices $\mathbf{p}$.

Equilibrium. In equilibrium, markets clear, $D_j(\mathbf{p}) = q_j\, x_j$.

6.2 Parameterization and Estimation

Estimation proceeds in three steps. First, we estimate a model of land use (following Souza-Rodrigues (2019)), which determines the empirical distribution of (Δ_j, ω_j). Second, we use the event study estimates in Section 5.1 to estimate the parameters governing credit issuance and entry. Third, we estimate *status quo* demand using the strategy in Section 5.3. Appendix F provides details about data and estimation.

Land Use. We parameterize agricultural profit on plot g for potential entrant j as

$$\pi^a_{jg} = \bar{\pi}^a_j - \kappa_j + \delta_{cb} + \zeta_j + \nu_{jg}, \tag{12}$$

where $\bar{\pi}_j^a$ is the annual average potential agricultural revenue across j's plots, calculated from GAEZ yields and FAO prices, κ_j is the cost of transporting agricultural output to the nearest city or port,[28] δ_{cb} is a country-biome shifter, ζ_j is project-specific unobservable, and ν_{jg} is a plot-level idiosyncratic shock distributed Type 1 Extreme Value (T1EV) with scale parameter σ_ν. At baseline, the share of land in forest is:

$$\theta_{j0} = \frac{1}{1 + \exp\left(\frac{1}{\sigma_v}\left(\bar{\pi}_j^a - \kappa_j + \delta_{cb} + \zeta_j\right)\right)}. \tag{13}$$

We construct potential entrants by defining a grid of the tropical forest region equal to the average size of an entrant in each sector (REDD+ or ARR). We estimate the parameter σ_ν using observations of land use in 2014 (prior to the scale up of the offset market). Regions with high unobservable agricultural returns ζ_j are more likely to have road infrastructure, which can attenuate ordinary least squares estimates of σ_ν. We instrument for transport costs with the straight line distance from each potential entrant to the nearest port or city.

Given an estimate of σ_ν, we then invert observed land use in 2023 to recover π_{jg}^a (Berry, 1994).[29] Given the estimated parameters and the offset revenue per forested plot $((\mathbb{E}[p_j] - \tau)\,\alpha\, m_j)$, we can calculate θ_{j1} and $(\omega_j,\, \Delta_j)$.

Credit Issuance and Entry. We parameterize the estimation error and entry cost as

$$\varepsilon_j \sim N(0, \sigma_\varepsilon^s) \quad \chi_j \sim N\left(\mu_\chi^s, \sigma_\chi^s\right) \tag{14}$$

which we allow to differ by sector s. We estimate the parameter vector $\left(\alpha, \sigma_\varepsilon^s, \mu_\chi^s, \sigma_\chi^s\right)$ by Simulated Method of Moments (SMM), given the entry condition in equation (10). For each sector, we match the entry rate ($\mathbb{E}[x_j]$), the event study estimates of forest cover ($\mathbb{E}[(\theta_{j1} - \theta_{j0})\, x_j]$) and CO_2 emissions ($\mathbb{E}[\Delta_j\, x_j]$), issued credits ($\mathbb{E}[q_j\, x_j]$), the correlation between issued credits and CO_2 emissions among entrants ($\rho(q_j\, x_j,\, \Delta_j\, x_j)$), and the mean CO_2 conversion factor among entrants ($\mathbb{E}[m_j\, x_j]$). The CO_2 conversion factor, m_j, is a shifter of the returns to offset market entry, which we assume is exogenous to opportunity and entry costs.

[28] We assume that crops are chosen to maximize potential agricultural annual revenue net of transport costs and calculate $\bar{\pi}_j^a - \kappa_j$ as the average of the plot-specific profit-maximizing crops within a potential entrant. Following Souza-Rodrigues (2019), transport costs are scaled by yields and are denominated in dollars per hectare per year.

[29] For non-entrants, θ_{j0} is directly observed. For entrants, our event study estimates measure the treatment effect $\theta_{j1} - \theta_{j0}$. Given observed non-forest cover θ_{j1} in 2023, we recover θ_{j0}.

In implementation, we assume a potential entrant is REDD+ if it is at least 80% forested and ARR otherwise, that price expectations ($\mathbb{E}[p_j]$) are equal to the average credit price per sector, and that the fee τ is equal to \$0.20 per credit.

Demand. We parameterize equation (11) as

$$u_{ijt} = \phi\, AE(\mathbf{z}_{jt}) + \mathbf{z}_{jt}\psi + \gamma p_{jt} + \xi_{jt} + \epsilon_{ijt}, \tag{15}$$

where p_{jt} is the offset transaction price at date t, γ governs price sensitivity, ξ_{jt} is an unobserved credit demand-shifter, and ϵ_{ijt} is an idiosyncratic shock distributed T1EV with scale parameter normalized to one. The outside option of no purchase is normalized to zero. The key parameters of interest are ϕ and γ.

We estimate equation (15) using data on credit transactions, closely following the hedonic regressions in Section 5.3. We use the saturated triple difference interactions from (7) to estimate ϕ and instrument for price using a project's distance to the nearest city and port. We define the total potential market size using industry projections (Blaufelder et al., 2021).

6.3 Estimates

Table 4 summarizes parameter estimates.

Land Use. Panel A of Table 4 reports estimates of $\frac{1}{\sigma_v}$ and the implied average deforestation elasticity with respect to price.[30] The estimated elasticities are high, between 2.6 and 4.7, consistent with other work studying land use change in the tropics (Souza-Rodrigues, 2019; Jayachandran et al., 2017; Araujo et al., 2026; Mishra, 2026). These results indicate that if subsidies could be perfectly and costlessly targeted to *net increases* in tropical forest cover, these subsidies would indeed reduce CO_2 at low cost, generating large potential gains from trade. Approximately 50 $GtCO_2$ could be abated cumulatively at a marginal cost of less than \$30 per net ton. However, net changes in forest cover Δ_j are not directly observed and must be estimated to trade. Trade may involve non-zero transaction costs χ_j. Whether these potential gains are realizable depends on these costs and information constraints.

[30] Appendix Table A.9 presents a comparison of OLS and IV estimates.

Credit Issuance and Entry. Panel B of Table 4 reports the parameter estimates that quantify these constraints to trade.[31] We estimate that issued credits q_j are best-rationalized as *unbiased* population-wide estimates ($\alpha \approx 1$) of net CO_2 reductions Δ_j that are subject to error ($\sigma_\varepsilon > 0$). In other words, credit issuance may be scientifically correct *on average*, but ultimately biased among those who select into trade (as highlighted in Section 3). Second, Panel B documents the relevance of entry transaction costs.

Figure 5 illustrates the implications of these estimates by plotting the average and marginal effectiveness of credits at varying market prices, $AE(Q(p))$ and $ME(Q(p))$. The estimates rationalize the low effectiveness ratios documented in Section 5.1 and highlight the relevance of adverse selection in this setting. The entry condition (equation (10)) clears only when credit revenue is large relative to cost. At a low price, this may require a large positive error ε_j or a low opportunity cost ω_j, which is positively correlated with Δ_j. Entry at low prices is therefore positively selected on q_j and negatively selected on Δ_j, yielding low values of $AE(Q(p))$ and $ME(Q(p))$. As prices rises, entering projects do not need as many excess credits to be profitable, so both curves slope upwards, as in Figure 2.

Demand. Panel C of Table 4 reports select parameter estimates of credit demand (equation (15)). First, ϕ, the coefficient on CO_2 reductions, estimated using the triple-difference interaction (see Section 5.3), is positive across specifications. Second, the coefficient on price is negative. These coefficients reflect related conclusions to Section 5.3. The willingness-to-pay for expected reductions in tons of CO_2 (units of Δ) under the *status quo* is positive but modest. Buyers are very price elastic (in units of q). The *status quo* is therefore characterized by a positive willingness-to-pay for credits above and beyond the (modest) willingness-to-pay for expected reductions in tons of CO_2. These results reflect the weak regulatory environment of the *status quo* and highlight the empirical relevance of the trade-sustaining force in the stylized theoretical framework (Figure 2b). However, these estimated values are not structural parameters, but rather a reflection of the *status quo* (weak) regulatory environment.

7 Counterfactual Market Designs

Our estimates imply that tropical forests offer substantial abatement at costs below current compliance market prices, yet the *status quo* market does not realize these gains from trade. Active proposals intend to reform market design to achieve this potential (see Appendix A). Regulatory instruments are available: alternative designs could enforce (1) *unbiased*

[31] Appendix Figure A.14 assesses model fit by comparing model-implied and data moments. Appendix Table A.10 reports estimates under alternative specifications.

trade and (2) valuations of CO_2 reductions benchmarked to higher compliance market prices. However, the potential for adverse selection persists across market designs. It is theoretically ambiguous whether proposed changes would realize or unravel the potential gains from trade.

7.1 The *Status Quo* versus an Opt-in Compliance Market

We use our parameter estimates to address this question. To do so, we construct an empirical analogue of Figure 2 (Figure 6), contrasting the *status quo* with a stylized counterfactual opt-in compliance market.

We calculate the (inverse) credit supply curve $P_S(Q)$ using the estimates in Table 4, separately by sector s. We suppress the sector s subscript. At each credit price p, we simulate credit issuance (equation (9)) and entry (equation (10)), calculate the credits issued to entrants, $Q(p) = \sum_{j \in s} q_j x_j(p)$, their true CO_2 reductions, $R(Q(p)) = \sum_{j \in s} \Delta_j x_j(p)$, and the average and marginal effectiveness of credits, $AE(Q) = R(Q)/Q$ and $ME(Q) = R'(Q)$. These correspond to the curves in Figure 5 with a change of variables. Buyers observe a credit's sector but not its project-specific effectiveness, so $AE(Q)$ pools all entrants within a sector. Figure 6 presents results for REDD+ (87% of credits in the *status quo*). Sub-figure (b) also plots the social value of a marginal increase in credits, $\Phi\, ME(Q)$, evaluated at $\Phi = \$60$. As in Figure 2, the sub-figures differ only in the demand curve implied by the regulatory environment.

Figure 6a plots the *status quo* estimated REDD+ demand curve implied by the estimates in equation (15),

$$P_D(Q) = \frac{\phi}{|\gamma|}\, AE(Q) + \underbrace{\frac{1}{|\gamma|}\left[\xi - \ln\left(\frac{Q}{M-Q}\right)\right]}_{\psi(Q)}, \tag{16}$$

where M denotes the potential market size and ξ is a sector-level demand shifter that reproduces the observed *status quo* market share of sector s. The coefficient on $AE(Q)$ is buyers' *revealed* willingness to pay for a one-ton reduction in CO_2 (Table 4, Panel C), and $\psi(Q)$ is the marginal buyer's willingness to pay for a nominal credit, which declines in Q with heterogeneous buyers. The intersection of $P_D(Q)$ and $P_S(Q)$ is the *status quo* equilibrium.

By contrast, Figure 6b plots demand in an opt-in compliance market. Covered emitters purchase a credit only if doing so is cheaper than paying a constant regulated price Φ, and unbiased trade is enforced through a trading ratio under which $1/AE(Q)$ credits offset one ton.[32] Demand is therefore $P_D(Q) = \Phi AE(Q)$.

[32] Holding fixed Φ means aggregate abatement increases with trade. A regulator could implement this by

The two sub-figures correspond to the two regimes highlighted in Section 3 and debated in current policy: first, the market as it operates today, and, second, the admission of offsets into compliance markets under rules intended to preserve *unbiased trade*.

***Biased Trade* vs. Unraveling.** Figure 6 highlights the empirical relevance of the tension between *biased trade* and unraveling. Table 5 reports additional outcomes and counterfactuals.

Figure 6a illustrates that the estimated *status quo* willingness to pay for nominal credits, independent of CO_2 reductions, lifts the demand curve and sustains trade at low prices and large quantities of credits Q, and modest true reductions Δ. This combines and rationalizes the observations in Section 5. Table 5 (column (1)) reports that although these reductions in CO_2 generate some environmental benefits (valued at Φ), participants incur costs that slightly outweigh them. Market participants gain from trading cheap credits.

Figure 6b presents outcomes under the stylized opt-in compliance market, which pivots the demand curve. Demand $\Phi AE(Q)$ is upwards-sloping due to the existence of adverse selection, and lies everywhere below the supply curve. At no price is it both profitable to enter and cost-effective to purchase a credit. Trade unravels. Of course, this result is stylized and involves substantial out-of-sample extrapolation. However, it highlights that reforms to improve the market by imposing unbiased trade may instead exacerbate the effects of adverse selection. Such reforms, when successful, eliminate the implicit subsidy facilitated by a willingness to tolerate biased trade $\psi \neq 0$ that made the marginal entrant profitable.

Further, the absence of trade is socially inefficient. This is apparent given the low-cost abatement opportunities highlighted by our land use model (Section 6.3) and others (Souza-Rodrigues, 2019; Jayachandran et al., 2017; Araujo et al., 2026; Mishra, 2026). It is also clear from the fact that the social value of a *marginally* supplied credit $\Phi\, ME(Q)$ exceeds its cost for some $Q > 0$ in Figure 6b. Low-cost reductions in CO_2 in principle exist but cannot be realized through opt-in trade alone. While trade may be socially efficient, it is not cost-effective due to adverse selection. Reforms that seek to improve the market may ultimately unravel it.

The graphs in Figures 6a and 6b are useful for intuition and to highlight a key tension, but focus on one specific contrast. Table 5 and Figure 7 explore the sensitivity of the conclusions. Switching between Figures 6a and 6b involves two changes: imposing unbiased trade ($\psi = 0$) and increasing ϕ to a high compliance market value, Φ. Imposing unbiased trade alone also

reducing the compliance market cap by the expected CO_2 reductions traded. In practice, this market design is referred to as "offsets under the cap."

causes the market to unravel (column (2) of Table 5). Figure 7 plots the total market size, in tons of true reductions in CO_2 traded, implemented by the opt-in compliance market at varying values of Φ. At valuations up to \$85, trade essentially unravels.

Policy Implications. We ground the stylized tension in Figure 6 in the context of an active policy proposal. The EU aims to meet 5% of its emissions reductions relative to 1990, or approximately 250 million tCO_2e, through the use of international offset credits (European Commission, 2025). This policy, along with similar examples in other countries, is motivated by cost-containment. Figure 8 plots three potential scenarios if the EU were to implement this goal through offset trade at Emissions Trading System (ETS) prices.

Scenario (1) considers a case in which credits are accepted at face value as tons of CO_2. Costs are low as intended (\$9 per ton, Figure 8b), but true reductions in CO_2 are only 10% of the stated cap met through offsets (Figure 8a). The EU would have met its cost-containment goal, but only through sacrificing some of its climate ambition. Scenario (2) implements outcomes imposing *unbiased trade*, requiring the mechanism to meet the policy's climate ambition. Trade occurs at ETS prices but is limited, meeting only 36% of the objective (Figure 8a). This preserves climate integrity but falls short on the cost-containment dimension.

The contrast between Figures 6a and 6b suggests a third potential scenario. The willingness to pay for a nominal credit acts as an implicit subsidy that sustains trade (the lifting of the demand curve). This implicit subsidy could be replaced by an *explicit subsidy* per ton of true abatement that makes the marginal entrant at the cap profitable. However, in contrast to offset trading mechanisms' frequent cost-containment motivations, this could entail a large fiscal cost. In scenario (3), we search for this explicit subsidy. Figure 8 highlights that implementing 250 million tons of CO_2 reductions would require a per-ton subsidy of \$15 (18%). This explicit subsidy replaces the implicit subsidy provided by a failure of carbon accounting.

7.2 Incentives for Information Acquisition

We conclude by asking why the extent of adverse selection is so severe. Estimating Δ_j, which depends in part on an unobserved counterfactual, is a genuinely challenging prediction problem. Anything short of a perfect predictor ($\sigma_\varepsilon > 0$) will introduce the potential for adverse selection. Still, innovations that improve prediction could potentially facilitate unbiased trade. Is the prediction problem intractable or is this innovation under-provided?

Figure 9 investigates this question under the *status quo* market design. It plots the payoffs,

relative to the *status quo*, to market participants as credit issuance becomes more precise ($\sigma_\varepsilon \to 0$). While improvements in precision are socially valuable, they are not privately valued under the *status quo* market design. Increasing precision lowers the number of over-credited projects who enter on an optimistic draw of ε_j and the excess credit revenue of those who enter. Buyers value the higher CO_2 reductions per credit less than the excess credits themselves. Fewer transactions reduce revenue to the intermediary. The private surplus from excess credits that sustains biased trade also reverses incentives for improvements in information.

Market designs that impose *unbiased trade* re-align these incentives. A market unraveled generates neither private surplus nor revenue to the intermediary. By committing to enforce an opt-in compliance market that imposes unbiased trade and a high valuation of CO_2 reductions, a regulator creates a large payoff (a "prize") for sufficient improvements in information. This payoff is the existence of demand, which creates private surplus and revenue to the intermediary, so long as adverse selection is mitigated sufficiently that unbiased trade can occur.

8 Conclusion

Our results suggest that offset markets have sustained trade because their design permits *biased trade*: buyers' willingness to pay for nominal credits acts as an implicit subsidy, at the cost of substitutability and potentially higher aggregate emissions. Overcrediting is therefore not a fixable accounting error but a structural feature of market design. Eliminating it would require a major, largely untested shift toward *unbiased trade*—and even then, trade could unravel or remain inefficiently low unless the implicit subsidy is replaced by explicit subsidies for genuine abatement or measurement improves substantially.

These reforms face a basic obstacle: overcrediting is precisely what makes offsets cheap. Buyers, sellers, and regulators seeking low-cost reductions all benefit from tolerating it, while explicit subsidies run counter to offsets' traditional appeal. Yet without such reforms, the gains from offset trade may not be realizable. Beyond offsets, our results highlight stark limits to voluntary, opt-in environmental action.

References

Akbar, Prottoy A., Victor Couture, Gilles Duranton, and Adam Storeygard, "The fast, the slow, and the congested: Urban transportation in rich and poor countries," Technical Report, Working Paper August 2023.

Akerlof, George A., "The Market for "Lemons": Quality Uncertainty and the Market Mechanism," *The Quarterly Journal of Economics*, 1970, *84* (3), 488–500.

AlliedOffsets, "AlliedOffsets API," https://alliedoffsets.com/api/ n.d. Accessed 25 August 2026.

Andreoni, James, "Privately Provided Public Goods in a Large Economy: The Limits of Altruism," *Journal of Public Economics*, 1988, *35* (1), 57–73.

_ , "Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence," *Journal of Political Economy*, 1989, *97* (6), 1447–1458.

_ , "Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving," *The Economic Journal*, 1990, *100* (401), 464–477.

Araujo, Rafael, Francisco Costa, and Marcelo Sant'Anna, "Efficient Conservation of the Brazilian Amazon: Estimates from a Dynamic Model," *The Review of Economic Studies*, 2026, *93* (1), 72–105.

Aspelund, Karl M. and Anna Russo, "Additionality and Asymmetric Information in Environmental Markets: Evidence from Conservation Auctions," *American Economic Review*, 2026, *116* (8), 2765–2805.

Assunção, Juliano J., Lars Peter Hansen, Todd Munson, and Jose Scheinkman, "Carbon Prices, Forest Conservation and Reforestation in the Brazilian Amazon," 2026. Working Paper.

Assunção, Juliano, Robert McMillan, Joshua Murphy, and Eduardo Souza-Rodrigues, "Optimal Environmental Targeting in the Amazon Rainforest," *The Review of Economic Studies*, 2023, *90* (4), 1608–1641.

Bénabou, Roland and Jean Tirole, "Individual and Corporate Social Responsibility," *Economica*, 2010, *77* (305), 1–19.

Bergstrom, Theodore, Lawrence Blume, and Hal Varian, "On the Private Provision of Public Goods," *Journal of Public Economics*, 1986, *29* (1), 25–49.

Berry, Steven T., "Estimating Discrete-Choice Models of Product Differentiation," *The RAND Journal of Economics*, 1994, *25* (2), 242–262.

BeZero Carbon, "Mapping the Market: BeZero's Global Boundaries Database," https://bezerocarbon.com/insights/mapping-the-market-bezeros-global-boundaries-database 2025. Accessed 25 August 2026.

_ , "BeZero Carbon Rating Methodology," https://bezerocarbon.com/ratings/resources n.d. Accessed 25 August 2026.

Blaufelder, Christopher, Cindy Levy, Peter Mannion, and Dickon Pinner, "A Blueprint for Scaling Voluntary Carbon Markets to Meet the Climate Challenge," McKinsey & Company January 2021. McKinsey Sustainability insights, January 29, 2021.

Calel, Raphael, Jonathan Colmer, Antoine Dechezleprêtre, and Matthieu Glachant, "Do Carbon Offsets Offset Carbon?," *American Economic Journal: Applied Economics*, 2025, *17* (1), 1–40.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The Effect of Minimum Wages on Low-Wage Jobs," *The Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

Chen, Qiaoyi, Nicholas Ryan, and Daniel Yi Xu, "Firm Selection and Growth in Carbon Offset Markets: Evidence from the Clean Development Mechanism," 2025. NBER Working Paper No. 33636.

Climate Impact Partners, "Setting the Pace: Carbon Credits and Climate Leadership in the Fortune Global 500," Report, Climate Impact Partners June 2026. Seventh annual analysis of Fortune Global 500 climate commitments; data collected 30 October 2025.

Conley, T. G., "GMM Estimation with Cross Sectional Dependence," *Journal of Econometrics*, 1999, *92* (1), 1–45.

Costinot, Arnaud, Dave Donaldson, and Cory Smith, "Evolving Comparative Advantage and the Impact of Climate Change in Agricultural Markets: Evidence from 1.7 Million Fields around the World," *Journal of Political Economy*, 2016, *124* (1), 205–248.

Cutler, David M. and Sarah J. Reber, "Paying for Health Insurance: The Trade-Off between Competition and Adverse Selection," *The Quarterly Journal of Economics*, 1998, *113* (2), 433–466.

Dappe, Matias Herrera, Mathilde Lebrand, and Aiga Stokenberga, *Shrinking Economic Distance: Understanding How Markets and Places Can Lower Transport Costs in Developing Countries*, Washington, DC: World Bank, 2024.

Dinerstein, Eric, David Olson, Anup Joshi, Carly Vynne, Neil D. Burgess, Eric Wikramanayake et al., "An Ecoregion-Based Approach to Protecting Half the Terrestrial Realm," *BioScience*, 2017, *67* (6), 534–545.

Einav, Liran, Amy Finkelstein, and Mark R. Cullen, "Estimating Welfare in Insurance Markets Using Variation in Prices," *The Quarterly Journal of Economics*, 2010, *125* (3), 877–921.

European Commission, "Proposal for a Regulation of the European Parliament and of the Council Amending Regulation (EU) 2021/1119 Establishing the Framework for Achieving Climate Neutrality," COM(2025) 524 final, 2025/0524(COD), Brussels, 2 July 2025 2025.

FAO, "FAOSTAT: Producer Prices," https://www.fao.org/faostat/en/#data/PP 2025. Accessed on 12 July 2026. Licence: CC-BY-4.0.

Farrokhi, Farid and Heitor S. Pellegrina, "Trade, Technology, and Agricultural Productivity," *Journal of Political Economy*, 2023, *131* (9), 2509–2555.

Fischer, G., F. O. Nachtergaele, H. T. van Velthuizen, F. Chiozza, G. Franceschini, M. Henry, D. Muchoney, and S. Tramberend, "Global Agro-Ecological Zones v4 – Model Documentation," Technical Report, FAO and IIASA, Rome 2021.

Forest Trends' Ecosystem Marketplace, "State of the Voluntary Carbon Market 2024: On the Path to Maturity," Technical Report, Forest Trends Association, Washington, DC May 2024.

Grupp, Tristan E., Prakash Mishra, Mathias Reynaert, and Arthur A. van Benthem, "An Evaluation of Protected Area Policies in the European Union," *The Review of Economic Studies*, 2026.

Guizar-Coutiño, Alejandro, Julia P. G. Jones, Andrew Balmford, Rachel Carmenta, and David A. Coomes, "A Global Evaluation of the Effectiveness of Voluntary REDD+ Projects at Reducing Deforestation and Degradation in the Moist Tropics," *Conservation Biology*, 2022, *36* (6), e13970.

Handel, Benjamin R., "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *American Economic Review*, 2013, *103* (7), 2643–2682.

Hansen, M. C., P. V. Potapov, R. Moore, M. Hancher, S. A. Turubanova, A. Tyukavina, D. Thau, S. V. Stehman, S. J. Goetz, T. R. Loveland, A. Kommareddy, A. Egorov, L. Chini, C. O. Justice, and J. R. G. Townshend, "High-Resolution Global Maps of 21st-Century Forest Cover Change," *Science*, 2013, *342* (6160), 850–853.

Hendren, Nathaniel, "Private Information and Insurance Rejections," *Econometrica*, 2013, *81* (5), 1713–1762.

Herbst, Daniel and Nathaniel Hendren, "Opportunity Unraveled: Private Information and the Missing Markets for Financing Human Capital," *American Economic Review*, 2024, *114* (7), 2024–2072.

Intergovernmental Panel on Climate Change, *Good Practice Guidance for Land Use, Land-Use Change and Forestry*, Hayama, Japan: Institute for Global Environmental Strategies (IGES), 2003. Edited by Jim Penman and others.

International Carbon Action Partnership, "Offset Use Across Emissions Trading Systems," Technical Report, Berlin, Germany 2023.

Jayachandran, Seema, Joost de Laat, Eric F. Lambin, Charlotte Y. Stanton, Robin Audy, and Nancy E. Thomas, "Cash for Carbon: A Randomized Trial of Payments for Ecosystem Services to Reduce Deforestation," *Science*, 2017, *357* (6348), 267–273.

Karnik, Akshata, John B. Kilbride, Tristan R. H. Goodbody, Rachael Ross, and Elias Ayrey, "An open-access database of nature-based carbon offset project boundaries," *Scientific Data*, April 2025, *12* (1), 581.

Kitzmueller, Markus and Jay Shimshack, "Economic Perspectives on Corporate Social Responsibility," *Journal of Economic Literature*, 2012, *50* (1), 51–84.

Kotchen, Matthew J., "Green Markets and Private Provision of Public Goods," *Journal of Political Economy*, 2006, *114* (4), 816–845.

_ **and Andrew Vogt**, "Who Wants Carbon Offsets? And Why It Matters," 2026. Working Paper.

_ **and Michael R. Moore**, "Private Provision of Environmental Public Goods: Household Participation in Green-Electricity Programs," *Journal of Environmental Economics and Management*, 2007, *53* (1), 1–16.

Kwon, Soonwoo and Jonathan Roth, "(Empirical) Bayes Approaches to Parallel Trends," *AEA Papers and Proceedings*, 2024, *114*, 606–609.

Macintosh, Andrew, Gregory Trencher, Benedict Probst, Shanta Barley, Danny Cullenward, Thales West, Don Butler, and Johan Rockström, "Carbon Credits Are Failing to Help with Climate Change—Here's Why," *Nature*, 2025, *646*, 543–546.

Meijer, Johan R., Mark A. J. Huijbregts, Kees C. G. J. Schotten, and Aafke M. Schipper, "Global Patterns of Current and Future Road Infrastructure," *Environmental Research Letters*, 2018, *13* (6), 064006.

Mishra, Prakash, "The Global Allocative Efficiency of Deforestation," 2026. Working Paper.

National Geospatial-Intelligence Agency, "World Port Index (Pub. 150), 21st Edition," Shapefile, https://msi.nga.mil/Publications/WPI 2011. Source shapefile dated 10 March 2011.

Nelson, Andy, Daniel J. Weiss, Jacob van Etten, Andrea Cattaneo, Theresa S. McMenomy, and Jawoo Koo, "A suite of global accessibility indicators," *Scientific Data*, November 2019, *6* (1), 266.

Pande, Rohini, Robin Burgess, Maryam Farboodi, Agnes Norris Keiller, and Lucy Page, "Draft Proposal for a Unified Carbon Market," 2025. Working Paper.

Polyakova, Maria, "Regulation of Insurance with Adverse Selection and Switching Costs: Evidence from Medicare Part D," *American Economic Journal: Applied Economics*, 2016, *8* (3), 165–195.

Probst, Benedict S., Malte Toetzke, Andreas Kontoleon, Laura Díaz Anadón, Jan C. Minx, Barbara K. Haya, Lambert Schneider, Philipp A. Trotter, Thales A. P. West, Annelise Gill-Wiehl, and Volker H. Hoffmann, "Systematic assessment of the achieved emission reductions of carbon crediting projects," *Nature Communications*, 2024, *15* (1), 9562.

Rambachan, Ashesh and Jonathan Roth, "A More Credible Approach to Parallel Trends," *The Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Renoster, "Renoster Carbon Project Ratings," https://mercury.renoster.co/ 2025. Accessed 10 April 2025.

Robinson, Nathaniel, C. Ronnie Drever, David A. Gibbs, Kristine Lister, Adriane Esquivel-Muelbert, Viola Heinrich, Philippe Ciais, Celso H. L. Silva-Junior, Zhihua Liu, Thomas A. M. Pugh, Sassan Saatchi, Yidi Xu, and Susan C. Cook-Patton, "Protect young secondary forests for optimum carbon removal," *Nature Climate Change*, July 2025, *15* (7), 793–800.

Souza-Rodrigues, Eduardo, "Deforestation in the Amazon: A Unified Framework for Estimation and Policy Analysis," *The Review of Economic Studies*, 2019, *86* (6), 2713–2744.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

The Guardian, "Revealed: More Than 90% of Rainforest Carbon Offsets by Biggest Provider Are Worthless, Analysis Shows," 2023. Joint investigation with Die Zeit and SourceMaterial. https://www.theguardian.com/environment/2023/jan/18/revealed-forest-carbon-offsets-biggest-provider-worthless-verra-aoe.

U.S. Department of the Treasury, U.S. Department of Agriculture, U.S. Department of Energy, and White House Office of Domestic Climate Policy, "Voluntary Carbon Markets Joint Policy Statement and Principles," Policy Statement, U.S. Department of the Treasury, Washington, DC May 2024.

Vancutsem, C., F. Achard, J.-F. Pekel, G. Vieilledent, S. Carboni, D. Simonetti, J. Gallego, L. E. O. C. Arag ao, and R. Nasi, "Long-term (1990–2019) Monitoring of Forest Cover Changes in the Humid Tropics," *Science Advances*, 2021, *7* (10), eabe1603.

Verra, "Verra Registry," https://registry.verra.org/ n.d. Accessed 25 August 2026.

Walters, Christopher, "Empirical Bayes Methods in Labor Economics," in "Handbook of Labor Economics," Elsevier, 2024.

West, Thales A. P., Jan Börner, Erin O. Sills, and Andreas Kontoleon, "Overstated Carbon Emission Reductions from Voluntary REDD+ Projects in the Brazilian Amazon," *Proceedings of the National Academy of Sciences*, 2020, *117* (39), 24188–24194.

_ , Sven Wunder, Erin O. Sills, Jan Börner, Sami W. Rifai, Alexandra N. Neidermeier, Gabriel P. Frey, and Andreas Kontoleon, "Action Needed to Make Carbon Offsets from Forest Conservation Work for Climate Change Mitigation," *Science*, 2023, *381* (6660), 873–877.

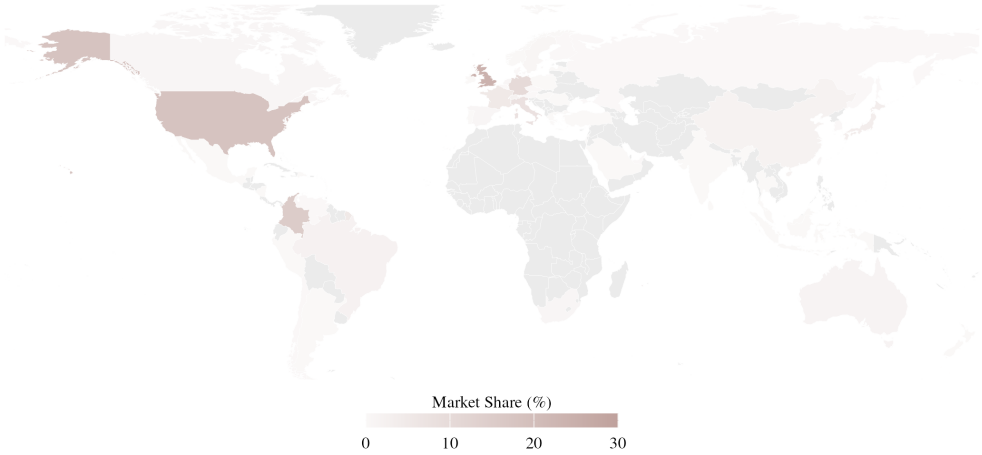
World Bank, "State and Trends of Carbon Pricing 2026," Technical Report, Washington, DC 2026.

Xpansiv, "Trusted Market Data for Environmental Commodities," https://www.xpansiv.com/data n.d. Accessed 25 August 2026.

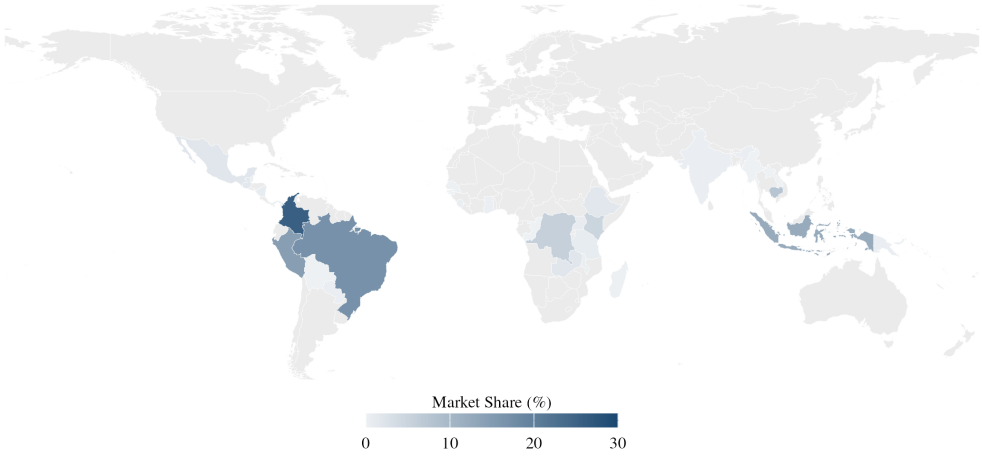
Figures and Tables

Figure 1: The Geography of Offset Credit Trade

(a) Buyers of Offset Credits

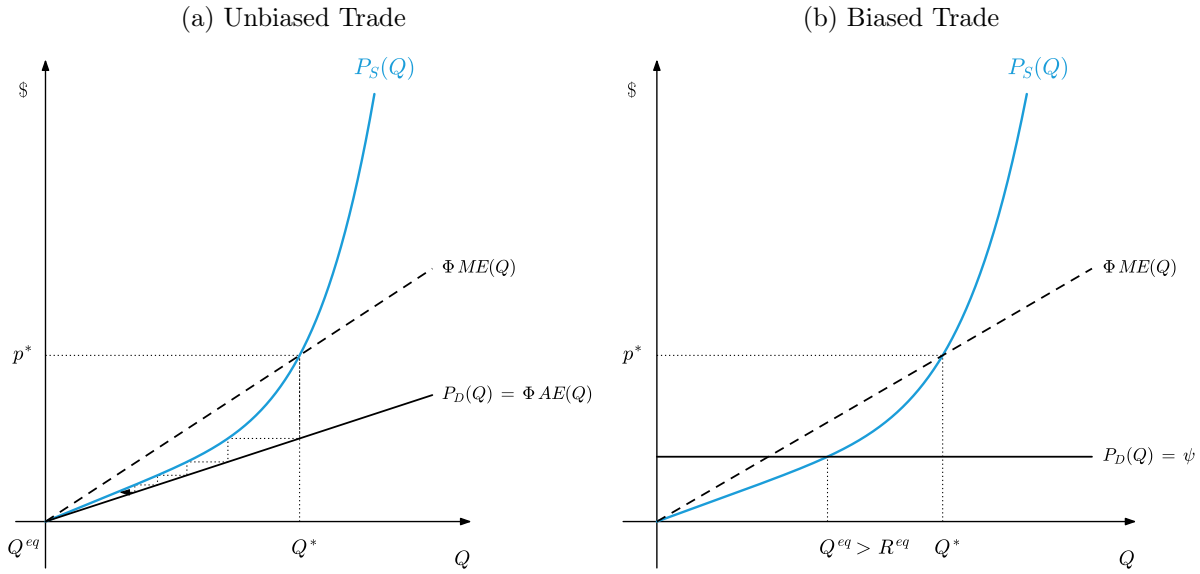


(b) Issuances of Offset Credits



Notes: Sub-figure (a) maps the market share of all offset credits purchased (retired) by country and sub-figure (b) maps the market share of credits issued by county for the tropical forest offset projects our sample. Sub-figure (a) aggregates retired credit volume by headquarter country of buyer.

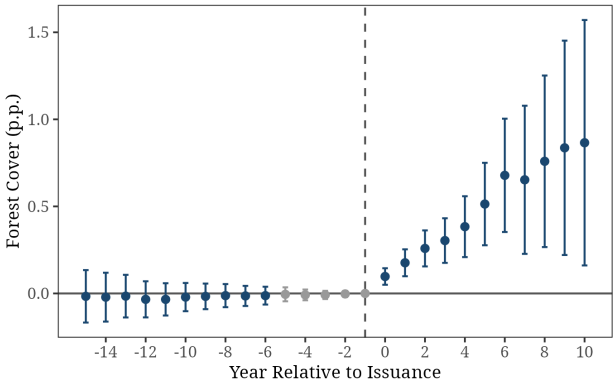
Figure 2: Unbiased versus Biased Trade



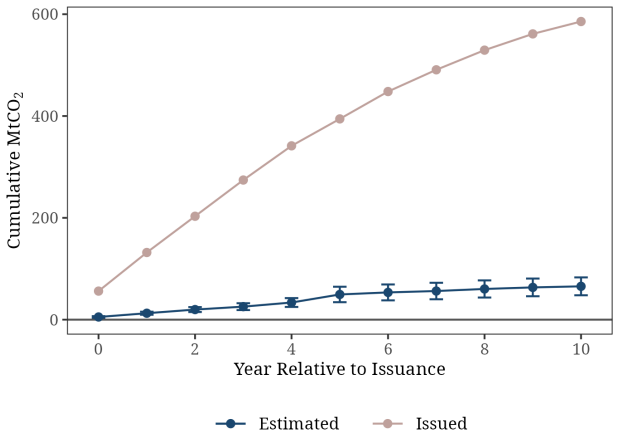
Notes: Figures illustrate equilibrium outcomes under two different regulatory environments. In both sub-figures, the x-axis is the quantity of credits Q and the y-axis is dollars per credit; $P_S(Q)$ is the inverse supply curve of credits, and $\Phi\, ME(Q)$ is the social value of the marginally supplied credit, where Φ is the social value per ton of CO_2. In both sub-figures, the socially efficient quantity is represented by Q^* at price p^*. Sub-figure (a) represents a scenario in which unbiased trade is enforced, so the demand curve is $P_D(Q) = \Phi\, AE(Q)$, the buyer's per-credit value of the expected reduction in CO_2; the illustration sets $\phi = \Phi$, so buyers and the social planner value CO_2 reductions equally and the two curves differ only in marginal versus average effectiveness. $P_D(Q)$ is upward-sloping because it conditions on the selection equation (2). In sub-figure (b), biased trade is tolerated and buyers value only nominal credits, so $P_D(Q) = \psi$, their willingness to pay per credit, which does not condition on the selection equation (2). In sub-figure (a), $P_D(Q)$ lies below $P_S(Q)$ at every $Q > 0$ and trade unravels to $Q^{eq} = 0$, traced by the dotted path descending from (Q^*, p^*). In sub-figure (b), trade occurs at $Q^{eq} > 0$, but the quantity of credits supplied Q^{eq} exceeds total CO_2 reductions R^{eq}.

Figure 3: Treatment Effects versus Credit Issuance

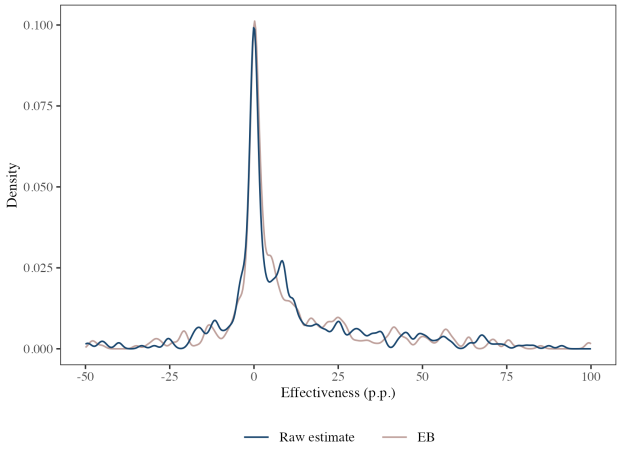
(a) Pooled Event Study Estimates on Forest Cover



(b) CO_2 Reductions (Δ) versus Issued Credits (q)



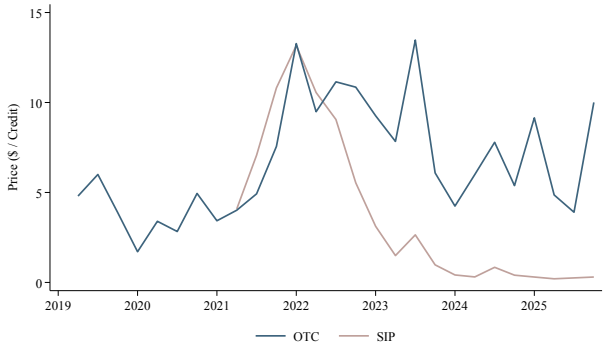
(c) Distribution of project effectiveness (Δ_j/q_j)



Notes: Sub-figure (a) plots estimates of the event study coefficients in equation (5), pooled across both REDD+ and ARR projects. Each observation is weighted by the area of the corresponding treated grid cell. Standard errors based on Conley (1999) using a uniform kernel and 50 km radius cutoff. Sub-figure (b) converts project-level treatment effects on tree cover (using coefficients from equation (6)), into estimated CO_2 reductions (see Appendix D) and compares them to issued credits. Sub-figure (c) plots the density of the ratio of estimated total CO_2 reductions to issued credits, displaying both the raw estimates and the estimates after the Empirical Bayes adjustment (detailed in Appendix E.2).

Figure 4: Evidence of Unraveling in Pooled (SIP) versus Separated (OTC) Trades

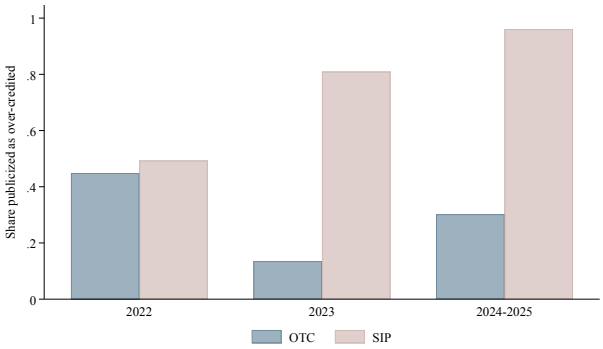
(a) Price Series



(b) Composition: Average OTC Price

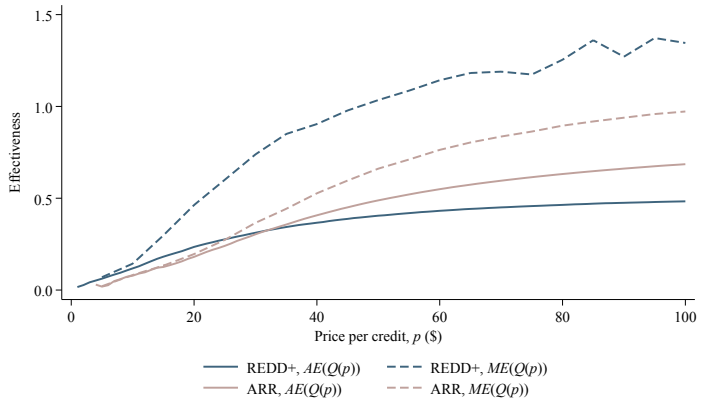


(c) Composition: Publicly Revealed Over-credited



Notes: Sub-figure (a) plots the credit-weighted average OTC price of SIP-eligible credits and the SIP (pooled contract) price. Sub-figures (b) and (c) compare the composition of credits selling OTC (separated) versus via the SIP (pooled). Sub-figure (b) plots the average credit-weighted OTC price of credits. Sub-figure (c) plots the share of credits transacted that were publicly revealed to be over-credited (see Section 5.3 for more details). Data on credit composition in the SIP begins in 2022. The years 2024 and 2025 are pooled given low transacted credit volumes in the SIP in 2025.

Figure 5: CO_2 Reductions vs. Issued Credits by Price

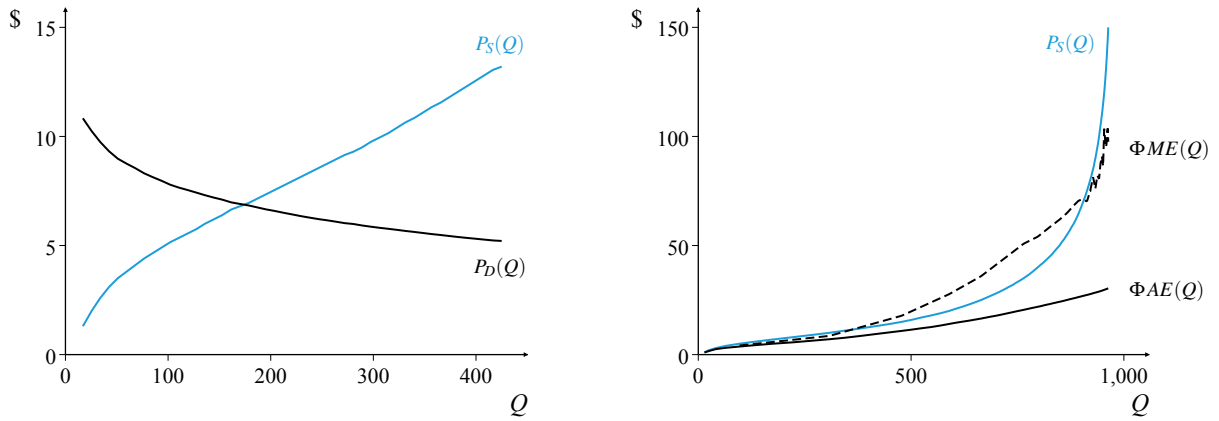


Notes: Figure plots simulated effectiveness, the ratio of true CO_2 reductions to issued credits, against the credit price p, separately for REDD+ and ARR. Solid lines plot average effectiveness among credits supplied at price p, $AE(Q(p)) = R(Q(p))/Q(p)$. Dashed lines plot the effectiveness of the marginally supplied credit, $ME(Q(p)) = R'(Q(p))$. At each price, credit issuance is simulated from equation (9) and entry from equation (10), given the parameter estimates in Table 4.

Figure 6: Empirically-Calibrated Graphs

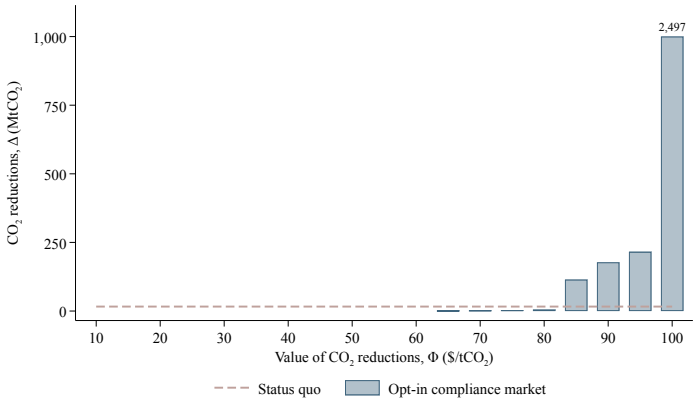
(a) Status Quo

(b) Opt-in Compliance Market



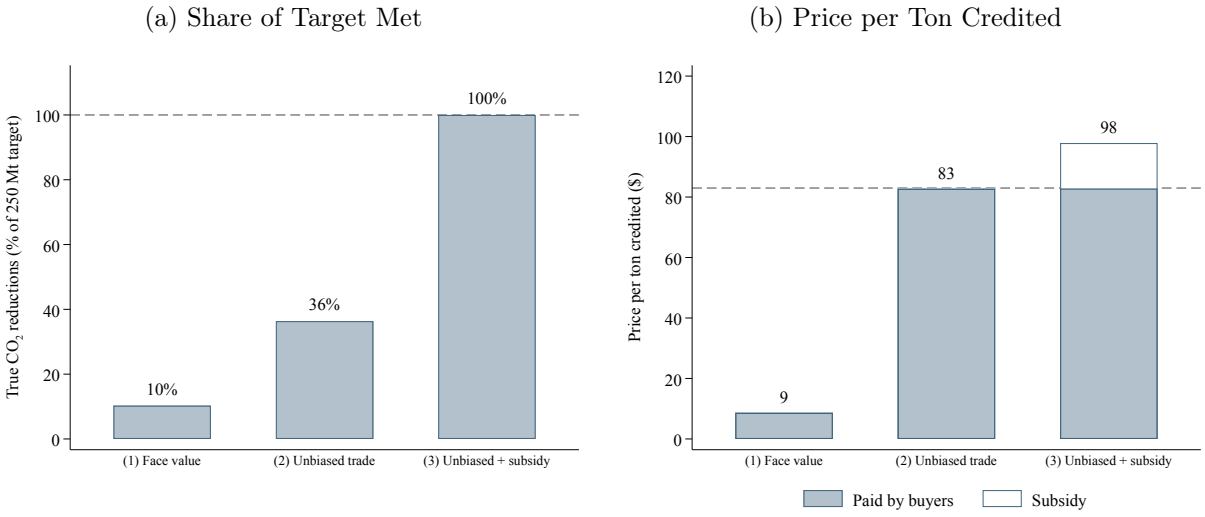
Notes: Figure plots the empirical analogue of Figure 2. In both sub-figures, the x-axis is the quantity of credits Q (millions) and the y-axis is dollars per credit. $P_S(Q)$ is the inverse supply curve of credits, constructed by simulating credit issuance (equation (9)) and entry (equation (10)) at each price, given the parameter estimates in Table 4. Sub-figure (a) plots status quo demand, $P_D(Q)$, from equation (15), using the demand estimates in Table 4. Results plotted for REDD+ (87% of credits in the *status quo*). Sub-figure (b) plots an opt-in compliance market in which unbiased trade is enforced, so demand is $\Phi\, AE(Q)$ with $\Phi = \$60$.

Figure 7: Market Size: Heterogeneity by Φ



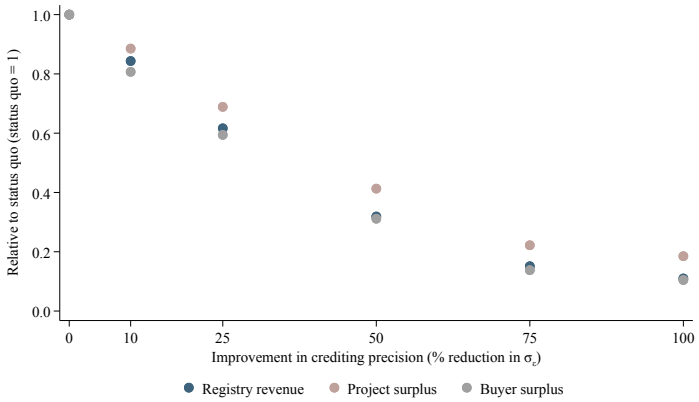
Notes: Figure plots market size, measured as true CO_2 reductions Δ ($MtCO_2$), at varying values of Φ in \$5 increments. Bars plot outcomes in the opt-in compliance market, in which unbiased trade is enforced and demand is $\Phi\, AE(Q)$, as in Figure 6b. The dashed line plots *status quo* market size, which does not depend on Φ. Results pool REDD+ and ARR. Bars are capped at 1,000 $MtCO_2$ for readability.

Figure 8: Application: EU International Offsets Policy



Notes: Figure presents outcomes of three scenarios in which the EU meets a target of 250 $MtCO_2e$ through international offset credits. Scenario (1) accepts credits at nominal face value, so each credit counts as one ton toward the target. Scenario (2) enforces unbiased trade in an opt-in compliance market at the EU ETS price. Scenario (3) adds the smallest per-ton subsidy at which the market supplies the target. Sub-figure (a) reports true CO_2 reductions as a share of the target; the dashed line marks the target. Sub-figure (b) reports the price per ton credited toward the target. In scenario (1) this is the credit price; under unbiased trade it is the price per true ton. The dashed line marks the EU ETS price. Credit supply is simulated from equations (9) and (10) given the parameter estimates in Table 4.

Figure 9: Incentives for Information Acquisition



Notes: Figure plots payoffs to market participants under the status quo market design as the precision of credit issuance improves. On the x-axis, an improvement of x percent multiplies the standard deviation of crediting error σ_ε by $(100 - x)/100$, so that $x = 0$ is the status quo and $x = 100$ is perfect information. Registry revenue is the per-credit registry fee ($0.20) times credits issued. Project surplus is credit revenue net of costs $(\omega_j + \chi_j)$ among entering projects. Buyer surplus is the estimated consumer surplus from credit purchase. Values are normalized to 1 at current precision. Results pool REDD+ and ARR and are simulated at the parameter estimates in Table 4.

Table 1: Summary Statistics

	Analysis Sample	Other Tropical Land	Transactions
	(1)	(2)	(3)
Panel A. Project Information			
Entry year	2015.1		
	(5.4)		
Credits issued (per year)	180,385		
	(482,003)		
Size (ha)	69,461		
	(236,074)		
Sector: REDD+ (%)	30.2		
Sector: ARR (%)	69.8		
N projects	450		
Panel B. Forest Cover and Geography			
Share forested	82.9	31.2	
	(32.7)	(39.8)	
Deforestation rate	0.04	0.10	
	(0.23)	(0.47)	
Travel time to nearest port or city (hours)	24.4	6.0	
	(16.8)	(10.1)	
Carbon / ha	119.4	61.8	
	(36.1)	(49.5)	
Carbon regrowth / ha	1.84	1.25	
	(0.50)	(0.59)	
Total area (ha)	31,257,373	3,514,016,759	
Panel C. Credit Transaction Prices			
Price per credit (USD)			7.02
			(4.03)
Credits per transaction			20,957
			(73,992)
N transactions			3,018

Notes: Table presents summary statistics (standard deviations in parentheses). Panel A column (1) reports project information about the 450 projects in our primary analysis sample. Panel B compares the project locations for these projects (column (1)) to all other land in the tropics (column (2)). Statistics in Panel B are area-weighted. Forest cover is measured in 2014. The deforestation rate is the annual decline in forest share from 2010 to 2014, reported as percentage points of forest-cover loss per year. Carbon stock and maximum regrowth are measured per hectare of land in Mg C per ha and Mg C per ha per year, respectively. Panel C reports summary statistics for the analysis sample of transaction data (column (3)).

Table 2: Event Study Treatment Effects *versus* Issued Credits

	All (1)	REDD+ (2)	ARR (3)
Panel A. Forest Cover			
Average post-period coefficient	0.543	0.457	2.067
	(0.147)	(0.149)	(0.925)
Final year coefficient	0.866	0.777	2.224
	(0.360)	(0.375)	(1.372)
Panel B. CO_2 Accounting			
Implied CO_2 reduction ($MtCO_2$)	69.2	62.3	6.8
	(9.0)	(9.0)	(0.5)
Total credits issued ($MtCO_2e$)	615.9	534.0	81.9
Panel C. Effectiveness Estimates			
Project-level	0.202	0.213	0.197
	(0.650)	(0.734)	(0.611)
Credit-level	0.078	0.078	0.083
	(0.278)	(0.281)	(0.260)
N projects	450	136	314

Notes: Table presents estimates of project treatment effects and compares them to issued credits for all projects (column (1)) and separately by REDD+ (column (2)) and ARR (column (3)). Panel A reports coefficient estimates from equation (5). The first row averages coefficients from event years 1 through 10. The final-year row reports event year 10. Conley standard errors reported in parentheses. Panel B reports aggregate implied CO2 reductions, calculated using estimates from equation (6), multiplied by the relevant CO_2 accounting parameters, summed across projects (implied CO_2 reduction). Panel B also reports the total credits issued among those projects. Parentheses report standard errors. Panel C reports mean project-specific effectiveness, Δ_j/q_j, after applying the Empirical Bayes procedure described in Appendix E.2. Project-level rows weight projects equally; credit-level rows weight by total issuance. Parentheses in Panel C report standard deviations.

Table 3: Hedonic Estimates of Willingness to Pay for CO_2 Reductions

	No info.	All sources		Academic research		Numeric only		Implied WTP
		Cont.	Binary	Cont.	Binary	Cont.	Binary	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Price response to published effectiveness								
$\hat{\phi}$	0.547	4.588	1.500	3.415	3.099	6.237	4.711	
	(2.502)	(1.475)	(1.212)	(2.039)	(1.453)	(2.732)	(1.037)	
Panel B. Implied willingness to pay ($/credit)								
$\hat{\psi}$								6.50
								(2.40)
$\hat{\psi} + \hat{\phi}$								11.09
								(2.40)

Notes: Table presents estimates of equation (7). Column (1) regresses price on project effectiveness, Δ_j/q_j in the sample without published information ($Pub_j \cdot Post_t = 0$). Columns (2)–(7) report $\hat{\phi}$. Columns (2)–(3) pool all published estimates of $\tilde{\Delta}_j^{pub}$; columns (4)–(5) restrict to academic research; columns (6)–(7) restrict to estimates reported directly as numeric values, dropping letter grades (which in columns (2)–(3) are converted to numeric values using published guidelines). Within each pair, the first column enters $\tilde{\Delta}_j^{pub}$ continuously; the second uses an indicator for above-median $\tilde{\Delta}_j^{pub}$, divided by the difference in mean $\tilde{\Delta}_j^{pub}$ between above- and below-median projects so that it is on the same scale as the continuous estimate. Column (8) reports $\hat{\psi} = \bar{p}_{ljt} - \hat{\phi}\tilde{\Delta}_j^{pub}$, using $\hat{\phi}$ from column (2) and averaging across post-period project-years; $\hat{\psi} + \hat{\phi}$ is the willingness to pay for a credit that reduces one ton of CO_2. All specifications include fixed effects for transaction year, country, vintage, and subsector. Sample is AFOLU credits in the OTC data; observations are weighted by credits transacted. Panel A reports standard errors clustered by project in parentheses; Panel B reports standard deviations in parentheses.

Table 4: Parameter Estimates

	Supply		Demand	
	REDD+	ARR	Continuous	Binary
	(1)	(2)	(3)	(4)
Panel A. Land Use				
Price sensitivity, $1/\sigma_v$	0.0012	0.0021		
	(0.0003)	(0.0005)		
Average elasticity	2.58	4.65		
Panel B. Credit Issuance and Entry				
Shading parameter, α	1.01	1.01		
	(0.12)	(0.12)		
Estimation error, σ^s_ε	4.01	4.86		
	(1.23)	(0.58)		
Entry cost mean, μ^s_χ	50.49	302.04		
	(19.78)	(38.81)		
Entry cost s.d., σ^s_χ	23.36	70.03		
	(13.75)	(10.35)		
Panel C. Demand				
CO_2 reduction, ϕ			0.923	1.684
			(0.494)	(0.663)
Price, γ			−0.718	−0.583
			(0.363)	(0.281)
WTP for 1 tCO_2 reduction			1.285	2.888
			(0.745)	(1.097)
Average price elasticity			−5.324	−4.327

Notes: Table presents parameter estimates. Panel A reports estimates of the land use model based on equation (13). The unit of observation is a potential entrant. Agricultural revenue net of transport costs ($\bar{\pi}^a_j - \kappa_j$) is instrumented using the straight-line distance to the nearest city or port (Appendix Table A.9 compares OLS and reduced-form estimates). Panel B reports SMM estimates of the credit issuance and entry model (based on equations (9), (10), and (14)), given estimates from the land use model. The shading parameter α is imposed to be equal for REDD+ and ARR; other parameters are allowed to differ by sector (see Appendix Table A.10 for sector-specific α). Panel C reports demand estimates (equation (15)). The coefficient ϕ is estimated using the triple-difference variation in publicized effectiveness. Columns (3) and (4) parameterize the variation in published effectiveness as a continuous shock and as a rescaled binary shock (above versus below median), respectively. Price is instrumented by the crediting entity's distance to a city or port. The WTP for a 1 tCO_2 reduction is defined as $\phi/|\gamma|$. Standard errors in parentheses.

Table 5: Counterfactual Outcomes

	Status quo	Unbiased trade ($\psi = 0$)	Opt-in compliance ($\psi = 0$, $\Phi = 60$)
	(1)	(2)	(3)
Panel A. Market outcomes			
Price ($/credit)	7.43	–	–
Credits issued, Q (millions)	192.0	0	0
CO_2 reductions ($MtCO_2$)	16.3	0	0
Effectiveness	0.085	–	–
Panel B. Welfare (million $)			
Value of CO_2 reductions	975	0	0
Project costs	1,018	0	0
Buyer surplus	369	0	0
Project surplus	409	0	0
Registry revenue	38.4	0	0

Notes: Table presents simulated market outcomes under alternative market designs, pooling REDD+ and ARR. Column (1) reports the *status quo*. Column (2) imposes unbiased trade ($\psi = 0$) holding fixed the status quo valuation of CO_2 reductions. Column (3) reports an opt-in compliance market that imposes unbiased trade and values CO_2 reductions at $\Phi = \$60$ per ton. Credit supply is simulated from equations (9) and (10). Effectiveness is total CO_2 reductions divided by total credits issued. CO_2 reductions are valued at $60 per ton in all columns. Costs are entrants' opportunity and entry costs, $\omega_j + \chi_j$. Buyer surplus is the estimated consumer surplus from credit purchase. Project surplus is credit revenue net of costs. Registry revenue is the $0.20 per-credit registry fee times credits issued.

A Institutional Appendix

Market Overview. A wide range of project types can be transacted in offset markets. Table A.1 illustrates, tabulating the market shares of all offset projects issuing credits on over twenty registries.[33] Land use offsets, our focus, represent approximately 20% of the market. Other leading categories are renewable energy and industrial manufacturing projects. Figure A.1 documents that the general patterns of offset trade (lower and middle-income countries selling, higher-income countries buying) hold beyond our focus on tropical forest offsets.

Some buyers purchase credits voluntarily to meet net zero commitments, and other buyers purchase credits to meet obligations in a compliance market. About 34% of all global listed companies, 50% of large companies (over $10 billion market capitalization), and 44% of Fortune Global 500 Companies refer to the use of carbon credits to meet emissions targets.[34] Figure A.2 documents that a large fraction of global compliance schemes allow the use of offsets. Offsets are more prevalent in compliance markets with low carbon prices.

The registry is a central non-profit intermediary responsible for both issuing and maintaining a ledger of credits. Several registries exist (see Table A.1), but a single registry (Verra) represents a large fraction of the market of tropical forest projects and almost all REDD+ projects. Verra charges projects a modest flat listing fee and a per-issued-credit fee. The magnitude of the per-issued-credit fee has varied over time between $0.10 to $0.23 per credit. Due to Verra's dominance in tropical land use offsets, we take this fee as given in estimation in Section 6. The registry charges a de minimus fee to buyers ($0.02 per credit).

[33] Aggregated by AlliedOffsets; Appendix C provides details on data sources.

[34] Based on results reported in the Morgan Stanley Institute for Sustainable Investment report "Voluntary Carbon Markets: Surveying Current and Future Corporate Participants" and the Climate Impact Partners "Setting the Pace: Carbon Credits and Climate Leadership in the Fortune Global 500" for more details.

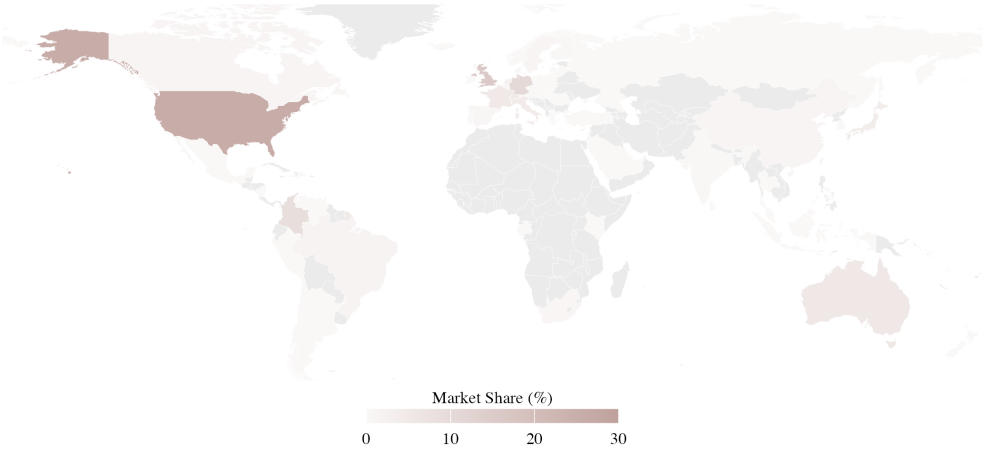
Table A.1: Market Shares by Project Category and Registry

	Issued Offset Credits (1)	Market Share (%) (2)	REDD+ Share (%) (3)	ARR Share (%) (4)
Panel A. Market shares by project category				
Renewable Energy	1,778,998,488	38.2		
Industrial Manufacturing	1,244,235,160	26.7		
Forestry and Land Use	897,346,200	19.3		
Waste Disposal	284,644,273	6.1		
Energy Efficiency	235,884,316	5.1		
Household Devices	164,774,623	3.5		
Other	47,592,989	1.0		
Agriculture	6,535,375	0.1		
Panel B. Registry market shares				
Verra			90.5	45.3
Clean Development Mechanism			0.0	17.1
Australian Carbon Credit Units			0.0	13.9
Gold Standard Impact Registry			0.0	9.1
ACR			0.0	8.2
BioCarbon Standard			9.5	6.4

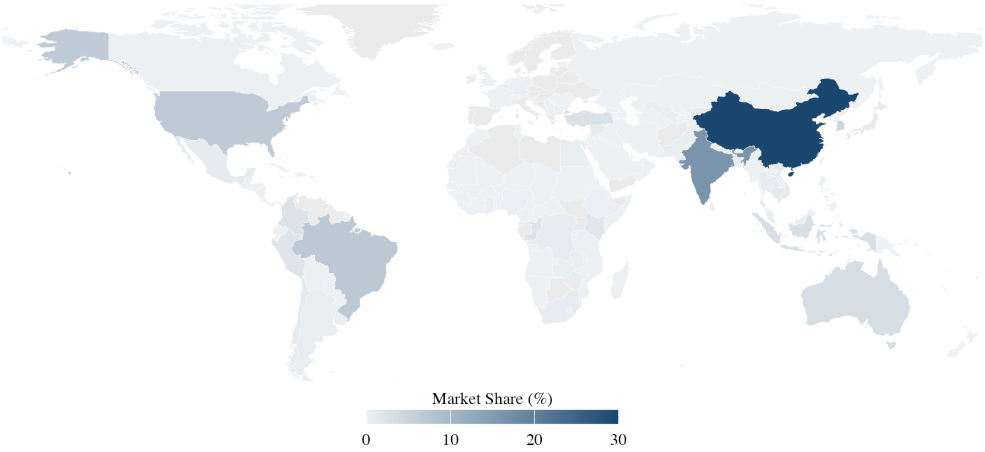
Notes: Panel A reports each sector's share of all issuances. Panel B reports registry shares within all REDD+ and ARR issuances. Shares are a percent of issued credits. Data from AlliedOffsets, which scrapes and aggregates data from over twenty major registries.

Figure A.1: The Geography of Offset Credit Trade: All Sectors

(a) Buyers of Offset Credits



(b) Issuances of Offset Credits



Notes: Sub-figure (a) maps the market share of all offset credits purchased (retired) by country and sub-figure (b) maps the market share of credits issued by county for all sectors. Sub-figure (a) aggregates retired credit volume by headquarter country of buyer.

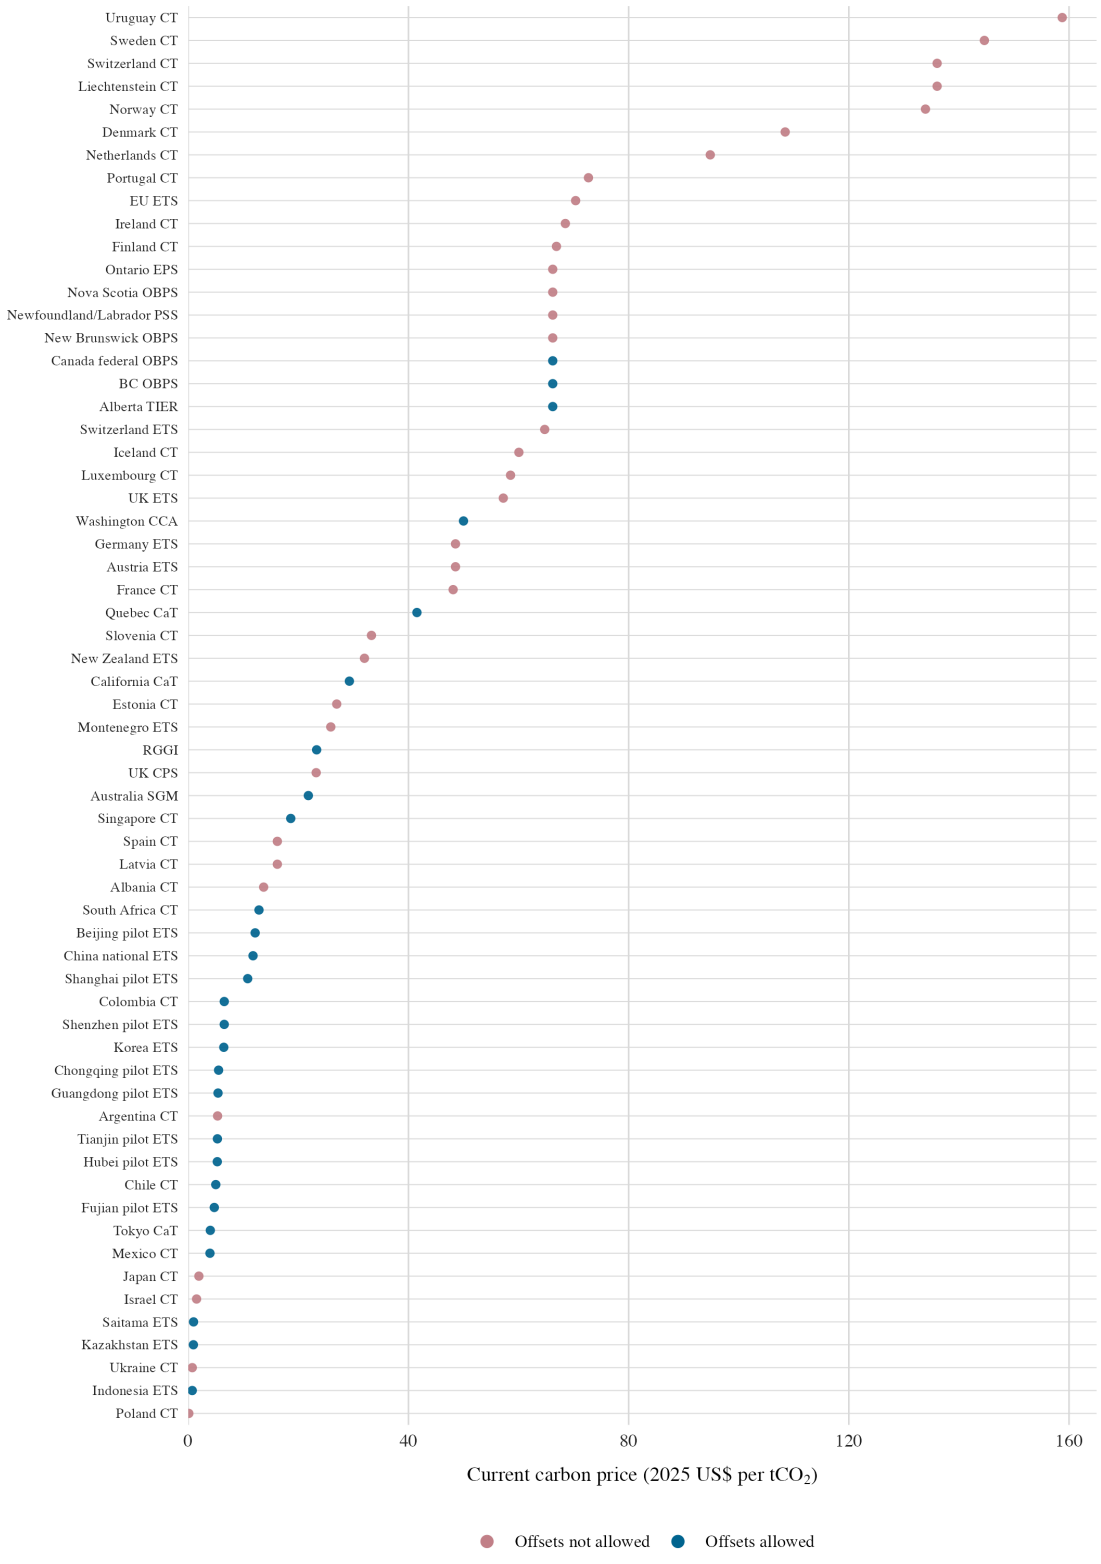


Figure A.2: Offset Eligibility in Compliance Markets

Notes: The figure reproduces the carbon-pricing comparison in Macintosh et al. (2025). Each point represents a carbon-pricing scheme. The horizontal axis reports the carbon price in 2025 US dollars per tCO_2e.

Trading: OTC vs. SIP. Offset credits typically transact over-the-counter (OTC), in which sellers and buyers of credits trade in decentralized markets.[35] Each credit is defined by a project-vintage (the year of the issued credit). A typical transaction will be for tens of thousands of credits (see Table 1). Financial platforms have emerged to facilitate trade. Xpansiv CBL, our data provider, is one such platform. It is self-described as "the world's largest and most active spot marketplace for environmental commodities." It provides information about credit prices and infrastructure to settle trades for OTC transactions. More than 300 $MtCO_2$ of carbon have been traded on Xpansiv CBL since 2020.

In April 2021, Xpansiv CBL launched a new product, the Nature Based Global Emissions Offset (N-GEO) under its Standard Instruments Program. The N-GEO was designed to be a high-quality, uniform, exchange-traded product. Specifically, the contract documents state:[36]

> "The contract provides Participants with a straightforward way to buy and sell offsets that meet specific criteria for quality and integrity. Each N-GEO delivers one Product meeting N-GEO Screening Criteria (as defined below) from an Approved Registry that is equivalent to the reduction of one tonne of carbon dioxide equivalent ("tCO2e"). The N-GEO spot contract simplifies carbon-offset procurement for market participants and enables a transparent, benchmark price signal for qualifying offset units."

These screening criteria are Verra's Verified Carbon Standard (VCS) AFOLU Project Type, CCB (Climate, Community, and Biodiversity) Label/Certification, and Project Vintage Eligibility. These screening criteria are designed to include only "high-quality" credits. The analyses in Section 5.2 condition on these eligibility criteria.

The N-GEO SIP is not a distinct product or derivative. It is a contracting infrastructure to physically deliver eligible credits that all trade at the N-GEO price. After the introduction of the N-GEO contract, both sellers and buyers of credits were presented a choice in their mode of trade. Both a seller and a buyer could choose to sell or buy their credits OTC or through the N-GEO SIP. A seller sees the price their credits could transact at selling OTC or through the N-GEO SIP. By contrast, a buyer purchasing credits through the N-GEO SIP only observes the price and the N-GEO label. Upon purchase, the buyer is delivered credits, differentiated by project-vintages, among the sub-set of credits being sold by credit sellers

[35] As in other OTC markets, brokers, a second type of intermediary, also play a role in facilitating trade. Because our transaction data are provided from the Xpansiv CBL platform which facilitates trade without brokerage, we do not consider the role of brokers in this paper's analysis.

[36] See www.xpansiv.com/wp-content/uploads/2025/12/Xpansiv_CBL_StandardInstruments.pdf for more details.

willing to sell via the N-GEO SIP at the N-GEO SIP price.[37]

Proposed Reforms. Despite—or in response to—concerns about the environmental integrity of offset trade, there exist a large number of governmental and quasi-governmental initiatives to *expand* trade through improved market design. Many of these proposals aim to move towards perfect substitutability within the existing offset-market architecture. Table A.2 provides a taxonomy of these enacted and proposed reforms with stated objectives of both perfect substitutability and expansion of trade.

[37]Sometimes a buyer and a seller will transact for specified credits (as in OTC trades), but will not want information about transacted prices to be revealed. These transactions are able to use the SIP infrastructure to hide these prices and reduce post-trade transparency. These transactions are excluded from our analysis.

Table A.2: Offset Policies with Stated Aims of Expanding Unbiased Trade

Reform	Institution	Year
Panel A: Compliance markets		
2040 target international credits	European Union	2026
Proposal for a Unified Carbon Market	Pande et al. (proposal)	2025
Article 6.4 mechanism standards	UNFCCC Supervisory Body	2024
CCER relaunch	China (MEE)	2024
International Carbon Credit Framework	Singapore (NEA)	2023
ACCU Scheme reforms	Australia (DCCEEW)	2023
Federal Greenhouse Gas Offset Credit System	Canada (ECCC)	2022
Emissions-unit eligibility (CORSIA)	ICAO	2019
Panel B: Voluntary market		
Principles for Voluntary Carbon and Nature Market Integrity	United Kingdom (DESNZ)	2024
Shared Principles for High-Integrity Credit Use	Coalition to Grow Carbon Markets	2025
Claims Code of Practice	VCMI	2025
Joint Policy Statement and Principles	US Treasury, USDA, DOE	2024
Carbon Removals Certification Framework	European Union	2024
Core Carbon Principles and Assessment Framework	ICVCM	2024
Purchase criteria	LEAF Coalition	2021

Notes: The table lists policies whose stated objective is to facilitate, expand or reform offset trade with credited quantities corresponding to true emissions reductions. Panels are split by the market a policy governs rather than by the issuing body, so government rules directed at the voluntary market appear in Panel B. Year denotes adoption or release. Several entries are multi-government initiatives. The Coalition to Grow Carbon Markets has eleven original endorsers: Canada, France, Luxembourg, New Zealand, Singapore, Switzerland, Kenya, Panama, Peru, Zambia, and the United Kingdom (Indonesia joined in January 2026). VCMI is endorsed by the United Kingdom government and the Children's Investment Fund Foundation, and its guidance has been supported by the governments of France, Panama, Kenya, Singapore, and the United Kingdom. The LEAF Coalition is financed by Norway, the United Kingdom, and the United States alongside corporate buyers.

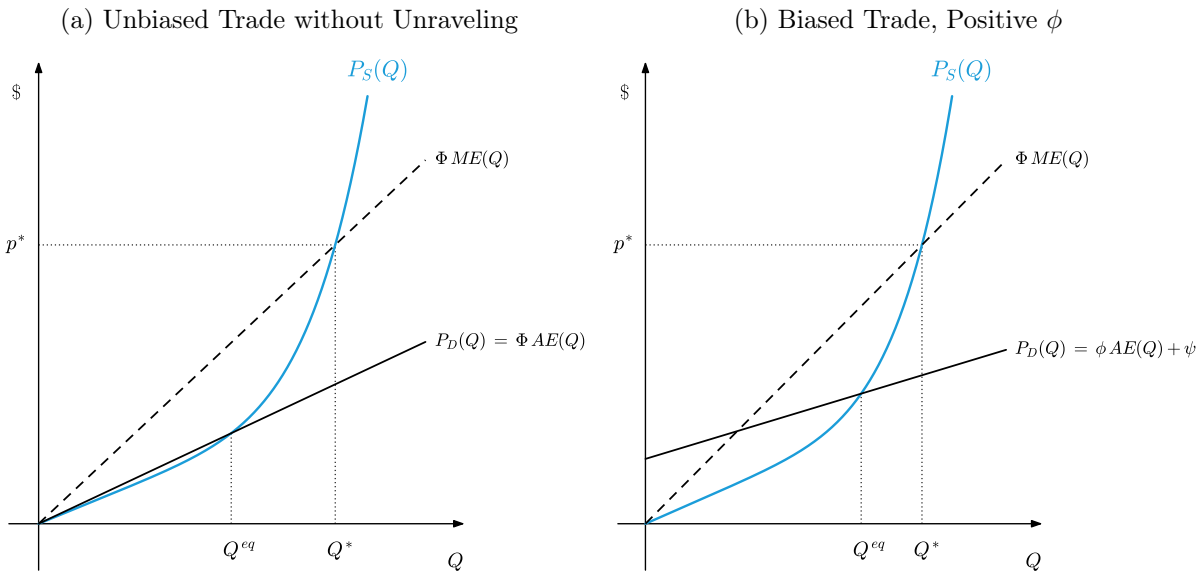
B Theory Appendix

B.1 Additional Graphical Analysis

Figure 2 in Section 3 displays outcomes under two different market designs. The illustration is intentionally stark: under unbiased trade (Figure 2a) the market unravels, while under

biased trade (Figure 2b) trade is sustained with low environmental integrity ($Q^{eq} > R^{eq}$). Figure A.3 illustrates the same juxtaposition of unbiased trade (Figure A.3a) and biased trade (Figure A.3b) under alternate parameterizations of supply and demand. In Figure A.3a, the unbiased trade demand curve $\Phi AE(Q)$ intersects the marginal cost curve at a positive value $Q^{eq} > 0$, so that the market does not completely unravel. However, trade is still limited: the equilibrium quantity (Q^{eq}, p^{eq}) is below the socially-optimal quantity and price (Q^*, p^*). Figure A.3b illustrates biased trade, with $\psi > 0$ and a non-zero willingness to pay for true CO_2 reductions $\phi > 0$. Similar to the case in Figure 2, ψ increases willingness to pay at every Q, resulting in more trade than in the unbiased case, but again, sacrificing perfect substitutability.

Figure A.3: Unbiased and Biased Trade Under Alternate Parameterizations



Notes: Figures illustrate equilibrium outcomes under two different regulatory environments. In both sub-figures, the x-axis is the quantity of credits Q and the y-axis is dollars per credit; $P_S(Q)$ is the inverse supply curve of credits, and $\Phi ME(Q)$ is the social value of the marginally supplied credit, where Φ is the social value per ton of CO_2. In both sub-figures, the socially efficient quantity is represented by Q^* at price p^*. Sub-figure (a) represents a scenario in which unbiased trade is enforced, so the demand curve is $P_D(Q) = \Phi AE(Q)$, the buyer's per-credit value of the expected reduction in CO_2; the illustration sets $\phi = \Phi$, so buyers and the social planner value CO_2 reductions equally and the two curves differ only in marginal versus average effectiveness. $P_D(Q)$ is upward-sloping because it conditions on the selection equation (2). In sub-figure (b), biased trade is tolerated and buyers value both expected reductions and nominal credits, so $P_D(Q) = \phi AE(Q) + \psi$, their willingness to pay per credit, where $\phi < \Phi$. In sub-figure (a), $P_D(Q)$ lies above $P_S(Q)$ for $Q < Q^{eq}$, so the market does not unravel and trade occurs at $Q^{eq} > 0$; because $P_D(Q)$ lies below $\Phi ME(Q)$, the equilibrium quantity falls short of Q^*. In sub-figure (b), demand for nominal credits raises willingness to pay at every Q.

B.2 Additional Details

Excess Crediting. It will be useful to define difference between credits issued and true CO_2 reductions, *excess crediting* $\mathcal{E}(Q)$:

$$\mathcal{E}(Q) = Q(p) - R(Q).$$

Proposition 1. *Assume the crediting error is conditionally independent of* (Δ_j, c_j) : $\mathbb{E}[\varepsilon_j|\Delta_j, c_j] = 0$ *with finite conditional variance. Holding fixed* $p > 0$*, (i) excess crediting is strictly increasing in the variance of* ε_j*,* $\mathbb{V}[\varepsilon_j]$ *(measurement error) and (ii) excess crediting is nonnegative.*

Proof. Excess crediting is given by the crediting error ε_j among entrants $\Theta(p)$:

$$\mathcal{E}(p) = \int_{\Theta(p)} (q_j - \Delta_j)\, d\mu(j) = \int_{\Theta(p)} \varepsilon_j \, d\mu(j).$$

Entry requires $p(\Delta_j + \varepsilon_j) \geq c_j$, i.e., a crediting error above the threshold

$$\varepsilon_j^* = \frac{c_j - p\Delta_j}{p},$$

which depends only on fundamentals. To parameterize an increase in $\mathbb{V}[\varepsilon_j]$, scale the error by $\sigma > 0$: conditional on (Δ_j, c_j), the standardized error ε_j/σ has fixed distribution G_j with density g_j, $\mathbb{E}[\varepsilon_j \mid \Delta_j, c_j] = 0$, and finite conditional variance, so that $\mathbb{V}[\varepsilon_j]$ is strictly increasing in σ. Then, per the law of iterated expectations:

$$\mathcal{E}(p) = \int_j \mathbb{E}\big[\varepsilon_j \, \mathbf{1}\{\varepsilon_j \geq \varepsilon_j^*\} \mid \Delta_j, c_j\big]\, d\mu(j) = \sigma \int_j M_j(\varepsilon_j^*/\sigma)\, d\mu(j),$$

where $M_j(t) \equiv \int_t^\infty u \, dG_j(u)$. Differentiating in σ, and writing $t_j \equiv \varepsilon_j^*/\sigma$:

$$\frac{\partial \mathcal{E}(p)}{\partial \sigma} = \int_j M_j(t_j)\, d\mu(j) - \int_j M_j'(t_j)\, t_j \, d\mu(j).$$

M_j' can be derived via the Fundamental Theorem of Calculus as $M_j'(t) = -t\, g_j(t)$. Substituting,

$$\frac{\partial \mathcal{E}(p)}{\partial \sigma} = \int_j M_j(t_j)\, d\mu(j) + \int_j t_j^2\, g_j(t_j)\, d\mu(j) \geq 0.$$

Positivity of the first term follows from $\mathbb{E}[\varepsilon_j \mid \Delta_j, c_j] = 0$, which gives $M_j(t) = -\int_{-\infty}^{t} u\, dG_j(u)$ for each j. For $t \leq 0$ the integrand is non-positive on $(-\infty, t]$, so $M_j(t) \geq 0$. For $t > 0$ the integrand of $\int_t^{\infty} u\, dG_j(u)$ is positive, so $M_j(t) \geq 0$ as well. The second term is a square times a density and is non-negative for any t_j. The sum is strictly positive whenever t_j falls in the interior of the support of G_j for a positive measure of projects.

Nonnegativity follows immediately from conditional independence:

$$\mathcal{E}(p) = \int_j \mathbb{E}\big[\varepsilon_j \mathbf{1}\{\varepsilon_j \geq \varepsilon_j^*\} \,\big|\, \Delta_j, c_j\big]\, d\mu(j) = \sigma \int_j M_j(t_j)\, d\mu(j) \geq 0,$$

□

Adverse Selection. Adverse selection occurs when $ME(Q) > AE(Q)$; a sufficient condition is that the marginal effectiveness curve slopes upward, $ME'(Q) > 0$ (see Einav et al. (2010)). Figure 2 (and Appendix Figure A.3) assumed the existence of adverse selection for illustration, though it need not be the case. The following proposition characterizes when the marginal effectiveness curve slopes upward in the model in Section 3.

Proposition 2. *Define $z_j \equiv c_j/q_j$ as the breakeven price of project j, so that $\Theta(p) = \{j : z_j \leq p\}$, and define $\varrho(p) \equiv \mathbb{E}[c_j \mid z_j = p]/\mathbb{E}[\Delta_j \mid z_j = p]$, the cost per true CO_2 reduction among the marginal entrants at price p. Then, for any $Q > 0$, evaluated at $p = P_S(Q)$,*

$$ME(Q) = \frac{P_S(Q)}{\varrho(P_S(Q))},$$

Marginal effectiveness slopes upward, $ME'(Q) > 0$, if and only if $\varrho(p)$ rises less than proportionally with the price: $\frac{d\ln \varrho(p)}{d\ln p} < 1$.

Proof. The marginal entrant at price p is indifferent about trade: $pq_j = c_j$ on $\{z_j = p\}$, so conditional on $z_j = p$, $q_j = c_j/p$ and $\mathbb{E}[q_j \mid z_j = p] = \mathbb{E}[c_j \mid z_j = p]/p$. Differentiating R and Q with respect to price,

$$ME(Q) = \frac{dR/dp}{dQ/dp} = \frac{\mathbb{E}[\Delta_j \mid z_j = p]}{\mathbb{E}[q_j \mid z_j = p]} = \frac{p}{\varrho(p)}.$$

Differentiating in p, $\frac{d}{dp} ME(Q(p)) = \frac{1}{\varrho(p)}\big(1 - \frac{d\ln \varrho(p)}{d\ln p}\big)$, which is positive if and only if the elasticity of ϱ is below one; since $dQ/dp > 0$, $ME'(Q) > 0$. □

Examples.

No Measurement Error. If $\varepsilon_j \equiv 0$, then $\varrho(p) = p$ and $ME \equiv 1$: there is no adverse selection.

Only Measurement Error. Suppose all projects are identical in fundamentals: $\Delta_j = \Delta > 0$ and $c_j = c > 0$, so projects differ only in ε_j. Project j enters when credit revenue covers costs, $pq_j \geq c$, or:

$$\varepsilon_j \geq \frac{c}{p} - \Delta.$$

At low prices, only projects with large positive errors find entry profitable. The marginal entrant holds exactly $q_j = c/p$ credits and reduces Δ tons, while every inframarginal entrant holds more, $q_j \geq c/p$. Therefore:

$$ME(Q(p)) = \frac{p\Delta}{c} \geq AE(Q(p)) = \frac{\Delta}{\mathbb{E}[q_j \mid \Theta(p)]}.$$

The ME curve slopes upward in price, as each price increase draws in entrants with smaller errors, and the market is adversely selected at every price. If $p < c/\Delta$, entry requires $\varepsilon_j > 0$: every entrant overcredits ($q_j > \Delta$) and $AE(Q(p)) < 1$.

C Data Appendix

C.1 Data: Overview

We use several datasets throughout the analysis. Table A.3 provides an overview.

Table A.3: Data Sources and Variable Construction

Data source	Measure and construction	Role in the analysis
A. Project and registry data		
AlliedOffsets project and credit records (AlliedOffsets, n.d.)	Licensed, proprietary records on project characteristics, registry affiliation, crediting dates, issuances, retirements. Scraped, aggregated, and cleaned from over twenty registries.	Defines the project universe, treatment dates, and credited quantities. Buyer records identify buyer characteristics and headquarters locations.
B. Project-boundary sources		

Continued on next page

Data source	Measure and construction	Role in the analysis
Karnik et al. carbon-project boundaries (Karnik et al., 2025)	Georeferenced project-area boundaries (publicly available).	Provides the preferred treatment geometry and determines project area and treated-cell coverage when available.
BeZero project boundaries (BeZero Carbon, 2025)	Georeferenced project-area boundaries (proprietary).	Supplies the second-priority treatment geometry when the preferred accounting area is unavailable.
AlliedOffsets shapefiles (AlliedOffsets, n.d.)	Georeferenced project-area boundaries (proprietary).	Supplies the third-priority treatment geometry.
Verra registry KML files (Verra, n.d.)	Project boundaries recovered from Verra project pages and supporting documents (publicly available).	Supplies the fallback treatment geometry.
C. Forest outcomes and carbon accounting		
JRC Tropical Moist Forest annual change	Annual tropical moist forest classes at 30 m by 30 m resolution. Class 1 is undisturbed forest, class 2 is degraded forest, and class 4 is forest regrowth. Within each cell-year, we also calculate the shares classified as deforested land, other land cover, and water. (Vancutsem et al., 2021)	Classes 1, 2, and 4 define annual forest share throughout the analyses.
IPCC land-use guidance (Intergovernmental Panel on Climate Change, 2003)	Default biomass-stock parameters by country and continent.	Supplies the stock parameter used in the REDD+ CO_2 accounting.
Chapman–Richards forest-regrowth parameters (Robinson et al., 2025)	Approximately 1-km global rasters of the Chapman–Richards asymptote, scaling factor, and growth-rate parameters for naturally regenerating forests.	Supplies the time-varying carbon-growth path used in the ARR CO_2 accounting.

Continued on next page

Data source	Measure and construction	Role in the analysis
D. Geographic variables		
RESOLVE Ecoregions 2017 (Dinerstein et al., 2017)	Global terrestrial ecoregions grouped into 14 biomes.	Restricts donor cells to the biome or biomes intersecting the project and defines country-by-biome groups.
Global Roads Inventory Project (Meijer et al., 2018)	Georeferenced road segments classified by road type. We calculate centroid distance to the nearest main road (types 1–3) and local road (types 4–5).	Main- and local-road distances enter the matching specification.
NGA World Port Index (National Geospatial-Intelligence Agency, 2011)	Port locations and characteristics. We calculate each cell's shortest distance to the nearest port.	Port distance enters the matching specification and is used to construct an instrument for transport costs and credit price in Section 6.
City locations	Coordinates for more than 1,200 large cities in 152 countries, obtained from the authors of Akbar et al. (2023). We calculate the shortest distance from each cell to the nearest city.	City distance enters the matching specification and is used to construct an instrument for transport costs in Section 6.
Global travel time to cities and ports	Mean travel time (as of 2015) from each cell to any city or port (Nelson et al., 2019).	Used to calculate transportation costs.
E. Agricultural yields and prices		
FAO/IIASA GAEZ v4 (Fischer et al., 2021)	Cell means of potential crop yield by crop, water regime, and input level. The main construction uses rain-fed, high-input yields. Conversion factors from Table 9-1 of the GAEZ v4 Model Documentation map yields reported in dry weight, oil, or sugar units into fresh-weight metric tons per hectare.	Measures agricultural productivity and the opportunity cost of candidate project locations.

Continued on next page

Data source	Measure and construction	Role in the analysis
FAOSTAT Producer Prices (FAO, 2025)	Country-crop farm-gate producer prices in US dollars per metric ton for 1998, 2014, and 2018.	Values potential crop yields to construct gross agricultural revenue and profit.
F. Transport costs		
World Bank trucking costs	Country-level road-freight costs in US dollars per metric ton-kilometer from Figure 1.6 (p. 28) of Herrera Dappe et al. (2024).	Combines with travel time to measure the cost of moving agricultural output to cities and ports.
G. Transaction data		
Xpansiv trades and deliveries (Xpansiv, n.d.)	Licensed transaction records containing trade and delivery identifiers, dates, instruments (SIP versus OTC), project identifiers when disclosed, vintages, prices, quantities, and delivered credits.	Used in unraveling evidence (Section 5.2), the hedonic analysis (Section 5.3), and demand estimation (Section 6).
H. Public effectiveness estimates		
Published "offset-achievement ratios" (Probst et al., 2024)	Project-level effectiveness estimates compiled from published evaluations of avoided-deforestation projects.	Measures publicly disclosed project effectiveness and dates of the information revelation used in the hedonic analysis (Section 5.3), and demand estimation (Section 6).
Renoster ratings (Renoster, 2025) and BeZero ratings (BeZero Carbon, n.d.)	Project-level third-party effectiveness assessments. Renoster values were recorded from the provider's ratings page, while BeZero ratings are included in the licensed Xpansiv extract. The ratings are harmonized to numerical measures in the analysis.	Measures publicly disclosed project effectiveness and dates of the information revelation used in the hedonic analysis (Section 5.3), and demand estimation (Section 6).

C.2 Data: Additional Details

We provide additional details about datasets included in the overview table above.

Project boundaries. Offset registries report project boundaries in digital format with measurement error.[38] We collect four sources of project-level boundary information, ordered by the degree of confidence in the resulting shapefiles: Karnik et al. (2025) hand-cleaned boundaries (cross-referenced with registry documentation), BeZero's proprietary cleaned shapefiles (cross-referenced with registry documentation), Allied Offsets proprietary cleaned shapefiles (cross-referenced with registry documentation), and scraped geometries from the Verra registry itself. We use the highest-confidence dataset available for each shapefile.

Projects included in our sample must adhere to four inclusion criteria. Projects must:

1. have a nonempty shapefile in at least one of our four spatial data sources;
2. have non-missing first issuing date and total credit volume;
3. be REDD+ or ARR, dropping other Agricultural, Forestry, or Other Land Use subcategories.
4. have $\geq 90\%$ of area overlap with TMF (Vancutsem et al., 2021) coverage, and have less than 25% area coded as class 0 (which encodes missing data).

Table A.4 begins with all forestry and land use projects from AlliedOffsets and reports the projects retained after sequentially excluding Australian Carbon Credit Unit Scheme (ACCU) projects, requiring a nonempty candidate boundary, requiring project metadata on issuing date and credits, restricting to REDD+ or ARR, imposing TMF coverage requirements, validating the selected geometry, and requiring a usable forest-cover outcome. The geometry screen retains 452 projects. We exclude two ARR projects whose treated and matched-control forest-cover series are identically zero, leaving 450 projects in the analysis sample.

[38] Evidence of measurement error includes: invalid topographies, mismatch between project descriptions and mapped areas, and project boundaries which are too generous (e.g., entire countries).

Table A.4: Project Sample Construction

Sequential sample screen	Retained	Dropped	Share of initial (%)	Share of positive-credit AFOLU projects (%)	Share of AFOLU issued credits (%)
All forestry and land use projects	8,252		100.0	100.0	100.0
Exclude ACCU projects	6,968	1,284	84.4	77.9	91.7
Nonempty boundary	2,231	4,737	27.0	44.7	85.4
Non-missing issuing date and credits	994	1,237	12.0	44.7	85.4
REDD+ or ARR	753	241	9.1	33.8	70.1
TMF coverage	471	282	5.7	21.2	55.5
Valid geometry	452	19	5.5	20.3	51.6
Usable forest-cover outcome	450	2	5.5	20.2	51.4

Notes: The starting set is all projects in the forestry and land use sector (source: AlliedOffsets). Screens are cumulative, and "Dropped" is relative to the preceding row. The positive-credit AFOLU-project denominator is the 2,226 forestry and land use projects with positive cumulative issued credits. The AFOLU-credit denominator is the 1,199,096,224 credits issued by those projects. The geometry screen selects the first available geometry in the fixed hierarchy, checks that it is spatially valid and at least one hectare, and removes boundaries identified as incorrect during geometry review. The final outcome screen excludes VCS1071 and VCS1663 because forest cover is constant at zero throughout their treated and matched-control panels. The analysis sample contains 136 REDD+ and 314 ARR projects.

Grid construction. We combine the various sources (listed above) by constructing a global gridded dataset covering the region of Tropical Moist Forests (TMF) data (Vancutsem et al., 2021) coverage. We loosely refer to the coverage region as the tropical band (there is some coverage outside the geographic tropics). We divide the tropical band into a 1 kilometer-by-1 kilometer grid. Grids are clipped along project boundaries, so every final geometry is either 100% inside a project or 100% outside a project. We refer to an observation as a cell or grid cell, though clipped boundaries may be non-square.

Forest outcomes and carbon accounting. The TMF (Vancutsem et al., 2021) data classify land cover annually at 30 m by 30 m resolution based on Landsat imagery. We calculate forest share as the share of all pixels that are either disturbed moist forest (class 1), degraded forest (class 2), or regrowing forest (class 4).

We use the TMF dataset for three reasons. First, we aim to measure forest regrowth dynamics:

comparable satellite products such as Hansen et al. (2013) do not record gain over time. Second, nearly all AFOLU projects take place in the tropical band. Third, it provides the longest pre-period of these commonly used satellite data products, dating back to 1990.

To estimate CO_2 from forest-regrowth, we use estimates of the Chapman-Richard forest-regrowth parameters from Robinson et al. (2025). The study team collated in-situ data with explicit age, geolocation, and carbon/biomass measurements from Australia, Canada, the Netherlands, Sweden and the USA as well as data from previous literature, including on tropical basins. This meta-analysis yields 172,585 measurements, binned into 20 five-year age classes and averaged within ∼1-km grid cells, yielding 109,708 final estimates. These growth curves are then projected onto climatological variables using a random forest and spatially extrapolated to achieve global coverage.

Geographic variables. Travel times in Nelson et al. (2019) are derived by using a machine learning algorithm to generalize millions of Google Maps queries on the driving distance from randomly sampled points on Earth's surface to cities and ports (previously used in Mishra (2026)).

Agricultural yields and prices. Potential yield data is commonly used in economics as a measure of crop yields under "best case" cultivation conditions (see, e.g., Costinot et al. (2016); Farrokhi and Pellegrina (2023)). The data is based on an agronomic model of yields trained on crops grown under laboratory conditions. The agronomic model is applied to gridded climate and soil data to get cell-level measures of best-case yields.

We calculate the potential gross agricultural revenue (before transportation costs) per crop, k. We restrict attention to the following crops: maize, rice, wheat, soybeans, sugar cane, and oil palm. We convert from dry-weight potential yields to fresh-weight using Table 9-1 of the GAEZ v4 Model Documentation (Fischer et al., 2021).[39] Following Costinot et al. (2016), we compute 2014 FAO producer prices US dollars per metric ton and impute missing producer prices using a country λ_c and crop λ_k fixed effect. Estimates are weighted by country-crop production (also from FAO):

$$\log(p_c^k) = \lambda_c + \lambda_k + v_c^k$$

so that for a missing p_c^k, the imputed value is $\exp(\hat{\lambda}_c + \hat{\lambda}_k)$.

[39] GAEZ reports a range of 0.87–0.90 for dryland rice: we use the midpoint, 0.885. To maintain variation in geographic suitability, we exclude highland maize as a genetic variant of maize.

Transport costs. We compute transport costs for grid g producing crop k in country c:

$$\kappa_g^k = \min(\text{travel time to city}_g, \text{travel time to port}_g) \times \text{freight rate}_c \times 60 \times \text{potential yield}_g^k$$

where freight rates are reported in Herrera Dappe et al. (2024) and 60 km/hour is our assumed average travel speed.

Transaction data. We licensed a proprietary historical extract of Xpansiv CBL's trade tape as of October 2025. Transactions are either over-the-counter (OTC) or via the Standard Instruments Program (SIP). The data includes project identifiers, prices, dates, and transaction volumes for all OTC trades. For SIP trades, the data include transaction volumes, dates, and prices. The data also record the credits delivered in each SIP transaction (unobserved to the buyer at the time of purchase). Table A.5 reports the transaction counts, number of unique projects, credit volumes, and credit-weighted average price in the raw dataset and our final analysis sample for each of the OTC and the SIP markets. 93% of SIP projects by volume have a corresponding, matched transaction of the same project on the OTC market (the delivery data row of Table A.5), allowing us to assign a counterfactual OTC price (a quantity-weighted, project-year level price per OTC trade).

Table A.5: Xpansiv transaction data summary statistics

	Transactions	Unique projects	Credit volume per transaction	Avg. price ($/credit)
	(1)	(2)	(3)	(4)
Panel A. Over-the-counter (OTC)				
Raw data	4,621	648	50,006	4.23
AFOLU analysis sample	1,646	111	35,890	7.05
Panel B. Standard Instruments Program (SIP)				
Raw data	3,879		22,281	4.67
N-GEO analysis sample	1,372		3,042	6.65
Matched to delivery data	1,335	23	1,526	
Panel C. With public effectiveness estimates				
Probst et al. (2024)	546	25	42,498	6.56
Renoster	628	12	41,632	8.23
BeZero	1,284	44	35,801	7.10
Any of the above	1,328	55	35,595	7.04

Notes: Table presents summary statistics of the Xpansiv transaction data sample. Panel A summarizes all raw OTC transactions and the analysis sample of AFOLU credits. Panel B summarizes all SIP transactions and the AFOLU SIP analysis sample (the N-GEO instrument). The SIP analysis sample restricts to on-market transactions. Panel B also summarizes the matched delivery data, the dataset of credits delivered via the SIP. Delivery data begins January 15, 2022. Panel C summarizes credit transactions with public effectiveness estimates and includes all transactions for projects that have a public effectiveness estimate at any point in the sample period.

Public effectiveness estimates. We compile public estimates of effectiveness, the ratio of CO_2 reductions to credits issued, from three sources.

1. Probst et al. (2024) is a publicly available (at `www.carboncredits.fyi`) meta-analysis of the effectiveness of 42 avoided deforestation projects. Merging into the OTC sample yields matched trades for 25 unique projects.

2. 25 projects have Renoster third-party ratings covering 12 unique projects in our sample.

3. BeZero ratings are reported in Xpansiv directly. We recode letter graded ratings following BeZero interpretation guidelines as quantitative effectiveness estimates: A through AA- are coded as 100% effective, B through BBB are coded as 50% effective, and all other ratings as 0%.

Panel C reports counts and average prices in the merged sample of OTC trades with each effectiveness source. When we run pooled regressions across effectiveness estimates in Section 5.3, we prioritize estimates as Probst et al. (2024), then Renoster ratings (which are explicitly quantitative), and finally BeZero ratings. Collective coverage, the final row of Table A.5, is 1,328 trades and 55 projects. We do not use our effectiveness estimates in this sample as they were not part of the information shock in 2022–2023.

D Carbon Accounting

We calculate avoided carbon emissions from avoided deforestation for REDD+ projects or carbon sequestration by regrowing trees for ARR projects.

REDD+. We calculate the stock of biomass (carbon) which would have been burned absent an offset project using Intergovernmental Panel on Climate Change (2003) (Good Practice Guidance for LULUCF, Table 3A.1.4), and convert it into CO_2. The Guidance assumes that without the offset project, deforestation amounts to a complete burn. For project j in country c, the biomass density h_c^{REDD} is measured in tons of dry biomass matter per hectare.[40] Define the change in forest cover attributable to a carbon offset project as $\Delta\theta_j$, measured in hectares. The corresponding change in carbon dioxide is Δm_j:

$$\Delta m_j^{REDD} = \frac{44}{12} \times 0.47 \times 1.25 \times h_c^{REDD} \times \Delta\theta_j.$$

The last two terms give us the additional units of dry biomass matter due to the offset project. The 1.25 factor is a "root-to-shoot" ratio which accounts for the fact that belowground biomass is also burned when trees are burned: the usual IPCC benchmark is to assume that an amount of 25% of aboveground biomass is removed from belowground carbon stocks. Finally, the first two terms convert from dry biomass to carbon dioxide emissions: 0.47 is the IPCC standard conversion factor from dry biomass to carbon, and $\frac{44}{12}$ is the ratio of molecular weights of CO_2 and carbon, so that the final metric Δm_j is measured in tons of CO_2.

ARR. We require a method which accounts for the gradual growth of young trees. Define the total area of young trees which is added by project j over T years as $\Delta\theta_j$. The carbon conversion factor $h_j^{ARR}(T)$ must account for a *growth curve* of young trees to accurately capture carbon sequestration, and thus depends on growth time T. Following Robinson et al. (2025), we use the Chapman-Richards (CR) growth curve to model concave carbon growth in

[40] For one country, Vanuatu, the IPCC provides no guidance on h_c^{REDD} so we use a continent-level mean.

tree age. The CR curve models biomass growth (not dry matter, in contrast to h_c^{REDD}) with three grid g-specific parameters:

$$B_g(t) \equiv A_g \cdot [1 - b_g \exp(-k_g t)]^{1/(1-m).}$$

Here, A_g is the asymptotic level of biomass; k_g is the growth rate; b_g is a scaling parameter; and m governs the degree of concavity. Robinson et al. (2025) estimate a 1 kilometer-grid level Chapman–Richards curve using spatially predicted above-ground carbon densities for naturally regenerating forests aged 1–100 years. $B_g(T)$ represents the total predicted aboveground carbon density (in metric tons carbon per hectare) accumulated up to time T by grid g.[41] The final step aggregates grid-level carbon to a project-level CO_2 measure:

$$\Delta m_j^{ARR} = \Delta\theta_j \times \sum_{g \in j} \frac{44}{12} \times 1.25 \times B_g(T) \times \frac{A_{jg}}{\sum_{j \in g} A_{jg}}.$$

E Empirical Appendix

E.1 Land-Use Event Studies

Staggered timing variation. Figure A.4a illustrates the timing variation across projects in the data. Figure A.4b reports the gap between the project's recorded founding year and its first issuance year. This delay, due to verification timelines and documentation delays, varies substantial around a median of six years. Delays are largely based on registry administrative capacity and the burden of evidence that the registry requests prior to first issuance.

Alternate specifications. Figure A.5 estimates equation (5) by separately sector. Figure A.6 estimates (5) with alternate sets of fixed effects. Figure A.7 re-weights projects so that every project receives equal weight. REDD+ projects are generally larger than ARR projects and dominate area-weighted effects in Figure 3a. Because REDD+ treatment effects are smaller, the unweighted treatment effect pooled across sectors is larger.

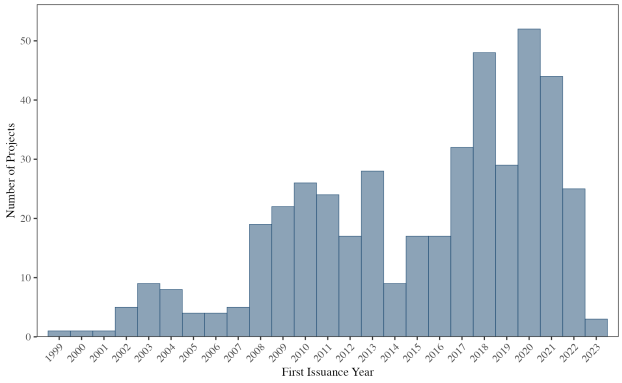
Heterogeneity in treatment effects. Figure A.8 plots a regression of our estimated effectiveness Δ_j/q_j estimates on several predictors.Observables explain only 5% of variation in estimated effectiveness. Projects which are closer to roads and thus more at risk of forest loss are no more effective, relative to the number of issued credits, than others (see also, Assunção et al. (2023); Grupp et al. (2026)). On average, ARR projects are more effective

[41] We maintain a root-shoot ratio of 1.25 to keep accounting consistent with REDD+. This assumes roots follow the same Chapman-Richards growth curve.

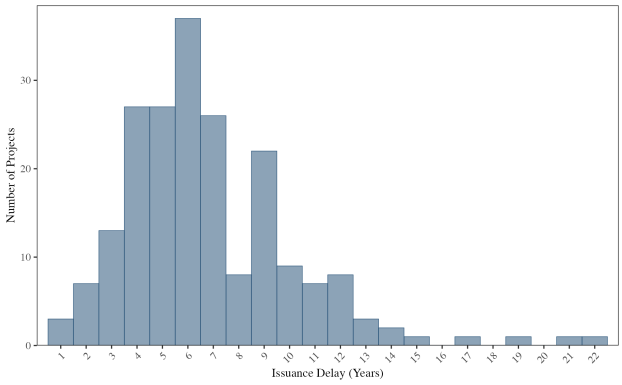
than REDD+ projects after controlling for observables. There is no evidence that project effectiveness has trended consistently over time in Figure A.9.

Figure A.4: Project Entry

(a) First issuance year



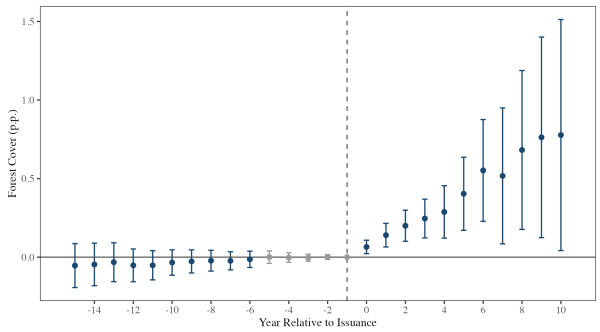
(b) Issuance Delay



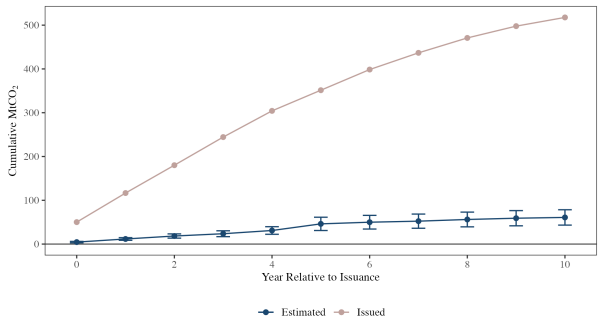
Notes: Figure A.4a plots the first year in which a project begins issuing credits. The sample comprises the 450 projects that enter event study (5). Figure A.4b plots the lag between the project's recorded founding year and its first registry issuance year.

Figure A.5: Treatment Effects by Sector

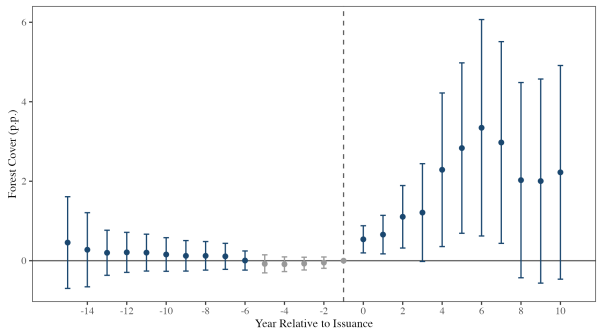
(a) REDD+: forest cover



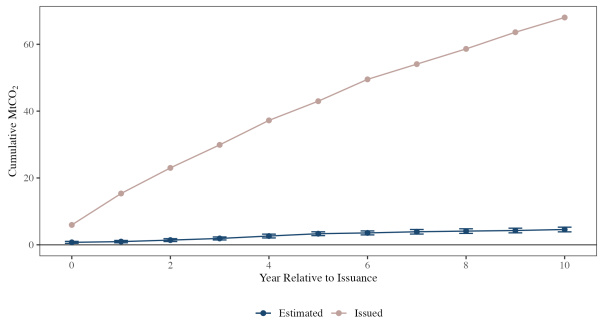
(b) REDD+: CO_2 accounting



(c) ARR: forest cover

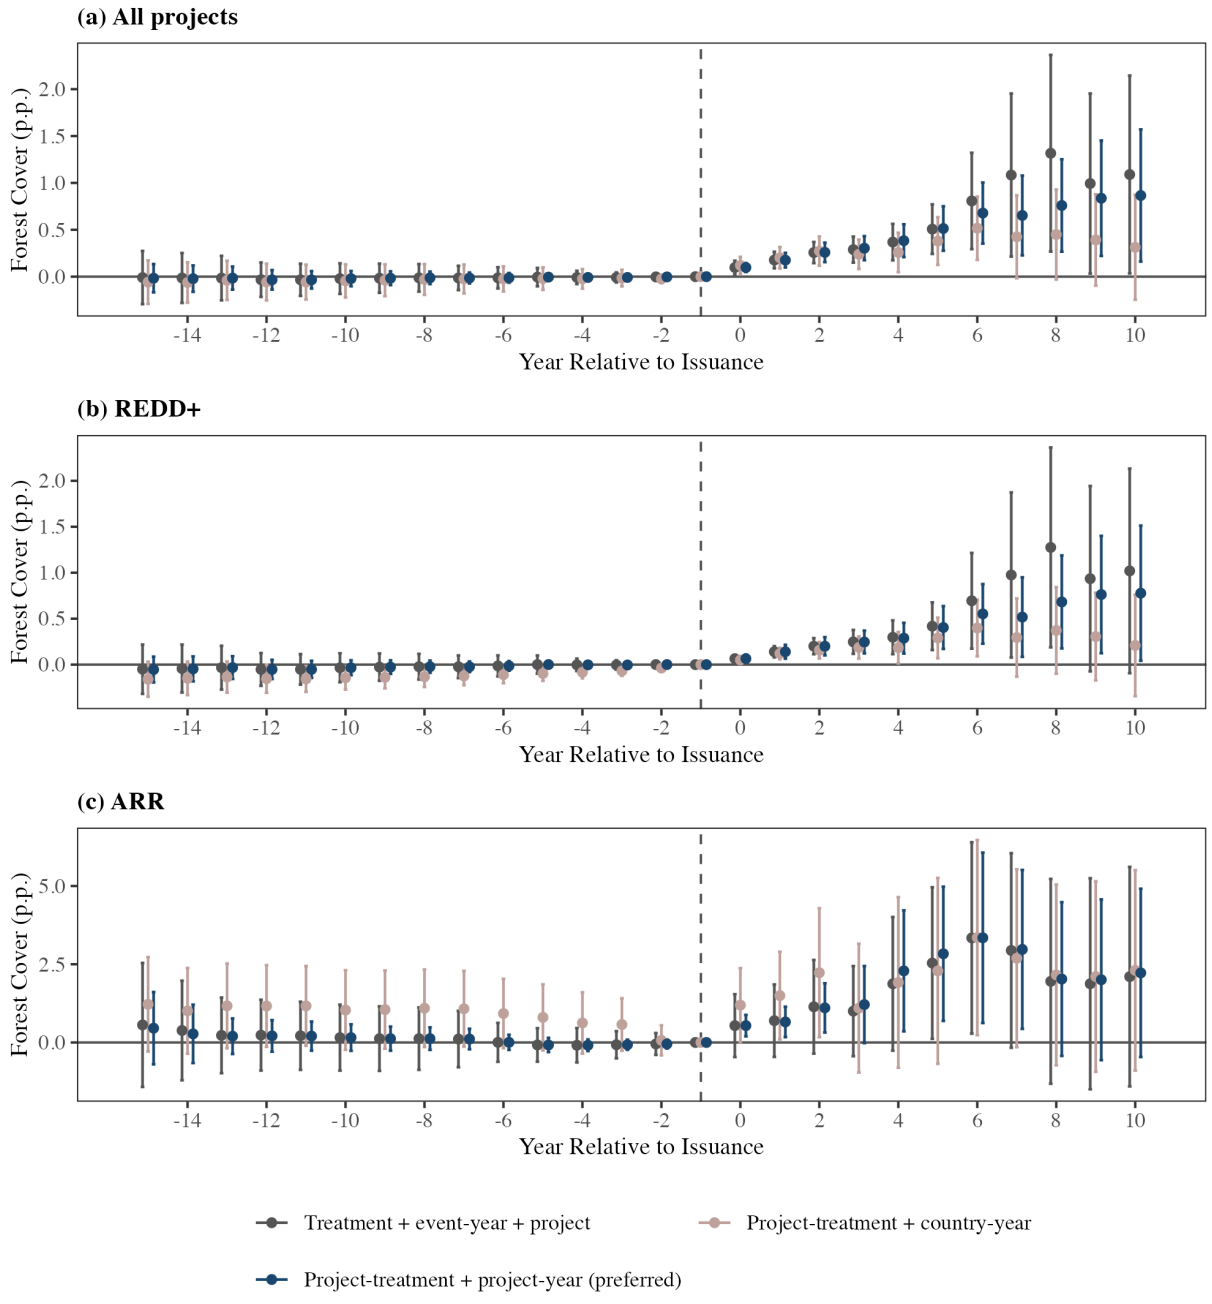


(d) ARR: CO_2 accounting



Notes: Subfigures (a) and (c) plot event-study estimates from equation (5) for REDD+ and ARR projects, respectively. Observations are weighted by the area of the corresponding treated grid cell. Standard errors are based on Conley (1999) using a uniform kernel and a 50 km radius cutoff. Subfigures (b) and (d) convert project-level treatment effects on forest cover from equation (6) into estimated CO_2 reductions (see Appendix D) and compare them with issued credits for REDD+ and ARR projects, respectively.

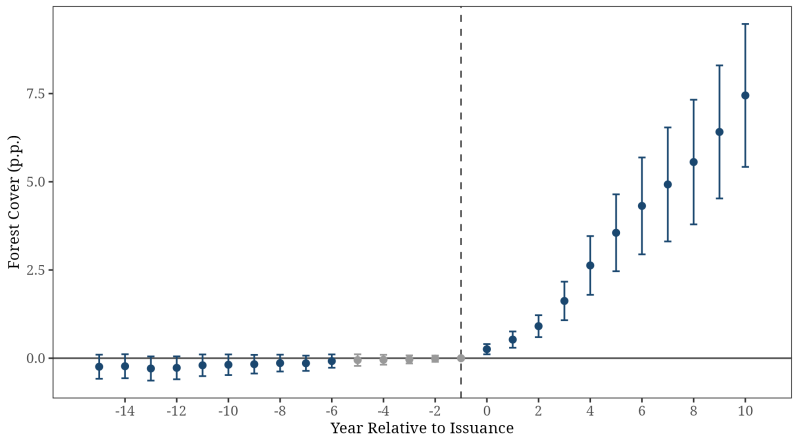
Figure A.6: Treatment Effects on Forest Cover, Alternate Fixed Effects



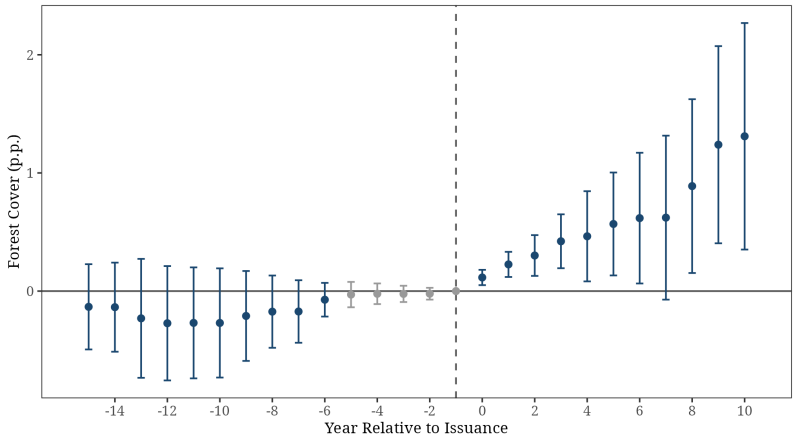
Notes: Figure plot estimates of the event study coefficients in equation (5) with alternative fixed effects. Each observation is weighted by the area of the corresponding treated grid cell. Standard errors based on Conley (1999) using a uniform kernel and 50 km radius cutoff.

Figure A.7: Treatment Effects on Forest Cover, Alternative Weighting

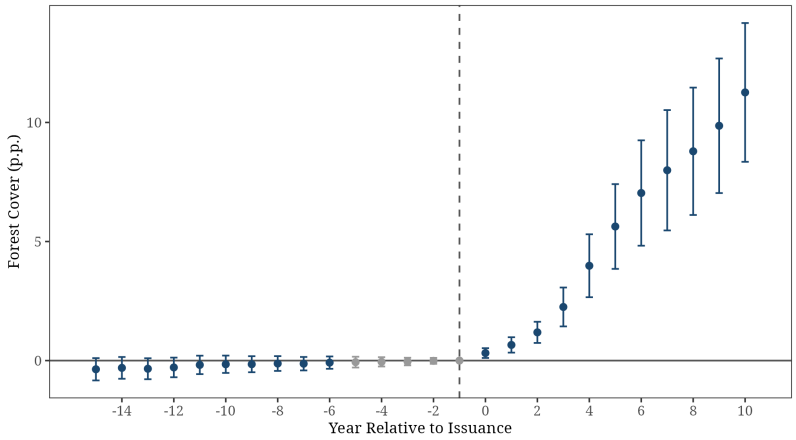
(a) All projects



(b) REDD+

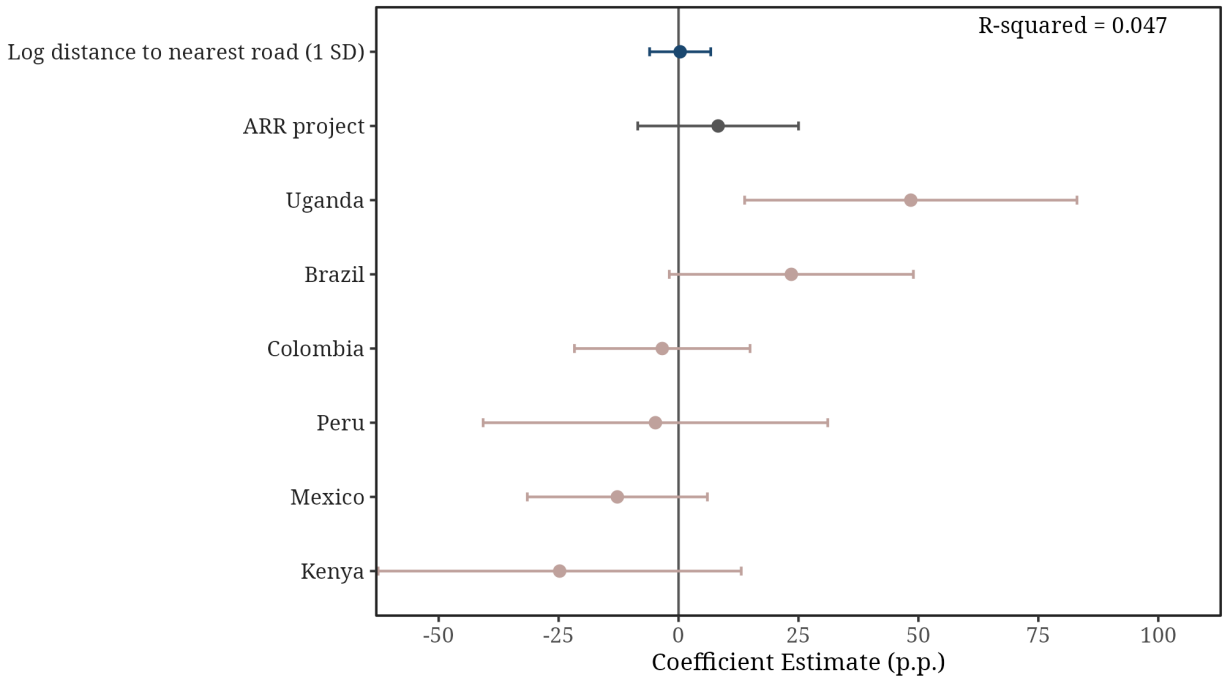


(c) ARR



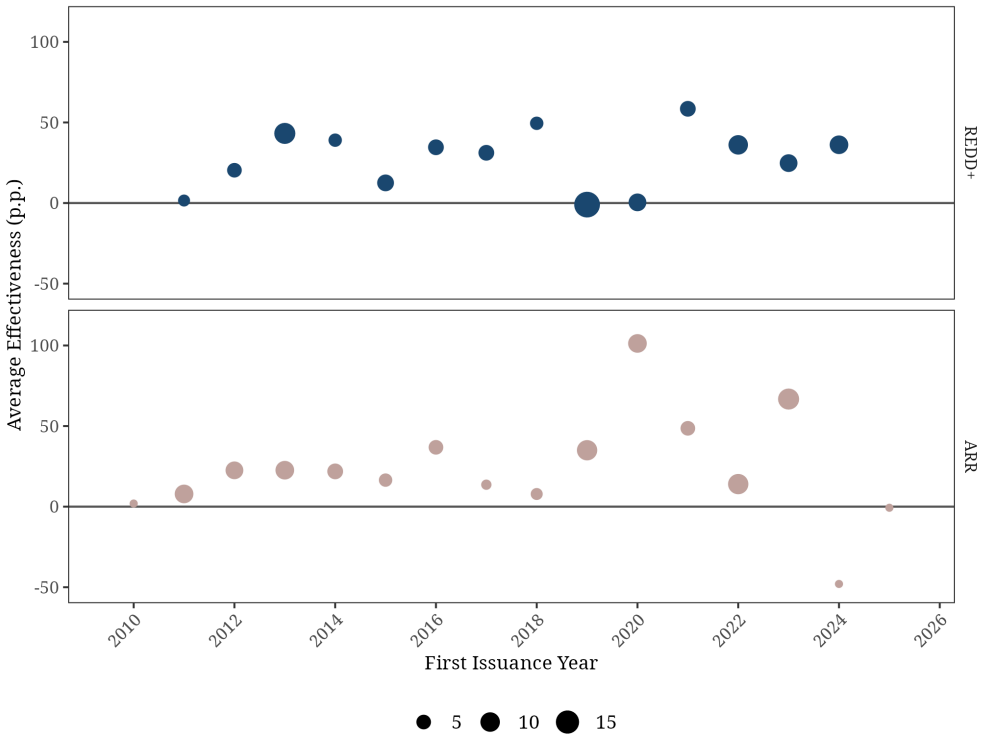
Notes: Subfigures (a), (b), and (c) plot event-study estimates from equation (5) for all projects, REDD+ projects, and ARR projects, respectively. Each project receives equal weight. Standard errors are based on Conley (1999) using a uniform kernel and a 50 km radius cutoff.

Figure A.8: Observable Heterogeneity in Effectiveness Estimates $\frac{\Delta_j}{q_j}$



Notes: The figure plots an unweighted multivariate regression of $\frac{\Delta_j}{q_j}$ on log distance to the nearest road (in standard deviations), an ARR sector indicator, and indicators for the six countries with the most projects in the sample. All remaining countries form the omitted Other countries category, and REDD+ is the omitted project type. A one-unit increase in a coefficient indicates that a project abated one percentage point more of its claimed CO_2 issuance.

Figure A.9: Estimated Effectiveness by Issuance Year



Notes: The top plot shows REDD+ effectiveness and the bottom plot shows ARR effectiveness. Each point is the project-weighted mean across projects with the same first issuance year. A 1 unit increase in the y-axis indicates a project abated 1 percentage point more of its claimed CO_2 issuance. Point area is proportional to the number of projects represented by the sector-year mean.

Table A.6: Evidence on selected outcomes and selection on crediting rules

	Baseline forest change			Deforestation		
	All	REDD+	ARR	Historic rate	DA percentile	Share issuing (%)
Panel A. Evidence: selected outcomes						
Documented (for issuance)	-0.98	-2.21	0.00			
Measured	-0.02	-0.20	0.12			
Panel B. Evidence: selection on rules						
Crediting based on historical baseline						
Yes				0.45		
No				0.32		
Crediting based on nearby reference regions						
REDD+ (common)					75.0	
ARR (never)					36.1	
Panel C. Scope for selection						
All projects who list in the registry						27.0

Notes: Panel A presents hand-collected information from project documents from a random sample of projects (20 REDD+ and 20 ARR). Documented (for issuance) reports the crediting-rule implied baseline forest cover change. Measured forest change is calculated for the matched controls from event year -1 through the final observed post-treatment event year. Baseline forest change is reported in percentage points per year. Panel B reports evidence on selection into crediting rules. The historic baseline row compares REDD+ projects whose crediting rules rely on historical baselines versus not, and reports mean annual forest-cover loss during the five years before first issuance, in percentage points per year. DA percentile refers to the percentile of a Deforestation Access (DA) measure of the project itself, relative to its matched controls. Deforestation Access (DA) is the inverse-distance-squared-weighted exposure to nearby forest-cover loss within 100 km and the same country-biome. Panel C reports the percent of Forestry and Land Use projects with any information recorded in a registry (had started the issuance process) that have any issued credits.

E.2 Empirical Bayes

There are two challenges to interpretation of the project-level effectiveness estimates in equation (6). First, project-specific estimates are subject to sampling error. Second, with many projects, even an empirical strategy that satisfies parallel trends on average may fail for a subset of projects. We develop a procedure that (1) adjusts project-specific event study estimates for potential pre-trend violations (Kwon and Roth (2024), building on Rambachan

and Roth (2023)), and then (2) shrinks project-level, pre-trend adjusted estimates towards the between-project average (Walters, 2024).

Let $\hat{\beta}_j^k$ indicate raw event study estimates from (6) and let $\breve{\beta}_j^k$ indicate the estimates after step 1 of 2. Let the $+$ superscript indicate post-periods $k \geq 0$, while $-$ indicates pre-periods $k < -1$: for the vector of coefficients $\hat{\beta}_j^-$ represents the subvector whose elements are $\hat{\beta}_j^k : k < -1$. The period -1 coefficient is normalized to 0. A similar convention applies to matrices. Define the variance-covariance matrix of the estimated coefficients as $\hat{\Sigma}_j$. Then, $\hat{\Sigma}_j^-$ is the pre-period submatrix, $\hat{\Sigma}_j^+$ the post-period, and $\hat{\Sigma}_j^\times$ the block off-diagonal covariance between pre- and post-period coefficients (dimension $|-| \times |+|$).

Pre-trend adjustment. We model violations of the parallel trends assumption, which imposes that the population pre-period coefficients are exactly 0, as a Gaussian random walk originating at the normalized period -1 and independently "walking" outward in the negative and positive directions. Define $\hat{\beta}_j^- \equiv \hat{u}_j^-$ for $\hat{\beta}_j^- = \{\hat{\beta}_j^k : k \leq -2\}$ from equation (6), which under no anticipation attributes all pre-period movement to the trend violation u_j. Under exact parallel trends, all pre-period coefficients would equal zero: $u_j^- = 0$. Normalizing $u_j^{-1} = 0$, pre-trend deviations are $d_j^k = u_j^k - u_j^{k-1} \sim \mathcal{N}(\mu_{d,j}, V_{d,j})$, i.i.d. across k; their sample analogues $\hat{d}_j^k$ additionally carry the sampling variance of $\hat{\beta}_j^-$.

We estimate $\mu_{d,j}$ and $V_{d,j}$ via maximum likelihood, yielding $\hat{\mu}_{d,j}, \hat{V}_{d,j}$. Because u_j^k accumulates $|k+1|$ deviations away from the reference period, these imply a prior on the level u_j with mean and variance

$$\hat{\mu}_{u,j}^k = (k+1)\hat{\mu}_{d,j}, \qquad \hat{V}_u[k,k'] = \begin{cases} \min\{|k+1|,\, |k'+1|\}\, \hat{V}_{d,j} & \text{if } \operatorname{sign}(k+1) = \operatorname{sign}(k'+1), \\ 0 & \text{otherwise,} \end{cases}$$

so that pre- and post-period violations are a priori uncorrelated, $V_{u,j}^\times = 0$. The posterior pre-period coefficient vector β_j^- has estimator:

$$\widehat{\mathbb{E}[\beta_j^- | \hat{\beta}_j^-]} = W_j^- \hat{\beta}_j^- + (I^- - W_j^-)\hat{\mu}_{u,j}^-,$$

where $W_j^- = ((\hat{\Sigma}_j^-)^{-1} + (\hat{V}_{u,j}^-)^{-1})^{-1}(\hat{\Sigma}_j^-)^{-1}$ and I^- is the identity matrix of the same dimension as the pre-period. W_j^- downweights pre-trends which are less precisely estimated. In practice, the pre-trend adjusted dynamic treatment effect $\breve{\beta}_j^{k^+}$ for $k^+ \in \{k : k \geq 0\}$ has estimator:

$$\breve{\beta}_j^+ = \hat{\beta}_j^+ - (\hat{\Sigma}_j^\times)^\top (\hat{\Sigma}_j^-)^{-1} \left(\hat{\beta}_j^- - \widehat{\mathbb{E}[\beta_j^- | \hat{\beta}_j^-]} \right) - \hat{\mu}_{u,j}^+.$$

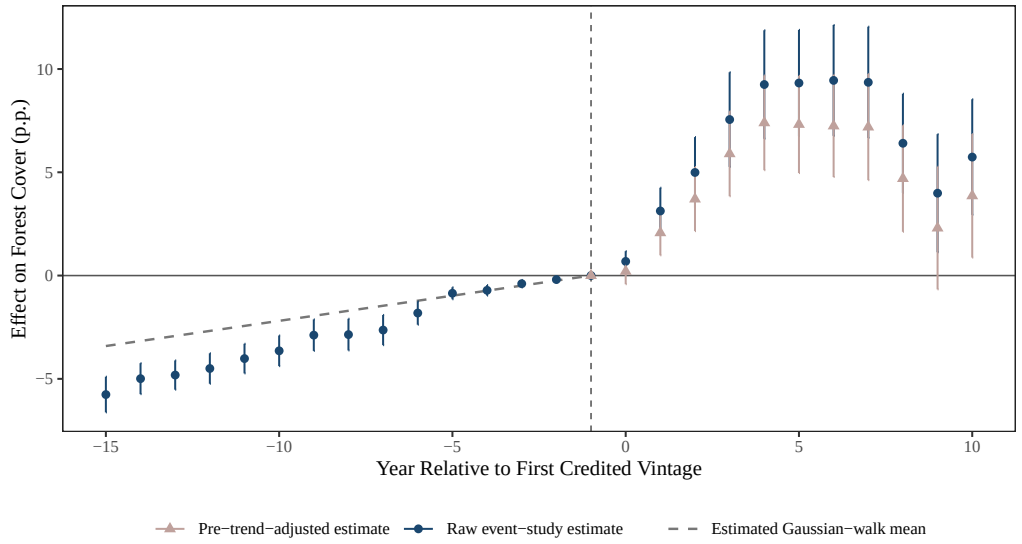
Loosely, the pre-trend adjusted coefficient can be thought of as the estimated coefficient less an extrapolated Bayesian line of best fit through pre-trend coefficients. The estimator for the posterior variance-covariance matrix of the pre-trend adjusted coefficients $\breve{\beta}_j^{k+}$ is

$$\breve{V}_j^+ = \hat{\Sigma}_j^+ + \hat{V}_{u,j}^+ - (\hat{\Sigma}_j^\times)^\top (\hat{\Sigma}_j^-)^{-1} \hat{\Sigma}_j^\times + [(\hat{\Sigma}_j^-)^{-1} \hat{\Sigma}_j^\times]^\top \left((\hat{\Sigma}_j^-)^{-1} + (\hat{V}_{u,j}^-)^{-1} \right)^{-1} [(\hat{\Sigma}_j^-)^{-1} \hat{\Sigma}_j^\times].$$

Proofs of these results can be found in the Online Appendix of Kwon and Roth (2024).

Figure A.10 illustrates the pre-trend adjustment for VCS1201, a REDD+ project with a smooth pre-period trend and a visible post-treatment adjustment. The dashed line plots the estimated mean path of the Gaussian random walk through the pre-period, anchored at zero in event year −1.

Figure A.10: Example of Pre-Trend Adjustment: VCS1201



Notes: Figure reports the raw project-level event-study estimates and the post-treatment estimates after the empirical Bayes pre-trend adjustment. The gray dashed line reports the estimated mean path of the Gaussian random walk through the pre-period, anchored at zero in the omitted event year −1. Error bars report 95% confidence intervals.

Shrinkage. Define the pre-trend adjusted effectiveness of project j is calculated by converting post-period coefficients to CO_2:

$$\breve{\Delta}_j/q_j = \frac{\Delta m_j^s}{q_j}\breve{\beta}_j^T,$$

where $k = T$ is the final event-time and Δm_j^s is a conversion factor for a project in sector s (REDD+ or ARR): see Appendix D for details. The corresponding variance of the adjusted effectiveness depends on the (T, T) element of $\breve{V}_j^+$, $\breve{v}_j^{TT}$, as $\breve{v}_j^{\Delta} = \left(\frac{\Delta m_j^s}{q_j}\right)^2 \breve{v}_j^{TT}$.

Our second stage is classic Empirical Bayesian shrinkage following Walters (2024). We calculate a prior for $\breve{\Delta}_j/q_j$ using the between-project mean and variance of pre-trend adjusted effectiveness estimates $\hat{\mu}_e, \hat{v}_e$. We then shrink potentially noisy project-specific estimates $\breve{\Delta}_j/q_j$ towards that prior:

$$\frac{\Delta_j}{q_j} \;\middle|\; \frac{\breve{\Delta}_j}{q_j}, \breve{v}_j^{\Delta} \sim \mathcal{N}\left(w_j \frac{\breve{\Delta}_j}{q_j} + (1 - w_j)\hat{\mu}_e, \frac{\hat{v}_e}{\breve{v}_j^{\Delta} + \hat{v}_e}\breve{v}_j^{\Delta}\right)$$

where, similar to step 1, $w_j = \frac{\hat{v}_e}{\breve{v}_j^{\Delta} + \hat{v}_e}$. Proofs of these results can be found in Walters (2024). This final set of posterior means are reported in Figure 3c.

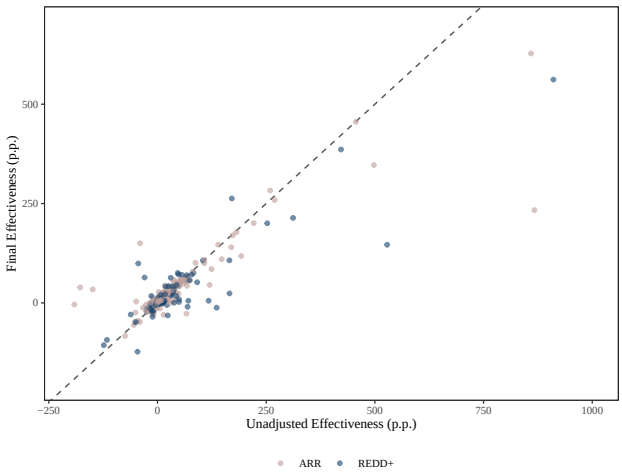
Appendix Table A.7 reports the empirical Bayes hyperparameters by sector and Appendix Figure A.11 compares the raw and final effectiveness estimates.

Table A.7: Hyperparameters for Empirical Bayes Procedure

	REDD+	ARR
Panel A. Pre-trend drift, area-weighted average		
$\hat{\mu}_{d,j}$	0.0016	-0.0197
	(0.0018)	(0.0199)
$\hat{V}_{d,j}$	0.0034	0.2402
	(0.0019)	(0.0668)
Panel B. Between-project shrinkage		
$\hat{\mu}_e$	0.4829	0.2184
$\hat{v}_e$	4.2496	0.6351

Notes: Panel A reports project-area-weighted averages of the mean and variance of the project-specific annual random-walk increments, d_j. Parentheses in Panel A report standard errors from 2,000 project-level bootstrap draws within sector. Panel B reports the mean of the pre-trend-adjusted effectiveness estimates and their cross-project variance net of average project-level estimation variance.

Figure A.11: Unadjusted and Final Effectiveness Estimates

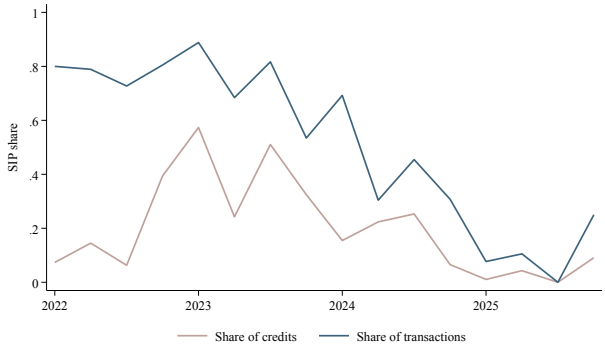


Notes: The figure compares unadjusted effectiveness with final effectiveness estimates following the Empirical Bayes procedure. The gray dashed line is the 45-degree line. For readability, the horizontal axis ends at 1,000 percentage points; one REDD+ project with an unadjusted estimate of 6,046 percentage points lies outside the displayed range.

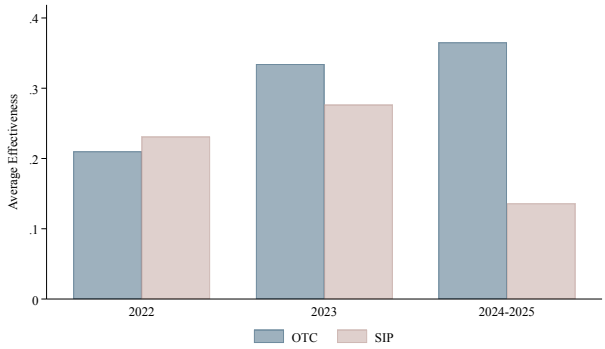
E.3 SIP Appendix

Figure A.12: Unraveling: Additional Evidence

(a) SIP Share of Transactions



(b) Composition: Publicly-Revealed Effectiveness



Notes: Sub-figure (a) plots the share of transactions and transacted credits that occurred via the SIP. Sub-figure (b) plots the composition of credits, measuring the average publicly revealed effectiveness among transacted credits (see Section 5.3 for more details). Data on credit composition in the SIP begins in 2022. The years 2024 and 2025 are pooled given low transacted credit volumes in the SIP in 2025.

E.4 Hedonics Appendix

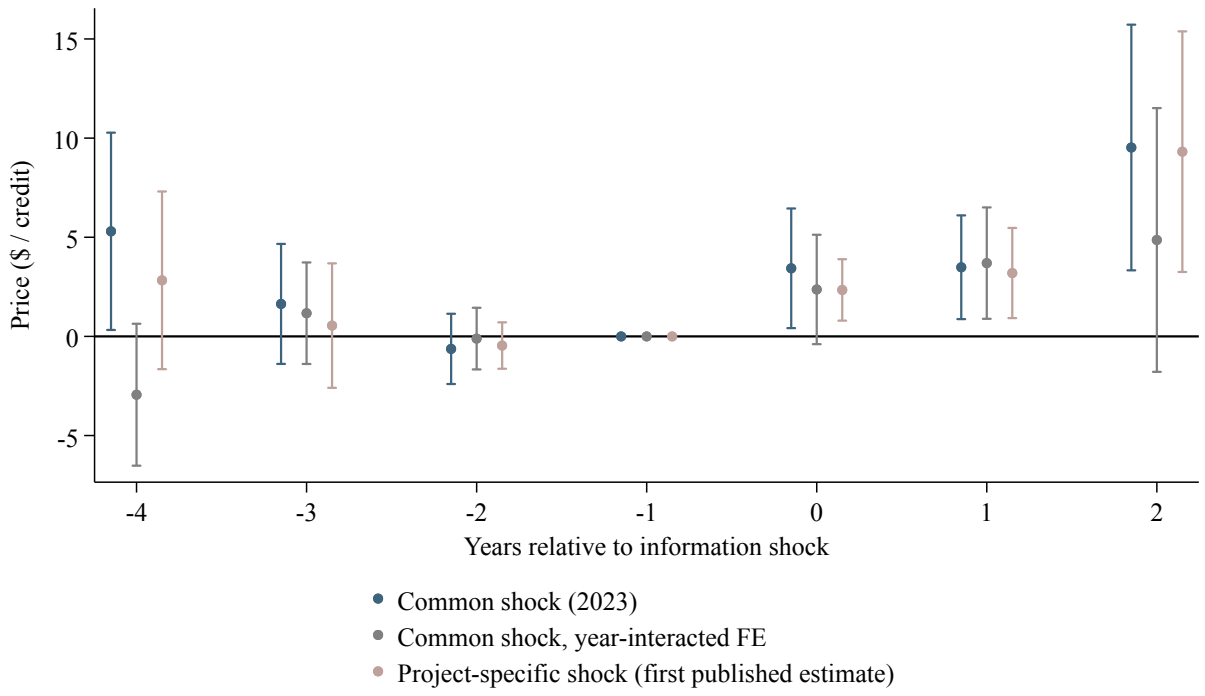
E.4.1 Figures and Tables

Table A.8: Robustness: Hedonic Estimates of Willingness to Pay for CO_2 Reductions

	Baseline FE		Year-interacted FE	
	Cont.	Binary	Cont.	Binary
	(1)	(2)	(3)	(4)
Panel A. Common shock (2023)				
$\hat{\phi}$	4.588	1.500	3.154	2.121
	(1.475)	(1.212)	(1.097)	(1.242)
Panel B. Project-specific shock (first published estimate)				
$\hat{\phi}$	5.203	2.189	3.810	3.194
	(1.115)	(1.008)	(1.212)	(1.211)

Notes: Table presents estimates of $\hat{\phi}$ from equation (7), pooling all published estimates $\tilde{\Delta}_j^{pub}$ (as in Table 3, columns (2)–(3)). Within each pair, the first column enters $\tilde{\Delta}_j^{pub}$ continuously; the second uses an indicator for above-median $\tilde{\Delta}_j^{pub}$, divided by the difference in mean $\tilde{\Delta}_j^{pub}$ between above- and below-median projects. Panel A defines $Post_t$ as 2023 onward for all projects; columns (1)–(2) reproduce Table 3, columns (2)–(3). Panel B defines the shock year separately for each project as the first year in which any published estimate of $\tilde{\Delta}_j^{pub}$ is available. Columns (1)–(2) include fixed effects for transaction year, country, vintage, and subsector. Columns (3)–(4) replace these with transaction-year fixed effects interacted with country, vintage, and subsector, where countries without an academic estimate are grouped into a single category. Sample is AFOLU credits; observations are weighted by credits transacted. Standard errors clustered by project in parentheses.

Figure A.13: Hedonics: Dynamic Effects



Notes: Figure plots event-study estimates of equation (7), pooling all published estimates $\tilde{\Delta}_j^{pub}$ and entering them continuously (as in Table 3, column (2)). Coefficients can be interpreted as the effect of moving from an ineffective (reduces zero tons of CO_2) to a perfectly effective (reduces one ton of CO_2) credit. The single post-shock interaction is replaced by interactions of $\tilde{\Delta}_j^{pub} \cdot Pub_j$ with indicators for each year relative to the information shock, omitting the year before the shock (-1). All specifications control for $\tilde{\Delta}_j^{pub}$ and for Pub_j interacted with transaction year. "Common shock (2023)" sets the shock in 2023 for all projects (event years -4 to $+2$ correspond to 2019–2025) and includes fixed effects for transaction year, country, vintage, and subsector. "Common shock, year-interacted FE" replaces these with transaction-year fixed effects interacted with country, vintage, and subsector, as in Table A.8, columns (3)–(4). "Project-specific shock" sets the shock year for each project to the first year in which any published estimate of $\tilde{\Delta}_j^{pub}$ is available, with baseline fixed effects. Event year -4 is estimated from very few transactions in projects with published estimates (0.2–0.4% of their credit volume), by credits transacted. Bars show 95% confidence intervals based on standard errors clustered by project.

E.4.2 Interpreting Estimates of $\hat{\phi}$

We characterize the interpretation of $\hat{\phi}$ in equation (7). It equals ϕ scaled by a factor that depends on how much buyers update their beliefs and on what they knew before the information shock.

Setup. Consider a project with a published estimate, $Pub_j = 1$. Let $\tilde{\Delta}_j \equiv \Delta_j / q_j$ denote its true, unobserved CO_2 reduction per credit and $\tilde{\Delta}_j^{pub} \equiv \Delta_j^{pub} / q_j$ its published value. Before the shock, under information set Pre, buyers' beliefs are $\mu_j \equiv \mathbb{E}[\tilde{\Delta}_j \mid Pre]$. After the shock, under $Post = Pre \cup \{\tilde{\Delta}_j^{pub}\}$, they are $\mu_j^\star \equiv \mathbb{E}[\tilde{\Delta}_j \mid Post]$. For simplicity, buyers share common beliefs and a common willingness to pay for CO_2, ϕ. Suppose buyers' prior is $\tilde{\Delta}_j \sim \mathcal{N}(\mu_j, \sigma_0^2)$ and they view the published estimate as a noisy signal, $\tilde{\Delta}_j^{pub} = \tilde{\Delta}_j + \eta_j$ with $\eta_j \sim \mathcal{N}(0, \sigma_\eta^2)$. Bayes' rule then gives

$$\mu_j^\star = (1 - \kappa)\, \mu_j + \kappa\, \tilde{\Delta}_j^{pub}, \qquad \kappa \equiv \frac{\sigma_0^2}{\sigma_0^2 + \sigma_\eta^2} \in [0, 1].$$

When $\kappa = 1$, buyers adopt the published estimates fully; when $\kappa = 0$, published estimates do not move beliefs.

Probability limit of $\hat{\phi}$. Under these assumptions,

$$\hat{\phi} \xrightarrow{p} \frac{\mathrm{Cov}\big[p_{j,\mathrm{post}} - p_{j,\mathrm{pre}},\, \tilde{\Delta}_j^{pub} \mid x_j\big]}{\mathbb{V}\big[\tilde{\Delta}_j^{pub} \mid x_j\big]} = \phi\, \kappa \left(1 - \frac{\mathrm{Cov}\big[\mu_j,\, \tilde{\Delta}_j^{pub} \mid x_j\big]}{\mathbb{V}\big[\tilde{\Delta}_j^{pub} \mid x_j\big]}\right), \tag{17}$$

where x_j denotes the fixed effects and lower-order interactions of Pub_j, $Post_t$, and $\tilde{\Delta}_j^{pub}$ in equation (7).

Proof. After partialling out x_j (Frisch–Waugh–Lovell), the triple-difference coefficient is the population slope from regressing the project-level change in price, $p_{j,\mathrm{post}} - p_{j,\mathrm{pre}}$, on $\tilde{\Delta}_j^{pub}$; this gives the first expression. Prices reflect buyers' expected CO_2 reductions, so $p_{j,\mathrm{post}} - p_{j,\mathrm{pre}} = \phi(\mu_j^\star - \mu_j)$. Substituting the updating rule, $\phi(\mu_j^\star - \mu_j) = \phi\kappa(\tilde{\Delta}_j^{pub} - \mu_j)$, and taking the covariance with $\tilde{\Delta}_j^{pub}$ yields the result. □

Discussion. Equation (17) shows two reasons $\hat{\phi}$ may differ from ϕ. First, buyers may not fully update to the published estimates ($\kappa < 1$), for example if they distrust them; this attenuates $\hat{\phi}$. Second, if buyers' prior beliefs were already correlated with the later-published estimates, part of the information was priced before the shock. A positive correlation attenuates $\hat{\phi}$.

The second term is testable. Before the shock, $p_{j,\mathrm{pre}} = \phi\, \mu_j$ up to x_j, so the coefficient from regressing pre-shock prices on $\tilde{\Delta}_j^{pub}$ equals $\phi\, \mathrm{Cov}[\mu_j, \tilde{\Delta}_j^{pub} \mid x_j] / \mathbb{V}[\tilde{\Delta}_j^{pub} \mid x_j]$. The pre-period coefficients in Figure A.13 are close to zero and Table 3 column (1) finds no relationship between prices and estimates of $\tilde{\Delta}_j$ in the sample without published information.

The first term, κ, is not separately identified from ϕ. Our baseline assumes the information shock was sufficiently informative relative to buyers' uninformed priors that κ is equal or close to one.

F Model Appendix

F.1 Land Use

Defining Potential Entrants. We construct a grid of potential entrants following the the steps below:

1. Create two grid meshes which cover the tropical band. Each divides the band into square bins. Bins have area equal to the mean observed project area of our in-sample REDD+ (first of two grids) or ARR (second of two grids) projects
2. Assign one-kilometer, non-ocean grid cells to these larger grids.
3. Cells intersecting any project up to 2014 are dropped, but cells intersecting projects founded after 2014 are eligible
4. Cells with missing GAEZ yield or travel time to port are dropped.

Finally, we take any potential entrant in the REDD+ grid with at least 80% forested area as a REDD+ potential entrant (based on the empirical forested area at baseline among REDD+ projects). All ARR potential entrants which do not intersect a REDD+ potential entrant are potential ARR entrants.

For each crop k we compute potential revenues for potential entrant j, π^k_{jg}:

$$p^k_c \times \text{potential yield}^k_{jg} - \kappa^k_{jg}$$

and aggregate to obtain the potential entrant level observables in equation (12), $\bar{\pi}^a_j - \kappa_j$:

$$\bar{\pi}^a_j - \kappa_j = \sum_{g \in j} \frac{A_{jg}}{\sum_{g \in j} A_{jg}} \max_k \{p^k_c \times \text{potential yield}^k_{jg} - \kappa^k_{jg}\}$$

That is, the agricultural return in potential entrant j is the area A_{jg}-weighted mean across cell-level, cross-crop, maximum returns.

Table A.9: Land Use Model: OLS versus IV Estimates

	REDD+		ARR	
	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)
Net agricultural revenue	−0.025	1.180	0.000	2.119
	(0.008)	(0.300)	(0.013)	(0.471)
Average elasticity	−0.06	2.58	0.00	4.65

Notes: Table compares OLS and IV estimates of the land use model (equation (13)). Net agricultural revenue ($\bar{\pi}_j^a - \kappa_j$, USD per ha) uses transport costs to the nearest port or city and, in columns (2) and (4), is instrumented with the straight-line distance to the nearest port or city. Coefficients are multiplied by 1,000. Columns (2) and (4) correspond to the estimates of $1/\sigma_v$ in Table 4. Average elasticity is the area-weighted average of the implied deforestation elasticity across potential entrants. Conley standard errors reported in parentheses.

F.2 Crediting and Entry

Our estimation procedure is as follows:

1. Draw standard normal shocks ε_{jn} and χ_{jn} for each potential entrant j and simulation draw $n = 1, \ldots, S$, held fixed across parameter guesses.

2. Given a guess of $(\alpha, \sigma_\varepsilon^s, \sigma_\chi^s)$, re-solve the land use model (equation (13)) at offset revenue $(p^s - \tau)\alpha$ per ton to obtain each potential entrant's true CO_2 reduction Δ_j and opportunity cost ω_j.

3. In each draw, compute issued credits $q_{jn} = \alpha(\Delta_j + \sigma_\varepsilon^s \varepsilon_{jn})$ and net revenue $\pi_{jn} = (p^s - \tau)q_{jn} - (\omega_j + \sigma_\chi^s \chi_{jn})$. The N^s potential entrants with the highest π_{jn} enter, where N^s is the observed number of projects; the mean entry cost μ_χ^s is the net revenue of the marginal entrant.

4. Compute moments among entrants, average across draws, and update the guess to minimize the sum of normalized deviations from the data moments.

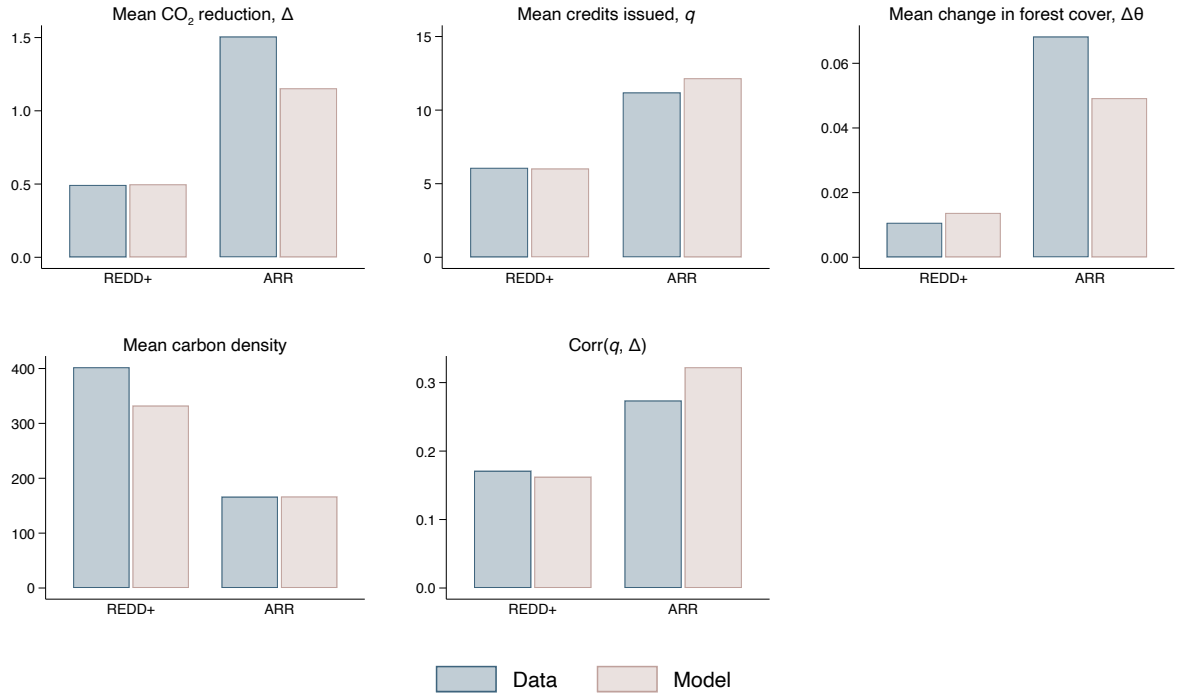
Standard errors are computed using the standard GMM sandwich formula.

Table A.10: Credit Issuance and Entry Estimates: Sector-Specific Shading Parameter

	Baseline		Sector-specific α		$\alpha = 1$	
	REDD+	ARR	REDD+	ARR	REDD+	ARR
	(1)	(2)	(3)	(4)	(5)	(6)
Shading parameter, α	1.01	1.01	0.91	1.12	1.00	1.00
	(0.12)	(0.12)	(0.16)	(0.09)		
Estimation error, σ_ε^s	4.01	4.86	3.64	5.57	3.88	4.73
	(1.23)	(0.58)	(3.35)	(0.51)	(0.99)	(1.16)
Entry cost mean, μ_χ^s	50.49	302.04	43.66	386.38	48.70	290.85
	(19.78)	(38.81)	(59.32)	(54.55)	(14.39)	(154.58)
Entry cost s.d., σ_χ^s	23.36	70.03	18.78	96.12	22.43	66.84
	(13.75)	(10.35)	(43.81)	(17.33)	(8.57)	(51.11)
α estimated	Yes	Yes	Yes	Yes	No	No
α equal across sectors	Yes	Yes	No	No	Yes	Yes

Notes: Table presents estimates of the credit issuance and entry model (based on equations (9), (10), and (14)) under alternative restrictions on the shading parameter α. Columns (1)–(2) reproduce Panel B of Table 4, in which α is estimated and imposed to be equal for REDD+ and ARR. Columns (3)–(4) estimate α separately for each sector. Columns (5)–(6) fix $\alpha = 1$, so that issued credits are unbiased population-wide estimates of CO_2 reductions. Standard errors in parentheses.

Figure A.14: Model Fit



Notes: Figure compares the moments targeted in estimation of the credit issuance and entry model (equations (9), (10), and (14)) in the data and in the model, separately for REDD+ and ARR. All moments are computed among projects that enter. Model moments are simulated at the baseline parameter estimates reported in Panel B of Table 4. Δ denotes CO_2 reductions, q denotes issued credits, and $\Delta\theta$ denotes the change in forest cover.

F.3 Demand

Under the T1EV assumption, equation (15) implies logit market shares. Inverting them (Berry, 1994) gives

$$\ln s_{jt} - \ln s_{0t} \;=\; \phi\, AE(\mathbf{z}_j) + \mathbf{z}_{jt}\psi + \gamma p_{jt} + \xi_{jt}, \tag{18}$$

where s_{jt} is the market share of credits from project j in year t and s_{0t} is the share of the outside option, given the potential market size M (Blaufelder et al., 2021).

We do not observe $AE(\mathbf{z}_j)$ directly. Following Section 5.3, we replace it with the triple-

difference interactions from equation (7):

$$\ln s_{jt} - \ln s_{0t} = \phi\, Post_t\, Pub_j\, \tilde{\Delta}_j^{pub} + \nu^1 Pub_j\, \tilde{\Delta}_j^{pub} + \nu^2 Pub_j\, Post_t + \nu^3 Pub_j + \gamma p_{jt} + \lambda_t + \eta_{c(j)} + \varsigma_{s(j)} + \vartheta_{v(j)} + \xi_{jt}, \quad (19)$$

where λ_t, $\eta_{c(j)}$, $\varsigma_{s(j)}$, and $\vartheta_{v(j)}$ are transaction-year, country, subsector, and vintage fixed effects, which together implement $\mathbf{z}_{jt}\psi$. As in Table 3, we enter $\tilde{\Delta}_j^{pub}$ either continuously or as a rescaled indicator for above-median $\tilde{\Delta}_j^{pub}$ (Table 4, columns (3) and (4)). Willingness to pay for a one-ton reduction in CO_2 is $\phi/|\gamma|$, and the own-price elasticity of demand is $\gamma p_{jt}(1 - s_{jt})$.

We estimate equation (19) by two-stage least squares, instrumenting p_{jt} with the distance from the crediting entity to the nearest city or port. These distances shift the cost of supplying credits but do not enter buyers' utility directly.