Yifan Wang
Department of Economics, London School of Economics
Houghton Street, London WC2A 2AE, United Kingdom
Email: y.wang390@lse.ac.uk | Phone: +44 7536 953648
Web: https://www.yifanwang.net

7 September 2026

Dear Professors Breza, Viviano, and Chandrasekhar,

I am writing to apply to the Workshop on Development Economics and Networks for Early-Career Researchers, to be held in Cambridge on 5 December 2026. I am a doctoral student in economics at the London School of Economics, where I am supervised by Oriana Bandiera, Robin Burgess, and John Van Reenen and where I work on development, environmental, and innovation.

I would like to present "The Global Reforestation Atlas" (joint with Linchuan Xu). Reforestation is now delivered through two very different mechanisms. On the one hand, governments plant directly, as in China's Grain for Green or AFR100. On the other hand, voluntary carbon markets pay private developers to reforest, and the market is already above two billion dollars a year. We still do not know whether either works, let alone which works better and where. We assemble a georeferenced database of reforestation projects from the primary carbon registries and from government sources, extract annual NDVI for every project polygon through Google Earth Engine, and estimate effects using synthetic difference-in-differences against a hundred randomly drawn control polygons of identical size in the same country. Preliminary estimates for sub-Saharan Africa put treated sites roughly five percentage points greener within five years, with wide variation behind that average: clear gains in the Gambia and Ghana, little or nothing in Gabon and Cameroon.

My complete contact details are given above. Robin Burgess (Professor of Economics, London School of Economics; r.burgess@lse.ac.uk) has agreed to provide a recommendation should I be shortlisted.

This single PDF contains, in order: this cover letter; my curriculum vitae, including a list of completed and current PhD coursework; and the paper I propose to present. Thank you for considering my application.

Yours sincerely,

Yifan Wang
PhD Candidate in Economics, London School of Economics

Yifan Wang

London School of Economics, Houghton Street, London WC2A 2AE, United Kingdom
+44 7536 953648 | y.wang390@lse.ac.uk | https://www.yifanwang.net

Education

PhD in Economics, London School of Economics 2024–2030
MSc in Economics, London School of Economics, with Distinction 2022–2024
BM in Management of Innovation and Entrepreneurship, Peking University 2020
BA in Biomedical English, Peking University 2020

PhD Coursework

Completed

EC441 — [Microeconomics], [2024/2025]
EC442 — [Macroeconomics], [2024/2025]
EC443 — [Econometrics], [2024/2025]
EC543 — [International Economics I], [2025/2026]
EC551 & EC552 — [Economics of Industry I & II], [2025/2026]
EC555 & EC556 — [Macroeconomics for Research Student I & II], [2025/2026]

Fields of Interest

Environmental Economics, Development Economics, Spatial Economics, Economics of Innovation

Working Papers

1. *Ray of Hope? China and the Rise of Solar Energy*
   with Ignacio Banares-Sanchez, Robin Burgess, Dávid László, Pol Simpson and John Van Reenen
   Revise and Resubmit, *The Quarterly Journal of Economics*
2. *Government and Nature: Evidence from the Distribution of Flood Damages in China*
   with Yuxiao Hu and Runhong Ma
   Reject and Resubmit, *The Review of Economics and Statistics*
3. *Weathering Poverty*
   with Clare Balboni, Oriana Bandiera, Robin Burgess, Anton Heil, Clément Mazet-Sonilhac and Munshi Sulaiman

Work in Progress

1. *Poverty in Space*
   with Oriana Bandiera and Robin Burgess
2. *Dams*
   with Guoju He and Yuxiao Hu
3. *Fires*
   with Agnes Norris Keiller
4. *Ghana Reforestation*
   with Maria Cedro
5. *Global Solar*
   with Ignacio Banares-Sanchez, Robin Burgess, Frederic Denker, Dávid László, Pol Simpson and John Van Reenen

6. *The Global Reforestation Atlas*
   with Linchuan Xu
7. *Rooftop Solar*
   with Robin Burgess, Bo Li and John Van Reenen
8. *South Africa Solar*
   with Qingyu Chen and Eddy Zou
9. *Transition Away from Fossil Fuels*
   with Robin Burgess, Michael Greenstone and John Van Reenen

Fellowships, Honors, and Awards

LSE Departmental Scholarship (LSE)	2024–2030
Excellence in Econometrics Award, EC221 (LSE)	2021
Excellence in Macroeconomics Award, EC210 (LSE)	2021
26th Challenge Cup: Award for Wusi Junior Scholar Competition (Peking University)	
Undergraduate Innovation Design (Peking University)	

Grants

The Weiss Fund (Travel), joint with Maria Cedro: $11,566	2026
Grantham Research Institute PhD Grant: £1,000	2026
STICERD PhD Grant, joint with Linchuan Xu: £5,000	2026
GSoS PhD Grant, joint with Linchuan Xu: £5,000	2026
IGC Summer Placement (Uganda): £5,000	2025

Employment

RA to Oriana Bandiera & Robin Burgess, LSE	2024–
RA (Predoc) to Oriana Bandiera, Robin Burgess & John Van Reenen, LSE	2022–2024
RA to Robin Burgess & John Van Reenen, LSE	2021–2022
RA to Ethan Ilzetzki, LSE	2021
RA to David Yang (Harvard), remotely	2021
RA at the Institute of New Structural Economics, Peking University	2019–2020
RA to Juanjuan Meng, Peking University	2019–2020

Policy Notes and Non-Refereed Publications

Ray of Hope? The Rise of Solar Energy in China
VoxDev, VoxEU, VoxChina — March 2026

Managing Flood Risk: Lessons from China's Flood Detention Basin Policy
VoxDev — August 2024

Presentations

2026 LSE EEE Work-in-Progress Seminar; NBER Chinese Economy Spring Meeting; SJTU Brownbag Seminar; Stanford SITE China in the Global Economy

2025 Columbia SIPA, New Thinking in Industrial Policy

2024 UCL Applied Micro Reading Group

Professional Activities

Refereeing: Econometrica; Journal of Development Economics; Management Science

Mentoring: PhD Applicant Mentoring Programme

Service:

LSE Environment Week, Organizing Committee	2025, 2026
LSE Environment Camp, Organizing Committee	2025, 2026
LSE EEE Seminar, Student Organizer	2024, 2025, 2026

Teaching Experience

EC1A3 (Microeconomics); EC1B3 (Macroeconomics), LSE	2025
Technology Commercialization of New Ventures, Peking University	2017–2018
Innovation Management for Entrepreneurship, Peking University	2017

Personal Information

Citizenship: China. Languages: Mandarin (native), English (fluent).

Statement of Motivation and Research Interests

Yifan Wang, London School of Economics

I am a doctoral student in economics at the London School of Economics, working in development and environmental economics with an interest in the economics of innovation. I am currently interest in using satellite and administrative data at scale, to evaluate policies that operate over wide areas and long horizons, where conventional survey evidence is thin or absent.

My current project, and the paper I propose to present, is "The Global Reforestation Atlas" (joint with Linchuan Xu). Reforestation is delivered both by governments planting directly, as in China's Grain for Green or AFR100, and by voluntary carbon markets paying private developers to plant trees. There is very little credible evidence on whether these reforestation projects raise forest cover and what's their impact on the local community. We assemble a georeferenced database of reforestation projects from the primary carbon registries and from government sources, extract annual vegetation series for every project polygon from satellite imagery, and estimate effects using synthetic difference-in-differences against control polygons drawn within the same country. Preliminary estimates for sub-Saharan Africa find treated sites are about five percentage points greener than their synthetic controls within five years, with wide variation behind that average: clear gains in the Gambia and Ghana, little or no effect in Gabon and Cameroon. We are now working on what explains that heterogeneity, and on what it implies for how restoration should be financed and where it should be sited.

A second strand of my work concerns clean energy and innovation. In "Ray of Hope? China and the Rise of Solar Energy", we ask how much of the global fall in solar costs is attributable to Chinese industrial policy, and what that implies for global welfare. Related projects extend this to the cost of capital for renewable projects across countries and to rooftop solar adoption.

A third strand concerns poverty and climate risk in low-income settings. This includes work on how flood and drought shocks affect household welfare and adaptation in rural Bangladesh, work on how individual poverty trap impacts the community level demand and wealth, and a field study I am designing on labour and contracting in Ghana's studying how reforestation affects the household.

The Global Reforestation Atlas

June 7, 2026

Abstract

Large-scale reforestation has become a central pillar of global climate policy, yet systematic causal evidence on program effectiveness remains scarce. We construct the first global-scale causal evaluation of reforestation programs by combining a newly released database of 45,628 georeferenced reforestation projects with annual satellite NDVI time series extracted via Google Earth Engine. Using a synthetic difference-in-differences design that matches each treated project polygon to 100 randomly drawn control polygons on pre-treatment vegetation trajectories, we find that treated areas become approximately 0.65 percent greener than controls on average. Effectiveness is higher in countries with stronger forest policy environments and lower where road and building density, surrounding cropland pressure, or nearby mining activity is high. Programs are also more likely to be located in wealthier, less agricultural areas, suggesting positive selection in program placement. A cost-benefit analysis using the social cost of carbon (\$185/t CO_2) implies that the average reforestation project generates 0.032 dollars of climate benefit per dollar spent, well below the break-even threshold. Finally, we develop a dynamic model of forest management featuring two externalities: a within-country common-pool externality and a cross-country carbon externality. We find that the decentralized steady-state forest stock falls 27 percent below the social optimum. A global carbon market partially closes this welfare gap by reallocating reforestation activity toward the Global South, where sequestration costs are lower.

JEL codes: Q23; Q28; O13; O44; C23

Keywords: Reforestation, Afforestation, Satellite Data, Voluntary Carbon Markets

1 Introduction

Forests cover roughly 30 percent of the Earth's land surface and store approximately 860 billion tonnes of carbon. Yet each year, deforestation releases around 4.8 billion tonnes of CO_2, accounting for nearly 10 percent of global greenhouse gas emissions. Against this backdrop, large-scale reforestation has become a central pillar of global climate policy. China's Grain for Green program has converted over 28 million hectares of cropland and degraded land into forest at a cumulative cost exceeding $100 billion. The African Forest Landscape Restoration Initiative (AFR100) has mobilised pledges to restore 100 million hectares across the continent. The UN Decade on Ecosystem Restoration (2021–2030) targets up to 1 billion hectares globally. Voluntary carbon markets, now exceeding $2 billion annually and projected to reach $50 billion by 2030, have directed billions more into reforestation as a mechanism for private firms to offset their emissions.

The scale of these commitments reflects the enormous potential value of successful reforestation. At the current US EPA social cost of carbon of $190 per tonne of CO_2, even modest vegetation gains across tens of thousands of project sites can translate into hundreds of billions of dollars in global climate benefit. Yet despite the resources at stake, remarkably little is known about whether these programs actually work. The existing empirical literature on forest conservation focuses almost exclusively on preventing deforestation. Rigorous causal evidence on reforestation effectiveness is limited to a handful of country-specific studies, and no prior work has attempted a globally representative, causally identified evaluation of whether planted forests deliver measurable vegetation recovery.

The problem is compounded by mounting evidence that the voluntary carbon market, which finances most of the reforestation projects we study, is not delivering on its promises. A growing empirical literature finds that the vast majority of carbon credits do not represent genuine emissions reductions: Probst et al. (2024) find that fewer than 16 percent of audited credits reflect real climate benefit, and West et al. (2020, 2023) show that forest conservation credits systematically overstate avoided deforestation. These findings raise a fundamental question: if the market mechanism is broken, are the underlying reforestation programs themselves delivering ecological value?

We construct the first *Global Reforestation Atlas*: a comprehensive, causally identified evaluation of reforestation programs spanning 45,628 georeferenced projects across the world. We document, country by country, where reforestation programs succeed, where they fail, and why. We address four interconnected questions. First, how do national forest policies shape investment in

reforestation? Second, how effective are programs at increasing forest cover? Third, what factors explain the large heterogeneity in program outcomes, and why do programs cluster in wealthier, less agricultural areas rather than where ecological need is greatest? Fourth, do reforestation programs generate sufficient climate benefit to justify their cost?

To answer these questions, we combine the newly released John et al. (2025) database of georeferenced restoration projects with annual satellite NDVI time series extracted via Google Earth Engine, LLM-coded national forest policy documents from the FAO FAOLEX database, satellite-derived cropland and mining data, and household survey data from the Demographic and Health Surveys. Our identification strategy exploits the staggered timing of project implementation. For each treated project polygon, we construct a donor pool by drawing 100 random control polygons of identical size within the same country. We then apply the synthetic difference-in-differences estimator of Arkhangelsky et al. (2021), which reweights these control units to match the pre-treatment vegetation trajectory of the treated site and estimates the causal effect of project launch on vegetation outcomes.

Our main empirical findings are as follows. First, on average treated sites increase NDVI by 0.0027, indicating they become 0.65 percent greener than their synthetic control group, though the most successful projects show a 36.54 percent increase. Second, effectiveness varies systematically with local economic conditions and governance: programs succeed in countries with stronger forest policy environments and fail where pre-existing human activity is high, agricultural pressure is greater, or mining activity is nearby. Third, programs are not placed where they are most needed: reforestation projects are disproportionately sited in wealthier, better-electrified areas with lower agricultural activity, consistent with voluntary carbon market incentives directing investment toward easier sites rather than areas of highest ecological need. The result is a classic adverse selection problem: the projects most likely to be registered are precisely those least in need of subsidy (Jayachandran et al., 2017; Ferraro and Pattanayak, 2006).

A back-of-envelope welfare calculation translates our vegetation estimates into climate benefits. Using a literature-based conversion from NDVI gains to carbon sequestration (Spawn et al., 2020) and the social cost of carbon ($\$185/tCO_2$), we estimate that the average reforestation project generates only \$0.032 of climate benefit per dollar spent. This benefit-cost ratio of 3 cents on the dollar is far below the break-even threshold of one. The number, however, masks substantial heterogeneity: successful programs deliver a BCR of \$1.80 per dollar spent, while ineffective programs deliver essentially zero. This suggests that the poor average return is not an inherent feature of reforestation as an intervention, but rather reflects the misallocation of investment to-

ward locations and institutional environments where programs are unlikely to succeed.

Finally, we develop a dynamic model of forest management to interpret these findings and study the scope for carbon market interventions. Within each country, two groups differ in their forest-amenity valuation ($\phi_H > \phi_L$) and productivity, and face competing incentives to deforest for income and reforest for amenity. The model generates an endogenous forest Kuznets curve: forest first falls as capital-scarce economies clear land to finance accumulation, then recovers as capital deepens and depleted forest becomes scarce. Laissez-faire fails through two distinct channels: a within-country common-pool externality (each group ignores the forest amenity it provides to the other) and a cross-country carbon externality (each country ignores its contribution to the global atmospheric carbon stock). A competitive carbon market corrects the carbon channel but leaves the amenity distortion uncorrected, because a single carbon price cannot simultaneously implement the planner's optimal deforestation and reforestation conditions. Quantitatively, the decentralized steady-state forest stock falls 27 percent below the social optimum; the carbon market partially closes this gap, with reforestation investment shifting toward the lower-cost South.

Related literature. This paper contributes to three strands of research.

Forests and development. A large empirical literature studies the political economy of deforestation and the effectiveness of conservation policy. Burgess et al. (2012) show that political incentives shape deforestation enforcement in Indonesia. Jayachandran (2013) examine how credit constraints drive deforestation among smallholders. Balboni et al. (2023) use satellite vegetation data to study long-run land clearing dynamics. Goldstein and Udry (2008) provide foundational evidence on how land tenure shapes investment incentives, with direct implications for maintenance of planted forests. On the structural side, Assunção et al. (2024) develop a spatial dynamic model of optimal land allocation between cattle ranching and reforestation in the Brazilian Amazon, showing that modest carbon prices can shift the Amazon from net emitter to net absorber over 30 years; Hansen and Scheinkman (2025) extend this framework to study biodiversity co-benefits of emission pricing. Most closely related to our welfare analysis, Jayachandran et al. (2017) use an RCT to evaluate a PES program in Uganda, finding a cost of approximately \$46/tCO_2 avoided. We complement this literature by shifting focus from deforestation prevention to active restoration at global scale, and by contributing the first global benefit-cost estimates for reforestation programs.

Voluntary carbon markets. The voluntary carbon market is projected to reach \$50 billion by 2030, yet its effectiveness is increasingly questioned. Calel et al. (2025) find that CDM offsets broadly fail to offset carbon. Probst et al. (2024) find that only around 16 percent of issued credits rep-

resent genuine emissions reductions. West et al. (2020, 2023) show that REDD+ forest projects systematically overstate carbon reductions. Rodemeier (2026) find that consumer willingness to pay for carbon mitigation is close to zero absent explicit disclosure. The common thread is information asymmetry: buyers cannot verify whether credits represent real sequestration. We address this directly with independent satellite-verified estimates of vegetation outcomes for reforestation projects, providing the first large-scale empirical test of whether voluntary carbon market reforestation credits deliver measurable ecological impact.

Causal inference using satellite data. Henderson et al. (2012) establish nighttime lights as a proxy for economic activity. Donaldson and Storeygard (2016) survey satellite data applications in economics. Burgess et al. (2012) and Balboni et al. (2023) use Landsat and NDVI data to study deforestation and vegetation dynamics. We contribute by constructing a global satellite panel for 45,628 project polygons and linking annual NDVI outcomes to both a causal identification strategy and welfare calculations.

The paper is organized as follows. Section 2 describes the data. Section 3 presents motivating facts. Section 4 presents the research design and empirical results. Section 5 reports the welfare calculation. Section 6 develops the theoretical model. Section 7 concludes.

2 Data

2.1 Reforestation Projects

Our primary data source is John et al. (2025), who compile a global database from 45,628 afforestation and reforestation projects[1] spanning 33 years. The database consolidates information from over 50 sources, including voluntary carbon market registries (Verra, Gold Standard, Climate Action Reserve), NGO project trackers, and corporate sustainability platforms[2]. For each planting site, the database records the geographic polygon, start year, site area, species planted (where reported), and the hosting platform. For each planting site within a project, the database provides a site-level polygon and, where available, satellite-augmented time-series information[3].

[1] In total, there are 1,289,068 planting sites.

[2] However, government-led projects are under-represented because they rarely publish machine-readable data publicly.

[3] John et al. (2025) also designed an Location Data Integrity Score (LDIS), which evaluates ten spatial indicators of geographic boundary reliability, including the presence of other land cover within the declared site, nesting and intersection with other polygons, and atypical geometric shapes such as perfect circles. Our results are robust to restricting the sample to high-integrity sites (LDIS ≤ 2).

Our main estimation sample covers Sub-Saharan Africa[4], where we observe 1,964 project sites across 25 countries. Africa is the natural starting point: it accounts for the largest share of projects in the database, the voluntary carbon market is particularly active there, and the national policy and household data needed for our heterogeneity analysis are well-populated by the FAOLEX and DHS databases. We plan to extend the analysis to all regions in subsequent work.

2.2 Satellite Vegetation Data

For each project polygon, we extract an annual time series of the Normalised Difference Vegetation Index (NDVI) using the MODIS Terra Surface Reflectance product (MOD13A3) via Google Earth Engine. NDVI, computed as the ratio of near-infrared to red reflectance, is a standard proxy for vegetation density and photosynthetic activity[5]. The MODIS product provides monthly composites at 500-metre resolution, available from 2000 to the present; we compute annual averages from the monthly values. For each treated project polygon, we draw 100 random control polygons of identical size within the same country. Control polygons are drawn at the beginning of the analysis period and held fixed. Section 4 describes how we use these control polygons to construct synthetic counterfactuals.[6]

2.3 Forest Policy Data

To measure the national policy environment, we use the FAOLEX database[7], one of the world's largest electronic collections of national laws, regulations and policies on food, agriculture and natural resources management, covering 220,261 country-year observations. FAOLEX contains legal and policy documents drawn from more than 200 countries and territories, originating in over 40 languages, and is updated with an average of 8,000 new entries per year. We restrict

[4]This is due to computational restrictions.

[5]We supplement NDVI with two additional satellite-derived variables. First, we extract annual tree cover gain and loss from the Hansen Global Forest Change dataset (Hansen et al., 2013), a 30-metre annual product derived from Landsat imagery that provides a more direct measure of canopy cover change. Second, we use CHIRPS daily precipitation at 5km resolution as a climate control, since inter-annual rainfall variation is a major driver of NDVI changes that could confound our estimates if not accounted for.

[6]As a robustness check, we re-draw control polygons excluding existing forests (Hansen tree cover $>$ 30% in 2000), protected areas, and other project sites. This restriction ensures control polygons are placed on land that is plausibly suitable for reforestation. We do not impose it in the main specification because doing so risks introducing measurement error, reducing the donor pool, and creating unintended selection bias by systematically excluding certain land types. That said, results are quantitatively similar.

[7]© FAO. FAOLEX Database. License: CC BY-NC-SA 3.0 IGO. Extracted from: www.fao.org/faolex/opendata. Date of Access: [26/05/2026]

attention to forestry-relevant documents and extract all country-year records in which a new forest policy, law, regulation, or strategy was enacted.

To extract structured information from the raw policy texts, we use a large language model (LLM) pipeline. For each document, the pipeline extracts the enactment year, the policy instrument type (law, regulation, strategy, or decree), the responsible ministry or agency, and the stated enforcement mechanism. All LLM-coded outputs are validated against a manually coded reference sample of 200 documents. From these variables, we construct a country-year count of active forest policies, and also a variable indicating the strength of the policy, ranging from 1-5 (very high; middle high; middle; middle low, low).

2.4 Household Survey Data

We use the Demographic and Health Surveys (DHS), which provide nationally representative household-level data on education, employment, wealth, and asset ownership for all Sub-Saharan African countries in our sample, with survey waves typically spaced four to five years apart. All DHS surveys geocode households at the cluster level. To protect respondent confidentiality, the DHS programme displaces cluster coordinates by up to 2km in urban areas and 5km in rural areas (with 1 percent of clusters displaced by up to 10km). We account for this displacement in our spatial matching by computing treatment exposure using buffer radii of 10km, 25km, and 50km around each cluster centroid, and showing that our results are robust to this choice.

We spatially match each DHS cluster to the nearest reforestation project in our sample. The primary variables of interest for our heterogeneity analysis are measures of agricultural land pressure: the share of households owning agricultural land, livestock counts (cattle and goats), and the share of working-age adults employed in agriculture, all computed within a 50km buffer around each project polygon. Wealthier and more urbanised areas are proxied by the DHS wealth index and access to electricity.

2.5 Additional Spatial Data

Cropland. We use a global annual cropland map at 30-metre resolution (Potapov et al., 2022), which provides the share of cropland within any given area for each year from 2000 to 2019. For each project polygon, we compute the cropland share within a 50km buffer in the years preceding project launch. This variable captures the intensity of agricultural land use near the project and serves as a measure of local land pressure.

Mining. We use the global database of mining areas compiled by Maus et al. (2020), which provides georeferenced polygons for large-scale mining operations worldwide. For each project polygon, we compute the log distance to the nearest mine and an indicator for whether any mine falls within a 50km buffer. Mining activity is a potential threat to forest survival, both through direct physical encroachment and through the economic incentives it creates for land conversion.

Building footprint and road data. For Ethiopia, we obtain building footprint data from Google Open Buildings V3, retaining structures with a detection confidence score of at least 0.7. We spatially join each footprint to the treated project polygons and compute, for each polygon, the total number of buildings and total built-up area. Road network data are extracted from OpenStreetMap. We spatially join them to project polygons, and compute total road length. Both variables serve as proxies for local human activities inside reforestation project sites.

3 Motivating Facts

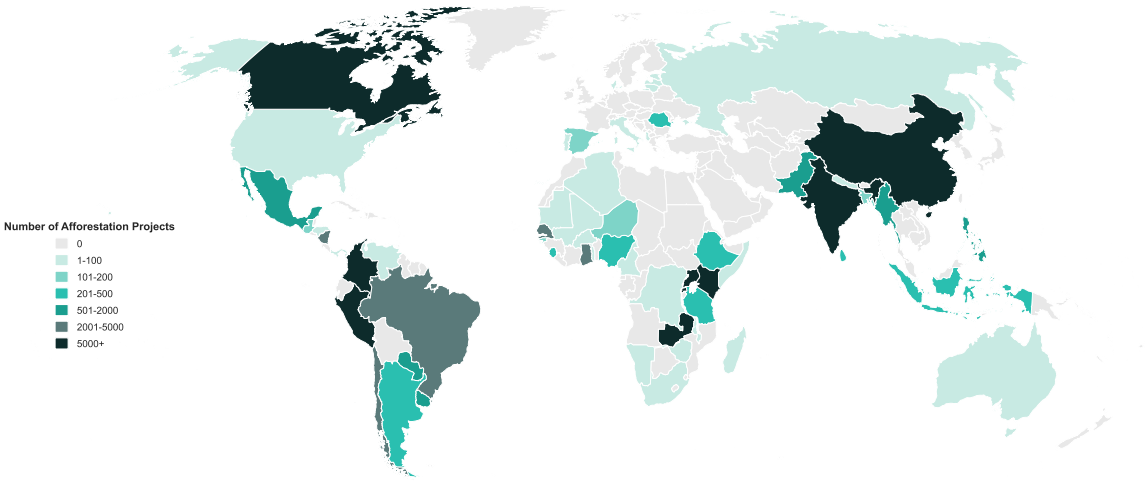
3.1 The Global Reforestation and Policy Atlas

Figure 1 presents two global maps that together motivate the paper. The top panel maps the cumulative number of afforestation and reforestation projects by country as of 2024, drawn from the John et al. (2025) database. The bottom panel maps the cumulative number of forest policy instruments enacted by country, drawn from the FAOLEX database.

Several patterns stand out. First, reforestation investment is geographically concentrated. Sub-Saharan Africa, South Asia, and parts of Latin America account for the bulk of recorded projects, reflecting the combination of high ecological demand (degraded land, deforestation pressure) and the active presence of voluntary carbon market registries in these regions. A small number of countries dominate: Nigeria, India, Ethiopia, and Brazil each have more than 2,000 recorded project sites, while the majority of countries globally have fewer than 100. Large parts of the world, including Central Asia, the Middle East, and much of Eastern Europe, have zero or near-zero recorded reforestation activity.

Second, the policy landscape follows a broadly similar geographic pattern. Forest policy activity is densest in the same regions that attract the most reforestation investment: Sub-Saharan Africa, South Asia, and Latin America. Western Europe and North America have extensive policy frameworks but relatively few registered projects, reflecting the dominance of voluntary carbon

(a) Global Reforestation Project Atlas



(b) Global Forestry Policy Atlas

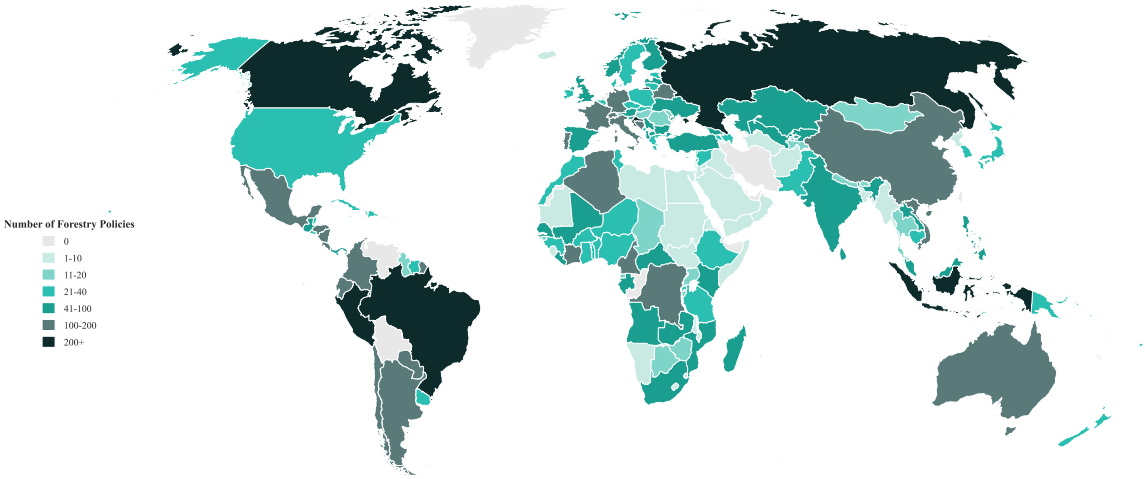


Figure 1. Global Reforestation Projects and Forest Policy

Notes: Panel (a) maps the cumulative number of afforestation and reforestation projects by country as of 2024, drawn from John et al. (2025). Panel (b) maps the cumulative number of forest policy instruments enacted by country, drawn from the FAO FAOLEX database.

markets in directing investment toward developing countries. Within the developing world, there is substantial variation: some countries with many reforestation projects have enacted relatively few forest policies, while others have built up strong governance frameworks without correspondingly high project counts.

Third, on balance, countries with more reforestation projects tend to have stronger policy environments. This positive correlation is consistent with voluntary carbon market developers

favouring countries where regulatory frameworks reduce uncertainty about land tenure and carbon credit validity. This pattern motivates our investigation of whether policy stringency not only predicts where projects are placed but also whether they succeed.

3.2 Reforestation and Structural Transformation

Figure 2 documents the relationship between a country's stage of structural transformation and its net change in forest cover.[8] The pattern is striking: higher agricultural GDP share is strongly negatively correlated with forest area growth, while higher services share is positively correlated. Countries where agriculture dominates the economy tend to lose forest, while countries that have shifted toward services tend to recover it. The relationship is driven by non-OECD countries (Appendix Figure B2), where agriculture shares are still high and structural transformation is ongoing; among OECD countries, which have already completed the transition, the relationship is substantially weaker (Appendix Figure B1). This is consistent with the empirical evidence in Foster and Rosenzweig (2003), who show that the decline in agricultural employment associated with development is the primary driver of forest recovery.

Taken together, these two motivating facts frame our empirical strategy. The strong positive correlation between policy environments and project counts suggests that reforestation investment flows toward countries with stronger governance, consistent with voluntary carbon market developers seeking regulatory certainty. The strong correlation between structural transformation and forest recovery points toward local agricultural pressure as a key determinant of whether planted forests survive. The heterogeneity results in Section 4.1 directly test whether these two forces, governance quality and agricultural pressure, explain the large variation in program effectiveness we document.

4 Method and Results

4.1 Research Design

Our unit of analysis is the project polygon. Each treated unit is a georeferenced planting site, characterised by its geographic boundary, start year, and country. The outcome variable is the annual average NDVI within the polygon, extracted from MODIS imagery via Google Earth Engine.

[8] GDP data from the World Bank World Development Indicators; country-level forest area data from the FAO Global Forest Resources Assessment.

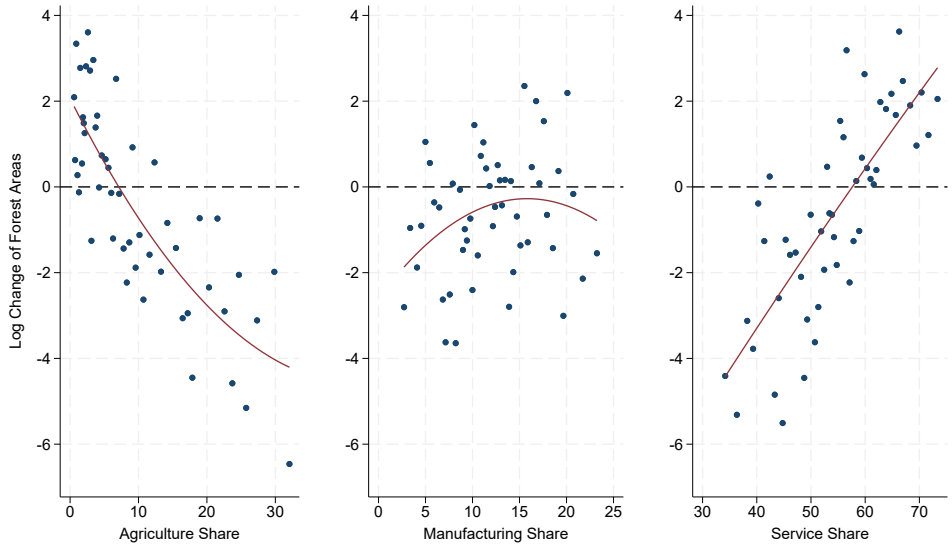


Figure 2. Forest Change and Sectoral GDP Share

Notes: Each panel shows a binscatter plot with a quadratic fit. The y-axis plots the change in forest area (asinh-transformed) between FRA benchmark years (1990, 2000, 2010, 2015, 2020, 2024). The x-axis plots the sectoral share of GDP (%) from the World Bank World Development Indicators. The three panels show agriculture, manufacturing, and services GDP shares respectively. Forest area data are from the FAO Global Forest Resources Assessment 2025.

The central challenge for causal identification is that reforestation projects are not randomly assigned. Projects are placed in locations chosen by implementing organisations for ecological, economic, and strategic reasons. Areas with naturally recovering vegetation may attract projects precisely because they offer cheap carbon credits; areas with high deforestation pressure may attract projects because the additionality is higher. Either form of selection would confound a naive before-after or treated-vs-untreated comparison.

Our identification strategy exploits the staggered timing of project launches across countries. For each treated project polygon, we draw 100 randomly placed control polygons of identical size within the same country. Restricting to the same country ensures that control units share the same macro-level policy environment, climate zone, and market conditions as the treated site, while random placement ensures they were not systematically selected for intervention. We extract the same NDVI time series for each control polygon, creating a balanced panel from 2000 to the present. We then apply the synthetic difference-in-differences estimator of Arkhangelsky et al. (2021), which reweights these control units to match the pre-treatment NDVI trajectory of each treated site before applying a DiD estimator. By construction, treated and control polygons follow parallel vegetation trends prior to project launch. Under the assumption that this parallel trend would have continued in the absence of treatment, any post-launch divergence in NDVI between

the two groups can be causally attributed to the reforestation program.

To document dynamic treatment effects and test for pre-trends, we complement this with a country-by-country event study that regresses NDVI on year-relative-to-treatment indicators. The formal estimating equations, identifying assumptions, and estimation details are provided in Appendix B.1.[9]

4.2 Policy Environment and Reforestation Investment

We begin by documenting the relationship between the national forest policy environment and the prevalence of reforestation investment. Figure 3 presents four complementary views.[10] Panel (a) plots the annual count of new reforestation programs registered globally over 2000–2024. New registrations grew steadily through the early 2010s, accelerated sharply around 2015–2021, and then declined somewhat toward the end of the period. The decline likely reflects a reporting lag rather than a genuine slowdown, since recently launched programs take time to be publicised and recorded in the database.

Panel (b) plots two time series for forest policy: the grey bars show the annual count of new forest policy instruments enacted globally, and the red line shows average policy strength, constructed using an LLM pipeline that codes each policy document for the strength of its enforcement provisions and binding commitments. Both series show a notable decline after approximately 2015: the count of new policies fell from around 300–400 per year to below 200 by the mid-2020s, and average policy strength has trended downward over the same period. This suggests that not only are fewer forest policies being enacted in recent years, but the ambition and enforceability of new policies has also weakened.[11]

[9] We assess robustness in four ways. First, we re-estimate the main specification using a standard two-way fixed effects (TWFE) estimator and the heterogeneity-robust estimators of Callaway and Sant'Anna (2021), Sun and Abraham (2021), and De Chaisemartin and d'Haultfoeuille (2020), which address potential bias from staggered treatment timing and treatment effect heterogeneity across cohorts. Second, we assess robustness to alternative matching methods, including entropy balancing (Hainmueller, 2012; Basri et al., 2021), propensity score matching, and Mahalanobis distance matching (Abadie and Imbens, 2006; Imbens, 2015). Third, we re-run the analysis on the full sample and a restricted sample limited to high-integrity sites to assess whether results are driven by projects with unreliable geographic boundaries. Fourth, we restrict our control polygons to exclude existing forests, protected areas, and other project sites.

[10] Since reforestation data are available from 2000, we focus on the post-2000 period.

[11] Our current policy strength measure is based on LLM scoring of document-level enforcement and commitment language. In future work, we plan to extend this pipeline to extract more granular policy characteristics, including the specific instruments used (carbon taxes, protected area designations, community forestry rights, logging bans), the responsible ministry, and whether the policy is linked to international frameworks such as the UNFCCC NDCs or the CBD. This richer coding will allow us to test which specific policy features predict reforestation investment and effectiveness, going beyond aggregate counts and strength scores.

Panels (c) and (d) plot cross-country relationships. Panel (c) shows that countries with more cumulative forest policies also tend to have more registered reforestation projects. Countries such as Spain, China, Brazil, India, and Peru combine extensive policy frameworks with large project portfolios. Russia is a notable outlier: despite having one of the largest bodies of forest legislation globally, it has very few registered reforestation projects, likely reflecting the reluctance of foreign developers and voluntary carbon market investors to operate there given geopolitical concerns. Panel (d) shows that the same positive relationship holds when policy count is replaced by policy strength: countries where forest legislation is more stringent and enforceable also attract more reforestation investment. Kenya, Algeria, and Morocco stand out as countries with both high policy strength and high project counts; Libya, Namibia, and Mauritania combine low policy strength with few projects.

Together, these patterns suggest that stronger policy reduces uncertainty about land tenure and carbon credit validity, making countries with robust forest governance more attractive to international voluntary carbon market investors seeking credible offsetting opportunities. This motivates our use of both policy count and policy strength as predictors of project effectiveness in Section 4.4.

We interpret this cross-sectional correlation with caution. Reverse causality is possible: countries with more reforestation activity may subsequently enact more forest legislation in response. Omitted variables are also a concern, since institutional quality, income level, and country size are all potentially correlated with both policy and projects.[12]

4.3 Project Effectiveness: The Reforestation Atlas

Figures 4, 5, and Appendix Figure B3 present the core empirical results: country-by-country SDID event studies for all 24 Sub-Saharan African countries in our sample, comprising 1,964 project sites. Following the atlas format of Kleven et al. (2025), each panel plots the estimated treatment effect by years since project launch, with 95% confidence intervals. Pre-treatment coefficients (years < 0) test the parallel trends assumption; post-treatment coefficients (years ≥ 0) trace the vegetation recovery trajectory. Pooling across all 1,964 sites, the average ATT is approximately 0.0027 NDVI units, corresponding to roughly 0.65 percent of the pre-treatment mean. Overall, we find the countries fall into three groups.

Successful programs (Figure 4; 9 countries). Burkina Faso, Central African Republic, Gambia,

[12] In the Appendix, we show that this relationship holds after controlling for log country size and log GDP per capita.

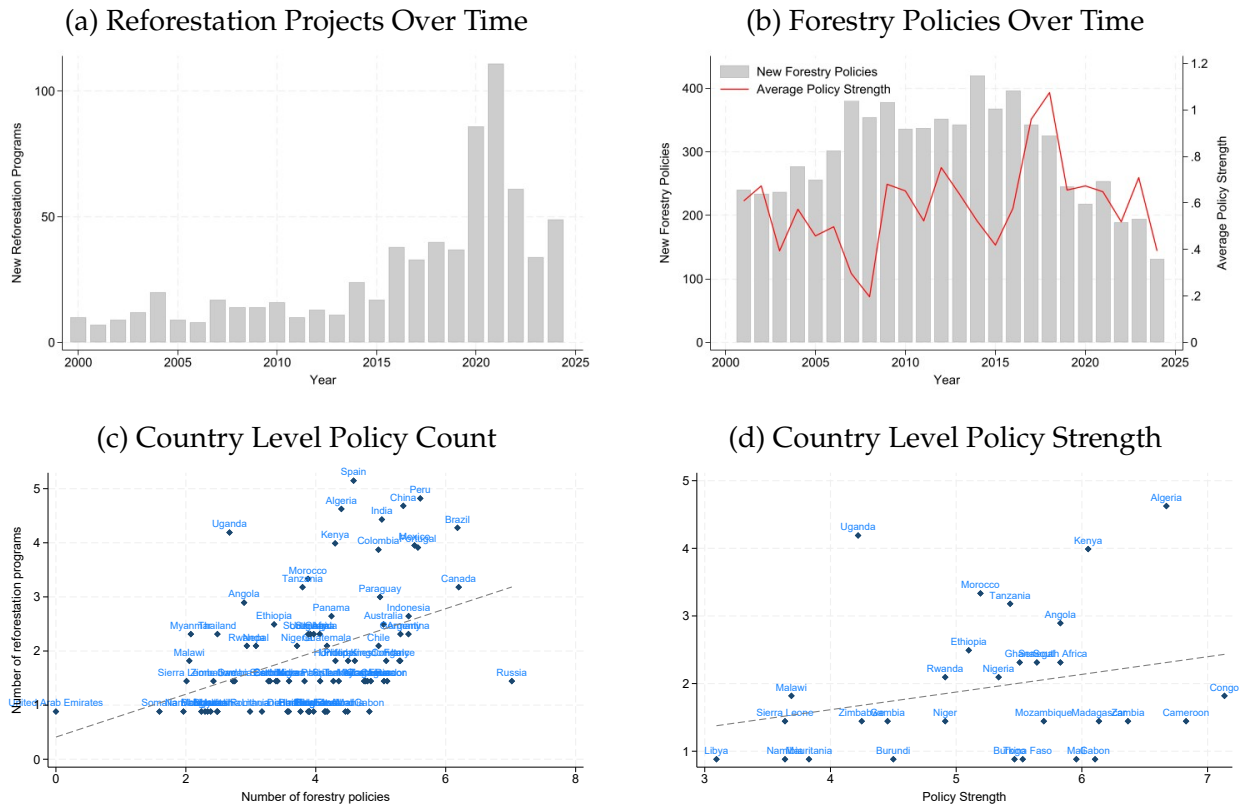


Figure 3. Forest Policy and Reforestation Investment

Notes: Panel (a) plots the annual count of new reforestation programs globally over 2000–2024. Panel (b) plots the annual count of new forest policy instruments enacted globally (bars, left axis) and the average policy strength constructed from LLM-coded document characteristics (line, right axis) over 2000–2024. Panel (c) plots the cross-country relationship between the cumulative number of forest policies (x-axis, asinh-transformed) and the total number of reforestation programs (y-axis, asinh-transformed) as of 2024. Panel (d) plots the same relationship using cumulative policy strength on the x-axis. Country labels are shown for selected countries. Forest policy data are from FAO FAOLEX; Reforestation project data are from John et al. (2025).

Guinea-Bissau, Guinea, Madagascar, Mali, Nigeria, and Zambia all show a clear and statistically significant NDVI increase following project launch, with flat pre-trends confirming the success of the synthetic control match. Effects typically emerge within one to two years. In countries with longer post-treatment windows, such as Burkina Faso, Gambia, Guinea-Bissau, and Nigeria, effects continue growing for at least ten years, suggesting initial planting is followed by sustained forest growth. Countries such as Central African Republic, Guinea, Mali, and Zambia have shorter observation windows because their projects were launched more recently, but even in the short run there is detectable vegetation gain relative to the control group.

Ineffective programs (Figure 5; 9 countries). Angola, Burundi, Cameroon, Chad, Congo, Côte d'Ivoire, Dem. Rep. Congo, Malawi, and Mozambique show post-treatment coefficients centred near zero with wide confidence intervals, indicating no detectable vegetation gain on average.

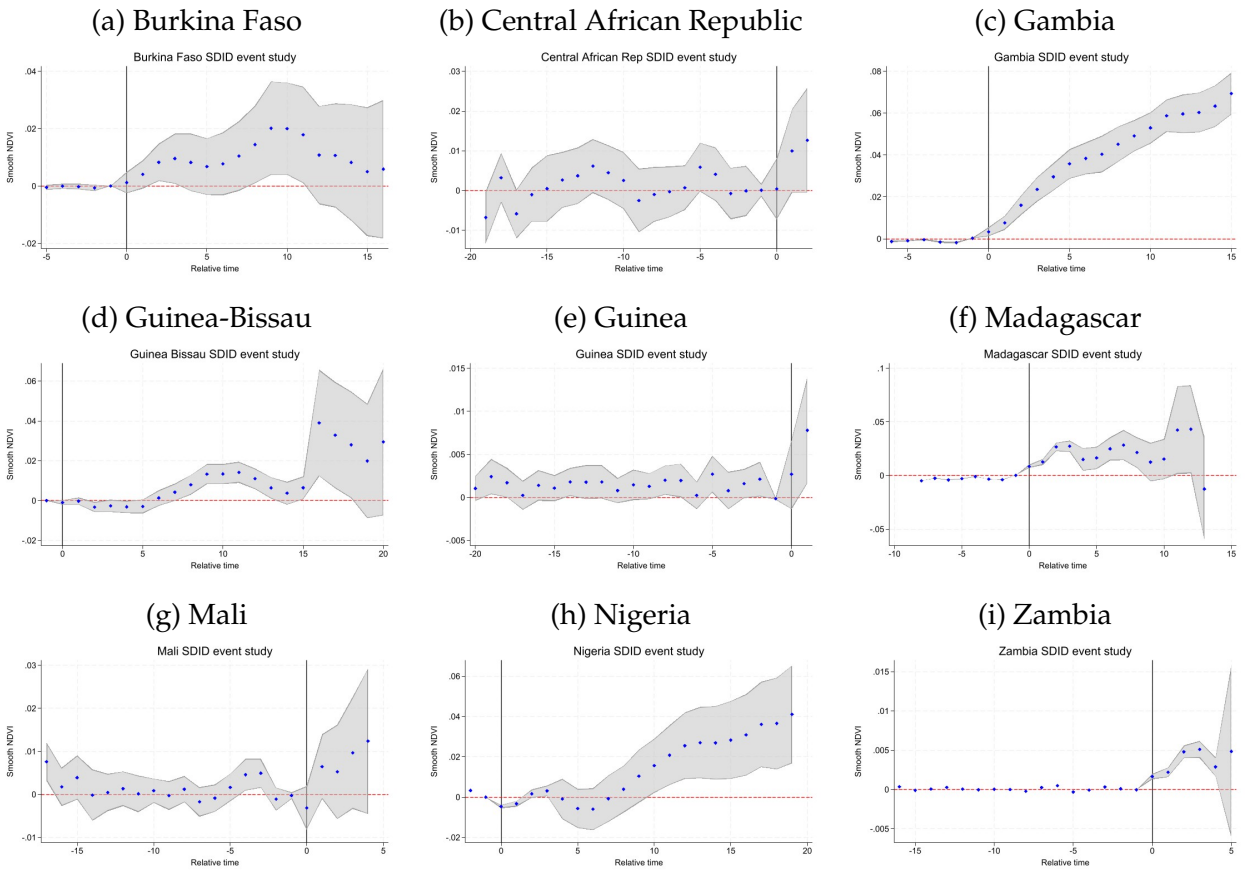


Figure 4. SDID Event Study: NDVI Effects — Successful Countries

Notes: Each panel plots the SDID event study coefficients for one country, with placebo-based inference. The outcome variable is NDVI. Each treated project site is matched to 100 control sites drawn from GEE satellite extracts. The x-axis shows years relative to project launch (treatment year). The shaded band represents 95% confidence intervals. NDVI data are from Google Earth Engine satellite extracts; project treatment years are from the ATLAS afforestation project database.

Pre-trends are generally flat across this group, with the exception of Angola where we observe a mild violation.

The absence of treatment effects could reflect three distinct mechanisms. First, project developers may have reported inaccurate site locations, so that our satellite measurements are not capturing the actual planting area. Second, even if site locations are correct, projects may have been registered but never implemented on the ground. Both of these are consistent with the well-documented additionality concerns in the voluntary carbon market literature (West et al., 2020; Probst et al., 2024). We assess these two mechanisms directly by restricting to sites with high LDIS scores (John et al., 2025), which flags projects with suspicious geographic boundaries. We will also use the richness of project documentation as a proxy for location reliability, under the assumption that projects with more detailed and well-written documentation are also more likely to have provided accurate geographic boundaries and have put actual effort on it. We are cur-

rently collecting project PDFs and using an LLM pipeline to score documentation quality, and will test whether restricting to well-documented projects yields higher average treatment effects. Third, even genuinely planted sites may fail to retain tree cover if surrounding communities rely on the land for agriculture or livestock: cattle and goats grazing on young saplings, or farmland expanding into planted areas, would produce the near-zero treatment effects we observe. These mechanisms carry very different policy implications. The first two call for stronger regulation and independent verification of voluntary carbon market projects. The third calls for better community engagement, profit-sharing arrangements, and local governance of planted areas. We return to this distinction in the heterogeneity analysis below.

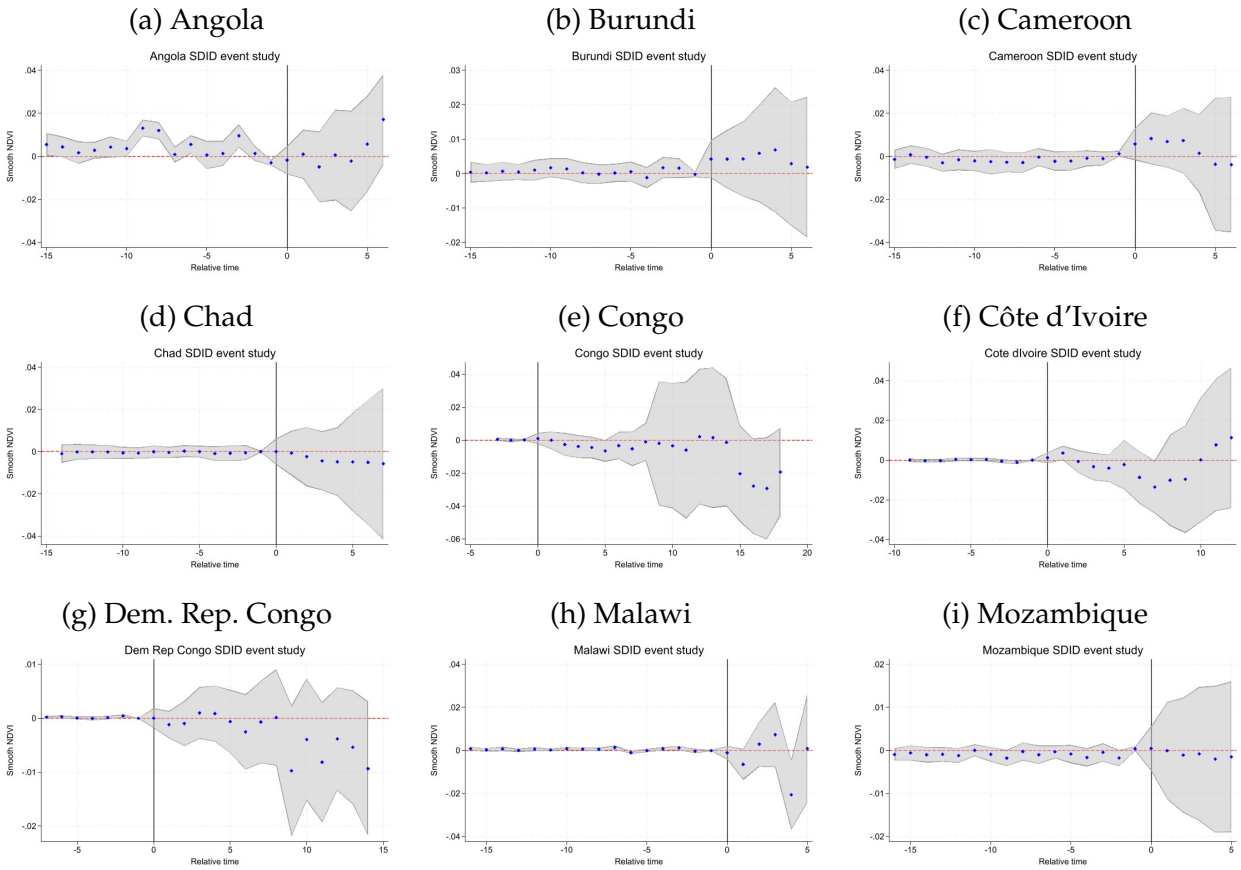


Figure 5. SDID Event Study: NDVI Effects — Ineffective Countries

Notes: Each panel plots the SDID event study coefficients for one country, with placebo-based inference. The outcome variable is NDVI. Each treated project site is matched to 100 control sites drawn from GEE satellite extracts. The x-axis shows years relative to project launch (treatment year). The shaded band represents 95% confidence intervals. NDVI data are from Google Earth Engine satellite extracts; project treatment years are from the ATLAS afforestation project database

Others (Appendix Figure B3; 6 countries). This group is heterogeneous and we discuss each case separately. Lesotho and Mauritania exhibit significant pre-trend violations, indicating that the synthetic control matching is not achieving balance for these countries. We are exploring

alternative matching approaches such as entropy balancing and propensity score matching to address this. Gabon has only one year of post-treatment observations, preventing any inference about program effectiveness.

Ethiopia, Namibia, and Niger show a different pattern: negative treatment effects, at least in the short run, meaning reforestation sites become less green than controls after project launch. Ethiopia's coefficient gradually converges back toward zero over time. Niger's effect is initially negative but eventually turns positive after approximately ten years, suggesting a delayed recovery once the initial disturbance subsides. Namibia's negative effect is less precisely estimated. A natural explanation for this pattern is that reforestation projects attract people to move in or start farming nearby, precisely because the area is expected to become more productive after the reforestation effort. This additional human activity around treated sites would reduce vegetation relative to control polygons in the short run, generating the negative treatment effects we observe.

To illustrate how dramatically human encroachment can degrade vegetation, Appendix Figures D1 and D2 present Google Earth satellite imagery for a project site in Ghana. The imagery shows the site transitioning from bare land to dense vegetation following planting, and then, in later periods, the visible emergence of cleared patches and farming activity within the project boundary. This anecdotal evidence illustrates the mechanism: once people move into a planted area and begin farming or grazing, vegetation loss can be rapid and substantial, consistent with the negative NDVI treatment effects we document in Ethiopia, Namibia, and Niger. The heterogeneity analysis in Section 4.4 examines this mechanism more systematically across all countries.

4.4 Heterogeneity and Mechanisms

Forest policy. Figure 6 examines whether the national forest policy environment predicts reforestation effectiveness. Panel (a) plots the country-mean ATT against the cumulative count of forest policy instruments from FAOLEX, asinh-transformed. The relationship is positive but with a very small magnitude: a log point increase in the policy count is associated with an ATT increase of 0.0009 NDVI units. Panel (b) plots the country-mean ATT against a policy strength index constructed using an LLM pipeline that scores each policy document for the strength of its enforcement provisions and binding commitments. A log point increase in the policy strength index is associated with an ATT increase of 0.0016 NDVI units. Gambia stands out as a notable high performer with both moderate policy strength and the highest country-mean ATT in the sample. Countries at the high end of the policy strength distribution, such as Congo, Niger, and Ethiopia,

cluster near zero ATT, which tempers the overall slope.

The two panels look broadly similar, which is unsurprising for two reasons. First, countries that enact more policies mechanically accumulate higher total strength scores, since both measures are cumulative.[13] Second, countries that invest more in forest governance tend to do so on both margins simultaneously: they both enact more laws and write more stringent ones. The two measures are therefore highly correlated in the cross-section. Despite this, the strength index carries independent information: conditional on policy count, the strength index still predicts higher ATT, suggesting that the quality of legislation matters over and above its quantity. The pattern is consistent with the view that it is not just the number of policies that matters, but their enforceability.

One potential interpretation is that stronger policy frameworks reduce encroachment risk by creating credible deterrents against illegal logging and land clearing near planted sites, making it harder for surrounding communities and firms to degrade planted areas. Investment without governance therefore appears insufficient to deliver durable forest recovery. However, we cannot rule out that the relationship reflects selection rather than causation: countries with stronger governance may attract better-resourced and better-monitored projects, so the policy-effectiveness correlation could partly reflect project quality differences rather than a direct effect of institutions on tree survival.[14]

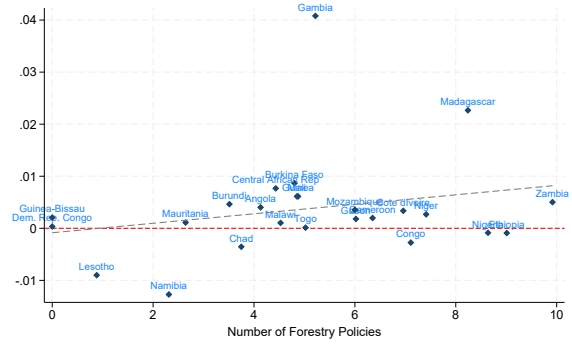
Pre-treatment human activity. We measure the presence of human activity within each project polygon using two variables: building density (area of buildings divided by polygon area, from the Google Open Buildings) and road density (total road length divided by polygon area, from OpenStreetMap). Due to computational constraints, we currently implement this analysis for Ethiopia, which has the largest number of project sites in our sample. We treat this as illustrative evidence and plan to extend the analysis to all countries as the pipeline scales.

Figure 7 documents the results. Panels (a) and (b) plot the mean ATT by building density and road density respectively, splitting sites into high and low groups at the median. The pattern is intuitive: sites with pre-existing human activity deliver substantially lower (negative) treatment

[13] As an alternative, we could also construct a policy strength measure that takes the average instead of summing across all scores, and a measure that uses only the most recent five years of policies to capture the current policy environment rather than its full history.

[14] In future work, we plan to use variation in policy enactment timing within countries to construct a more credible difference-in-differences test of whether strengthening forest legislation causally improves reforestation outcomes, and to use the LLM-coded policy characteristics to test which specific enforcement provisions are the active ingredient.

(a) Number of Forest Policies



(b) Policy Strength Index

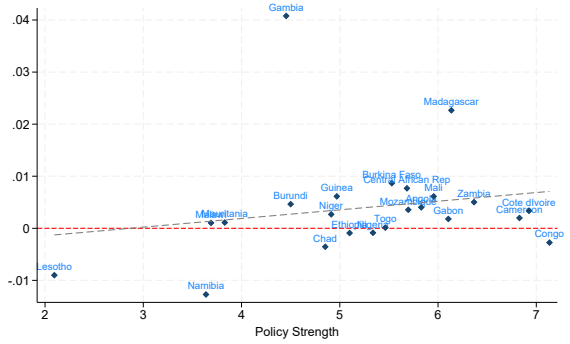
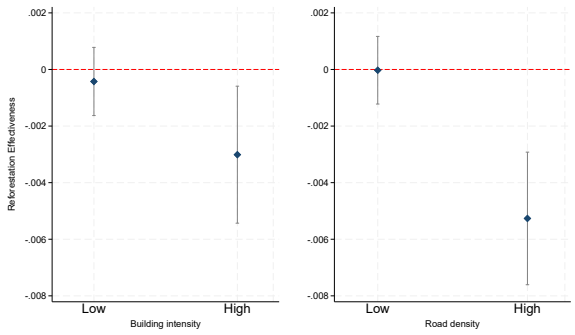


Figure 6. Forestry Policy and Reforestation Effectiveness

Notes: Each dot represents a country, labelled by country name. The y-axis plots the country-mean SDID ATT on NDVI, averaged across all sites within the country. Panel (a): x-axis plots the cumulative number of forest policy instruments, asinh-transformed. Panel (b): x-axis plots a policy strength index constructed from LLM-coded document characteristics, measuring the average enforcement stringency and binding commitments of enacted forest policies. The dashed horizontal line marks ATT = 0. The fitted line is estimated by OLS. Forest policy data are from FAO FAOLEX; ATT estimates are from the SDID analysis.

effects. Sites with high building density show an average ATT of approximately -0.003, while those with high road density show an average ATT of approximately -0.005. Panel (c) reinforces this with kernel density plots of the full ATT distribution separately for sites with and without any human activities. The distribution for sites with human activities shifts noticeably to the left.

(a) Building and Road Density



(b) Any Human Activities

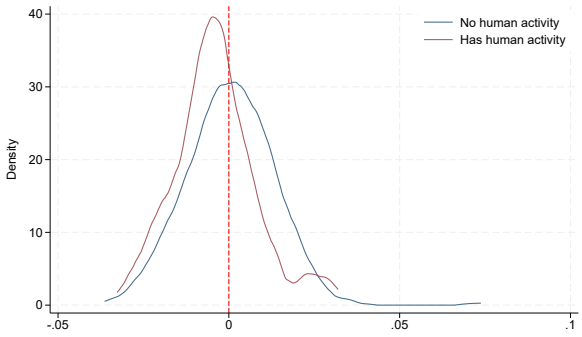


Figure 7. Treatment Effect Heterogeneity by Pre-Treatment Human Activity

Notes: The sample covers Ethiopian afforestation sites with SDID ATT estimates. Panel (a) plots the mean SDID treatment effect on NDVI for sites above and below the median pre-treatment building intensity (left) and road density (right). Building intensity is from the Google Open Buildings; road density is from OpenStreetMap. Error bars represent 95% confidence intervals based on the standard error. Panel (b) plots the kernel density of site-level SDID ATT estimates separately for sites with any human activity presence (buildings or roads) and sites with none.

The mechanism is straightforward. Sites with pre-existing roads and buildings are already embedded in areas of active human use. After planting, these areas face ongoing pressure from land clearing, agricultural encroachment, and livestock grazing that erodes any initial vegetation

gain. The infrastructure itself may also directly reduce the planted area available for tree growth. Together, these results suggest that a simple pre-screening of project sites for pre-existing human activity could substantially improve the average ecological return on reforestation investment, at low cost to program administrators.

Agricultural pressure from DHS. Figure 8 uses household survey data from the DHS to measure agricultural activity near each project site. For each project polygon, we compute three measures within a 50km buffer: a composite agricultural pressure index[15], agricultural land area, and goat ownership. We split sites into high and low groups at the mean of each measure and compare ATTs across the two groups.

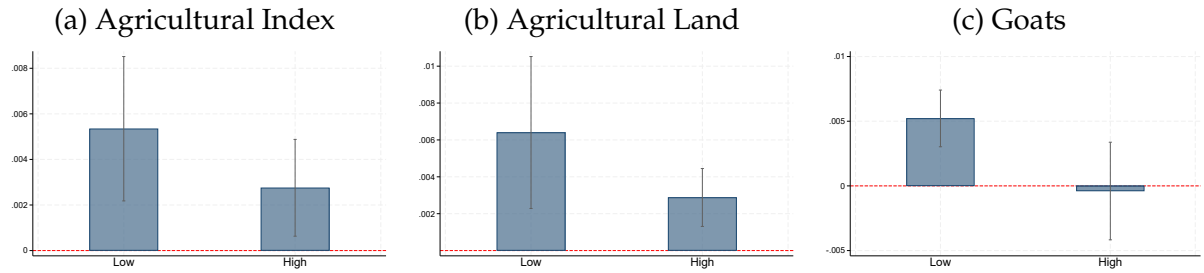


Figure 8. Reforestation Effectiveness by Household Agricultural Activities

Notes: Each panel compares the SDID ATT on NDVI for sites above vs. below the median of agricultural pressure measure. Dots show the mean ATT within each group; error bars represent 95% confidence intervals based on the standard error of the mean. Panel (a) uses a composite agricultural index, constructed as the average of four standardised measures: agricultural land ownership rate, agricultural land area, cattle ownership, and goat ownership, all measured within a 50 km buffer around each project polygon using DHS survey data. Panel (b) uses agricultural land area and Panel (c) uses goat ownership, both within a 50 km buffer.

Panel (a) shows that the composite agricultural index is strongly negatively correlated with ATT: projects in high-pressure areas deliver substantially lower vegetation gains than those in low-pressure areas, with the difference statistically significant at conventional levels. Panel (b) shows that agricultural land area alone is a strong predictor: sites surrounded by more farmland systematically underperform. Panel (c) shows that projects surrounded by more goats deliver near-zero treatment effects. Goats are well known to graze on young saplings and are among the primary biological threats to reforestation in dryland and semi-arid environments. The result that goat density predicts program failure, measured via satellite NDVI, provides micro-level evidence linking household-level livestock ownership to aggregate ecological outcomes. Additional disaggregated results for cattle ownership and land ownership are reported in Appendix and tell

[15]the average of four standardised variables: agricultural land ownership rate, agricultural land area, cattle ownership, and goat ownership

a consistent story.

It shows that households that depend on land for farming or livestock have a direct economic incentive to encroach on planted areas, either by grazing animals on young saplings or by converting planted land back to agriculture. The more land-dependent the surrounding population, the harder it is for planted forests to survive. This result has a clear policy implication: reforestation programs in agriculturally dense areas are unlikely to succeed without complementary interventions that reduce communities' dependence on the planted land, whether through profit-sharing arrangements, formal land tenure, or alternative livelihood programs.

Cropland pressure from Satellite. Figure 9 panel (a) corroborates the DHS results using an independent, satellite-derived measure of agricultural pressure. Rather than relying on household survey data, we use the Potapov et al. (2022) global cropland map to compute the share of cropland within a 50km buffer around each project polygon in the years prior to treatment. This measure captures the physical extent of agricultural land use rather than households' reported ownership or livestock counts, providing an independent test of the same underlying mechanism.

The results reinforce the DHS findings. Projects surrounded by higher pre-treatment cropland shares deliver significantly lower ATTs. The effect is not driven by any particular crop type or farming system but reflects the general intensity of agricultural land use near the project. Taken together with the DHS results, the cropland evidence strengthens the case that it is the local density of agricultural activity that determines whether planted forests survive.

Mining activity. Figure 9 panel (b) shows that proximity to mining activity is another predictor of program failure, operating through a different channel than agricultural pressure. Projects located within 50km of an active mine show systematically lower vegetation gains. The mechanism likely operates less through direct physical encroachment than through broader regional channels: mining booms raise land values and food demand across surrounding areas, expanding agricultural activity into forest margins (Aragón and Rud, 2013), while road and infrastructure networks built to service extraction sites open previously inaccessible land to encroachment and settlement (Giljum et al., 2022). Mining regions are also associated with weaker property rights enforcement and higher levels of informal economic activity (Burgess et al., 2012), conditions that reduce the institutional capacity to sustain forest restoration programs.

Taken together, these heterogeneity results tell a consistent story. Reforestation succeeds where the local economic incentive to convert or degrade planted land is low. Strong national forest gov-

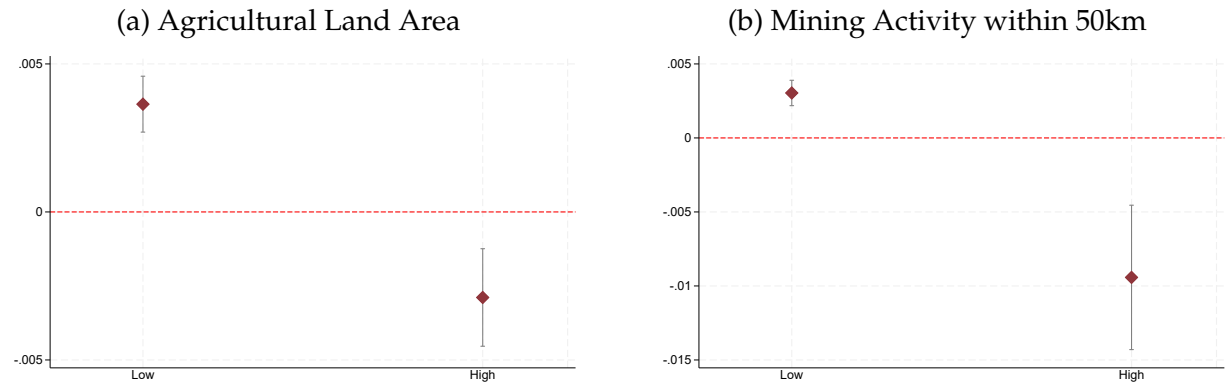


Figure 9. Reforestation Effectiveness by Agricultural Pressure and Mining Activity

Notes: Each panel plots the SDID ATT on NDVI for sites above vs. below the median of a site-level pressure measure. Dots show the mean ATT within each group; error bars represent 95% confidence intervals based on the standard error of the mean. Panel (a) uses the asinh of the expected cropland expansion pressure. Panel (b) uses a binary indicator for whether any active mining site falls within 50 km of the treated polygon. the "Low" group has no mining within 50 km and the "High" group has at least one mine. Cropland pressure data are from Potapov et al. (2022); mining data are from Maus et al. (2020).

ernance, low pre-existing human activity within project boundaries, limited agricultural dependence in surrounding communities, and absence of nearby extractive industries all predict higher treatment effects. Conversely, where any of these conditions fails, programs tend to deliver near-zero or even negative vegetation gains, regardless of the scale of investment or the identity of the implementing organisation.

The common thread across all four moderators is the opportunity cost of land. Where land is intensively used for farming, grazing, or extraction, the marginal value of keeping it forested is low relative to the returns from alternative uses. Planted trees face ongoing encroachment pressure that erodes any initial gain. National policy can partially offset this pressure by raising the cost of encroachment and providing legal protection for planted areas, which explains why governance predicts success even after controlling for local economic conditions. But policy alone cannot substitute for a local economic environment in which communities have a stake in forest survival. The implication for program design is clear: effective reforestation requires either targeting locations where land pressure is already low, or changing the incentives of surrounding communities so that maintaining the forest is in their economic interest, for example through profit-sharing arrangements, alternative livelihood programs, or formal land tenure that gives communities a direct stake in forest survival.

4.5 Program Location Selection

A natural question is whether programs are placed in locations that maximise the probability of ecological success, or whether voluntary carbon market incentives generate systematic misallocation. We assess this by regressing an indicator for treated project sites on pre-treatment DHS characteristics, comparing treated polygons to the randomly drawn control polygons within the same country.

Figure 10 presents the regression coefficients. The results reveal a clear and consistent pattern of positive selection toward easier, wealthier locations. On the agricultural side, all four livestock and land ownership measures are negative and statistically significant: DHS clusters near project sites have fewer households owning agricultural land (-0.044), fewer livestock owners (-0.064), fewer cattle (-0.085), and substantially fewer goats (-0.248). The goat coefficient is the largest in magnitude among the agricultural measures, consistent with our earlier finding that goat ownership is the single strongest predictor of program failure: developers appear to actively avoid areas where goat grazing poses a direct threat to young saplings. On the wealth side, households near project sites have significantly greater access to electricity ($+0.085$) and higher standardised wealth index scores ($+0.260$). Across all six measures, the pattern is consistent: projects are systematically located in wealthier, better-connected areas with lower agricultural and livestock pressure.[16]

Project developers therefore appear to systematically favour locations that are wealthier, better electrified, and under lower agricultural and livestock pressure. This is rational from a carbon credit issuance perspective: planting in areas with lower encroachment risk maximises the probability of verified vegetation gains and successful credit issuance. But it implies a fundamental misalignment between private incentives and social objectives. The areas most in need of reforestation intervention, those facing high deforestation pressure from dense agricultural and livestock activity, are precisely the areas that the voluntary carbon market systematically avoids.

This selection pattern has two implications. First, our ATT estimates identify the effect of reforestation for the sample of projects that are actually implemented, in locations that have al-

[16]These correlations should be interpreted as descriptive evidence of selection rather than causal estimates. Reverse causality is possible: if reforestation projects raise local wealth or expand electricity access over time, then the positive coefficients on wealth and electricity could partly reflect treatment effects on economic outcomes rather than selection. We partially address this by measuring all DHS characteristics prior to treatment, but the DHS survey waves are spaced several years apart and some post-treatment contamination cannot be ruled out. A fully credible test of program placement would require an instrumental variable that shifts project location independently of local characteristics. We view this as an important direction for the next step.

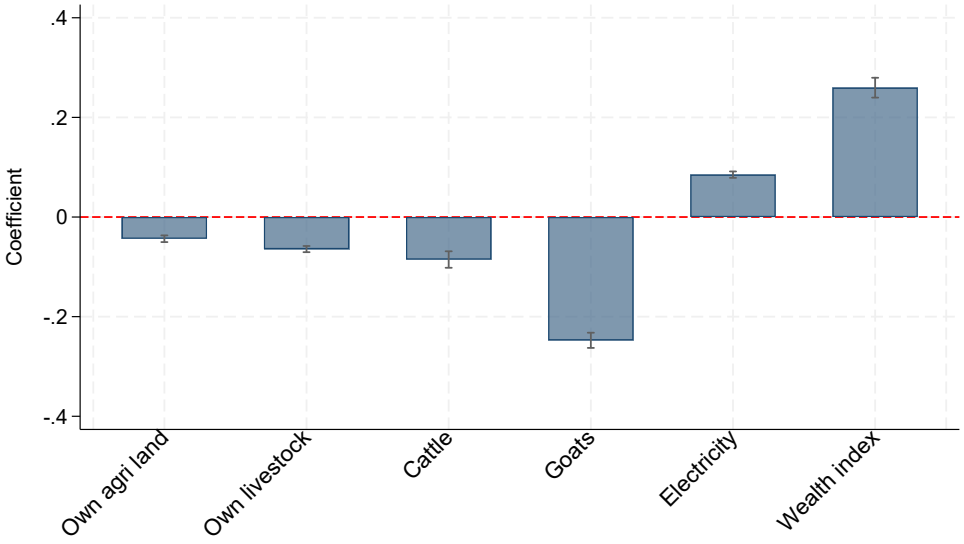


Figure 10. Reforestation Project Selection

Notes: OLS regressions of DHS cluster-level outcomes on an indicator for residing within 100km of a reforestation project site, with country and survey year fixed effects. Robust standard errors. The six outcomes are: agricultural land ownership rate, livestock ownership, cattle ownership, goat ownership (all levels), electricity access, and a standardised wealth index. Bars and error bars show point estimates and 95% confidence intervals. DHS data are from the Demographic and Health Surveys; project locations are from John et al. (2025).

ready been screened for viability. Effects may be smaller, or even negative, in the harder locations that are systematically avoided. This suggests our estimates provide an upper bound on what a scaled-up reforestation program, targeting a broader range of locations, would deliver on average. Second, the voluntary carbon market is not directing investment to where it is most needed ecologically. Areas with high deforestation pressure and high land demand are precisely the locations where successful reforestation would generate the largest additionality, yet these are the locations that projects systematically avoid. A better-designed policy would provide stronger financial incentives for projects in harder environments, whether through higher carbon credit prices for high-risk locations, performance bonuses for survival in high-pressure areas, or direct subsidies for community engagement and profit-sharing arrangements that reduce local land demand.

4.6 Discussion and Next Steps

The results above establish that reforestation programs deliver measurable vegetation gains in roughly one third of the countries in our sample. Program's effectiveness is systematically predicted by the national policy environment, pre-existing human activity within project sites, surrounding agricultural pressure from both satellite-derived cropland data and household survey

measures, and proximity to mining activity. Programs are also not placed where they are most needed: investment flows toward wealthier, less agricultural areas, consistent with voluntary carbon market incentives favouring sites with lower encroachment risk. Several important questions remain open and we discuss them here.

Causal identification of program effectiveness and placement. Our heterogeneity and program selection analyses are descriptive: they document systematic patterns but cannot fully establish causality. For program effectiveness, we regress site-level ATT estimates on a set of predictors, including national policy environment, pre-existing human activity, surrounding agricultural pressure, and proximity to mining. While these correlations are informative, they are subject to omitted variable bias. For example, areas with strong forest governance tend to also have better-funded implementing organisations, more experienced project managers, and greater community support and engagement, all of which could independently drive higher ATTs. It is therefore difficult to isolate whether governance and agricultural pressure affect program success through the mechanisms we propose, or whether they simply proxy for underlying ecological and institutional quality that we cannot directly observe. For program placement, we document that projects cluster in wealthier, less agricultural areas, but there can be reverse causality as we discussed in footnote 16. A more credible test of each would require a source of exogenous variation. For effectiveness heterogeneity, exploiting policy enactment timing within countries would allow a difference-in-differences test of whether strengthening governance causally improves outcomes. For placement, a natural candidate is an instrumental variable that shifts project location independently of local ecological conditions. We view these as important directions for the next step.

Local economic effects. Our analysis focuses on vegetation outcomes as measured by NDVI. But reforestation programs may also affect the welfare of nearby communities through multiple channels: employment in planting and maintenance activities, income from carbon credit revenues where profit-sharing arrangements exist, changes in local ecosystem services such as water availability and soil quality, and potential negative effects from land-use restrictions imposed on smallholders. Testing whether reforestation projects raise local household consumption or improve employment prospects using the DHS data described in Section 2.4 is a natural next step. The spatial linkage between DHS clusters and project polygons constructed for the heterogeneity analysis provides the infrastructure for this welfare analysis. Similarly, whether reforestation generates broader economic activity in surrounding areas, as proxied by nighttime lights or google

building data, remains an open question for this paper to explore in the next step.

Scaling to global coverage. Our main results cover a subsample of Sub-Saharan Africa. However, our database also contains projects across Asia, Latin America, and Europe, where the institutional environment, ecological conditions, and carbon market dynamics differ substantially. Extending the satellite extraction pipeline and the SDID analysis to all regions would allow us to test whether the mechanisms we identify in Africa generalise globally, and to provide a truly worldwide Reforestation Atlas. There is also a selection concern worth noting: our current sample is constrained to countries with relatively few projects due to computational reasons, which may be precisely the countries that program developers strategically avoid because they anticipate lower ecological returns. The countries with the most reforestation activity, such as India, Brazil, and China, may exhibit very different treatment effects and cost-benefit ratios. If high-activity countries were included, the average effectiveness and welfare returns we estimate could look substantially larger. Resolving this question requires the full global sample.

5 Back-of-Envelope Welfare Calculation

In this section, we translate our NDVI treatment effect estimates into a welfare calculation by converting vegetation gains into carbon sequestration, valuing sequestration at the social cost of carbon, and comparing the resulting climate benefit against project costs.

5.1 From NDVI to Carbon Sequestration

Our outcome variable is the estimated average treatment effect on the treated $\hat{\tau}$, which needs to be converted to tonnes of CO_2 equivalent to be valued at the social cost of carbon. We proceed in three steps.

Step 1: NDVI to aboveground biomass. We adopt the regression coefficient from Poudel et al. (2023), who regress field-measured aboveground biomass on NDVI and obtain a linear slope of $\alpha = 52.9$ tDM/ha per NDVI unit ($R^2 = 0.83$). Aboveground biomass gain per hectare is then:

$$\Delta \text{AGB}_i = \alpha \times \hat{\tau}_i = 52.9 \times \hat{\tau}_i \quad \text{(tDM/ha)}.$$

We use this coefficient as a rough empirical anchor rather than a precisely calibrated estimate.[17]

[17] The coefficient is derived from subtropical montane forest in the Chure region of Nepal and is unlikely

Step 2: Aboveground to total biomass carbon. Following the IPCC Tier 1 guidelines (Eggleston et al., 2006), we apply two standard adjustments. First, a root-to-shoot ratio of $R = 0.37$, the recommended value for dry tropical forest and savanna woodland, implies total biomass is $1.37\times$ aboveground biomass.[18] Second, a carbon fraction of $CF = 0.47$ converts dry biomass to carbon stock, reflecting the standard assumption that 47 percent of dry biomass weight is carbon.[19] Total carbon gain per hectare is:

$$\Delta C_i = \Delta \text{AGB}_i \times (1 + R) \times CF = \hat{\tau}_i \times 52.9 \times 1.37 \times 0.47 \approx \hat{\tau}_i \times 34.1 \quad \text{(tC/ha)}.$$

Step 3: Carbon to CO_2 equivalent. Carbon stock is converted to CO_2 equivalent using the molecular weight ratio $44/12 \approx 3.67$, the standard factor in IPCC and carbon market reporting protocols:

$$\Delta \text{CO}_2\text{e}_i = \Delta C_i \times \frac{44}{12} \approx \hat{\tau}_i \times 34.1 \times 3.67 \approx \hat{\tau}_i \times 125.1 \quad (\text{tCO}_2\text{e/ha}).$$

Combining the three steps, each additional NDVI unit gained on one hectare of treated land corresponds to approximately 125 tCO_2e sequestered.

5.2 Social Benefit

We value each tonne of sequestered CO_2 at a social cost of carbon of $\$185/\text{tCO}_2$, taken from Rennert et al. (2022), who combine an updated damage function, updated socioeconomic scenarios, and a Ramsey discount rate to obtain a central estimate of \$185 (95% CI: \$44–\$413) per tonne of CO_2 emitted in 2020.[20] The climate benefit for project is:

$$\text{Benefit}_i = \hat{\tau}_i \times 125.1 \times 185 \approx \hat{\tau}_i \times 23{,}144 \quad \text{(USD)}.$$

to transfer directly to sub-Saharan African savanna and open woodland, where NDVI saturates at lower biomass levels and the NDVI–biomass slope tends to be shallower and more variable across vegetation types. A more defensible approach would be to extract baseline aboveground biomass for each treated polygon from the pantropical biomass map of Avitabile et al. (2016), which provides field-calibrated estimates at 1 km resolution across Africa, and to infer biomass gains by scaling proportionally to the observed NDVI change. We retain the Poudel et al. (2023) slope for the present calculation and treat the result as an order-of-magnitude illustration only.

[18] IPCC (2006), Volume 4, Table 4.4.

[19] IPCC (2006), Volume 4, Section 2.2.

[20] This figure is broadly consistent with the US EPA's 2023 central administrative estimate of $\$190/\text{tCO}_2$ at a 2% near-term discount rate. Other recent estimates include Barrage and Nordhaus (2024), who obtain $\$176/\text{tCO}_2$ for 2020 using the DICE-2023 integrated assessment model with 2% discount rate, and Moore et al. (2024), who obtain $\$280/\text{tCO}_2$ once additional structural channels are incorporated.

For our sample average treatment effect $\bar{\tau}_{\text{ATT}}$, the average climate benefit is approximately $\bar{\tau}_{\text{ATT}} \times$ 23,144 USD.

5.3 Project Cost

We proxy project costs using implementation cost estimates from Busch et al. (2024), who compile 355 cost observations from 99 source documents across 26 low- and middle-income countries and use regression models to predict implementation costs spatially as a function of local wage rates, land accessibility, and terrain. We focus on implementation cost rather than their total abatement cost measure, which additionally incorporates opportunity cost (the foregone agricultural income from converting land to forest). From a voluntary carbon market project perspective, implementation cost is the appropriate cost concept, as opportunity cost is borne by landholders rather than project funders.[21]

Implementation costs differ substantially by project type. Busch et al. (2024) report a median of \$140/ha (interdecile range \$47–183/ha) for natural regeneration and \$3,729/ha (interdecile range \$1,530–11,675/ha) for plantations, the latter inclusive of replanting across harvest cycles. Since our sample contains both project types, we use a simple average of the two medians, \$1,935/ha, as our baseline cost estimate.[22]

5.4 Benefit-Cost Ratio and Implied Cost per Tonne

The average benefit-cost ratio per hectare is:

$$\text{BCR} = \frac{\text{Benefit}}{\text{Cost}} = \frac{\bar{\tau}_{\text{ATT}} \times 23{,}144}{1{,}935} = 11.96 \times \bar{\tau}_{\text{ATT}}.$$

this BCR depends only on the treatment effect and the implementation cost per hectare. With $\bar{\tau}_{\text{ATT}} = 0.0027$, the average BCR is equal to 0.032. This number is much smaller than 1, indicating that the social benefit of reforestation project is much smaller than it's cost. The contrast between project types is substantial: at the same treatment effect $\hat{\tau}_i$, a plantation project faces roughly 27

[21] Median opportunity cost in Busch et al. (2024) is \$4,770/ha in areas where natural regeneration is feasible and \$5,054/ha in areas where plantation is feasible. Combined with implementation costs of \$140/ha and \$3,729/ha respectively, total costs are approximately \$4,910/ha for natural regeneration and \$8,783/ha for plantations.

[22] A more precise approach would assign each project its corresponding cost based on recorded project type, allowing project-level benefit-cost ratios to reflect the true cost structure. We will work on this in the future.

times higher implementation costs than a natural regeneration project, implying a BCR approximately 27 times smaller. An equivalent and perhaps more intuitive metric is the implied cost per tonne of CO_2 sequestered:

$$\text{Cost per tCO}_2\text{e} = \frac{COST}{\bar{\tau}_{\text{ATT}} \times 125.1}.$$

With $\bar{\tau}_{\text{ATT}} = 0.0027$, the the implied cost per tonne of CO_2 sequestered is $\$5728/\text{tCO}_2$.

5.5 Comparison with the Literature

Our implied cost per tonne of CO_2 sequestered is $\$5{,}728/\text{tCO}_2$, which is substantially above existing benchmarks. Jayachandran et al. (2017) estimate a cost of approximately $\$46/\text{tCO}_2$ avoided for a payments-for-ecosystem-services programme in Uganda that compensated landowners to prevent deforestation. Busch et al. (2024) report median abatement costs of $\$23.80/\text{tCO}_2$ for natural regeneration and $\$23.00/\text{tCO}_2$ for plantations across low- and middle-income countries, based on modelled carbon accumulation assuming projects achieve their expected growth trajectories. Our estimate is roughly 125 times higher than these benchmarks, reflecting the fact that most projects in our sample deliver far smaller vegetation gains than engineering models assume.

The comparison with Busch et al. (2024) illustrates the value of causal identification. Their abatement cost estimates assume projects achieve modelled carbon accumulation rates derived from observed forest growth in favourable conditions. Our results suggest this assumption fails for a substantial share of projects: for the ineffective group, NDVI gains are indistinguishable from zero, implying an unbounded cost per tonne. Even for the most successful group in our sample, the implied cost per tonne is $\$101.9/\text{tCO}_2$, still more than four times the Busch et al. (2024) median benchmark of $\$23/\text{tCO}_2$. The gap between engineering projections and realised outcomes is stark: simply registering a reforestation project and planting trees does not guarantee that forest cover will increase. Without accounting for encroachment, local land pressure, and governance failures, modelled cost estimates will systematically overstate the climate return on reforestation investment. Our results suggest that engineering benchmarks, while useful for planning purposes, cannot substitute for independent satellite-verified evaluation of whether planted forests actually survive.

Four caveats apply. First, our benefit calculation captures carbon gains over the event study window only; sequestration continues as trees mature beyond our observation horizon, so our estimates are conservative on this dimension. Second, the NDVI-to-biomass conversion is imprecise, and the true slope for African savanna ecosystems may differ from the value we adopt. Third,

the cost figures from Busch et al. (2024) are global medians predicted from a regression model and may not accurately reflect costs in specific African country contexts. Fourth, we abstract from co-benefits including biodiversity, watershed protection, and local income effects, as well as potential welfare costs from land-use restrictions on smallholders, none of which are captured in our carbon-only calculation. We plan to address these properly in the future work.

6 The Model

6.1 Basics

We consider an economy in discrete time $t = 0, 1, 2, ...$ with a common forest stock. The aggregate forest stock at the beginning of period t is denoted by $F_t \geq 0$. The initial stock of forest is given as F_0. There are two types/groups of agents, indexed by $j \in \{H, L\}$. Agents derive utility from consumption c_{jt} and the stock of the forest F_t:

$$\sum_{t=0}^{\infty} \beta^t \left[u\left(c_{jt}\right) + \phi_j v\left(F_t\right) \right], \tag{1}$$

where $0 < \beta < 1$ and $\phi_H > \phi_L \geq 0$, i.e., type H agents intrinsically value the forest stock more than type L agents.[23] We assume that each type makes up an equal share of the population.

For each period t, agents of type $j \in \{H, L\}$ chooses to deforest $d_{jt} \geq 0$, which in turn generate income $A_j q\left(d_{jt}, a_{jt}\right)$, but also reforest $e_{jt} \geq 0$ with cost $m\left(e_{jt}\right)$ (in monetary terms), which adds up to the forest stock[24]. We assume that $q(d, a)$ is increasing and concave in d to capture the production technology with diminishing marginal returns to natural resource use, and $m(e)$ is increasing and convex in e to capture the planting and maintenance cost (i.e., the marginal cost of expanding reforestation within a period rises because the cheapest and most suitable sites are planted first and because nursery, seedling-supply, labour and monitoring capacity are congestible in the short run). Importantly, we also assume that production uses both physical capital input a_{jt} and natural capital d_{jt}. Agents face the inter-temporal consumption-saving problem with budget

[23] We adopt the framework of Acemoglu et al. (2012) in which economic agents derive utility directly from environmental quality. In our model, environmental quality is determined by the stock of forest and the carbon it sequesters.

[24] This formulation should be interpreted at the aggregate level. Each type j represents a large strategic unit, such as a community, region, or country-sector coalition, within which deforestation and reforestation may be carried out by different subgroups or projects. The variables d_{jt} and e_{jt} therefore capture net effort on the two margins rather than mutually exclusive choices of a single individual.

constraint:

$$a_{jt+1} = (1-\delta)a_{jt} + A_j q\left(d_{jt}, a_{jt}\right) - c_{jt} - m\left(e_{jt}\right), \tag{2}$$

where δ is the depreciation rate of physical capital and A_j is the productivity parameter. The forest stock evolve according to

$$F_{t+1} = G\left(F_t\right) + \sum_{j\in\{H,L\}} \left[\zeta e_{jt} - d_{jt}\right], \tag{3}$$

where $G(F)$ is the natural growth function of forest, and $\zeta \in (0,1)$ captures the overall effectiveness of reforestation programs, or the distinction between mature forest and planted young forests. We further impose $F_t \geq 0$ and $0 \leq d_{Ht} + d_{Lt} \leq F_t$ for all t, where the highest level of deforestation for each period is bounded above by the existing forest stock.

6.1.1 Decentralized Equilibrium

Each group chooses $\{c_{jt}, e_{jt}, d_{jt}\}_{t=0}^{\infty}$ to maximize (1) subject to (2) and (3), taking the other group's choices as given. The optimality conditions (derived in Appendix C.2) are the Euler equation

$$u'(c_{jt}) = \beta R_{j,t+1}\, u'(c_{j,t+1}), \qquad R_{j,t+1} \equiv 1 - \delta + A_j\, q_a(d_{j,t+1}, a_{j,t+1}), \tag{4}$$

the forest-management condition equating the marginal return to clearing, the marginal cost of planting, and the shadow value of forest:

$$A_j\, q_d(d_{jt}, a_{jt}) = \frac{1}{\zeta} m'(e_{jt}) = \frac{\mu_{jt}}{\lambda_{jt}}, \tag{5}$$

and the shadow prices

$$\lambda_{jt} = u'(c_{jt}), \tag{6}$$

$$\mu_{jt} = \beta\left[\phi_j v'(F_{t+1}) + \mu_{jt+1} G'(F_{t+1})\right]. \tag{7}$$

A higher ϕ_j raises μ_{jt}, reducing deforestation and increasing reforestation for group H. The critical limitation is that μ_{jt} reflects only group j's own preference ϕ_j: no group internalizes the forest amenity it provides to the other. This common-pool externality generates excessive deforestation and insufficient reforestation. Appendix C.1 characterizes the equilibrium formally and derives the corrective Pigouvian tax.

6.2 Forest and the Global Carbon Market

We extend to two countries $n \in \{N, S\}$, each containing groups $j \in \{H, L\}$, to study both the global carbon externality and whether a carbon market can restore efficiency. Country N is more productive: $A_{Nj} = \gamma A_j$ with $\gamma > 1$. The budget constraint is the country-level analogue of (2),

$$a_{nj,t+1} = (1-\delta)a_{njt} + A_{nj}q(d_{njt}, a_{njt}) - c_{njt} - m_n(e_{njt}), \tag{8}$$

where $m_n(e)$ is country-specific, reflecting cheaper reforestation in the South. Net carbon emissions from country n are

$$E_{nt} = \sum_{j\in\{H,L\}} \left[\xi_d d_{njt} - \xi_e \zeta e_{njt}\right], \tag{9}$$

and the global atmospheric carbon stock evolves as

$$X_{t+1} = (1-\rho_X)X_t + \sum_{n\in\{N,S\}} E_{nt}, \qquad \rho_X \in (0,1), \tag{10}$$

creating a cross-country externality: forest management in one country affects the climate faced by all.

6.2.1 The global social planner

A global social planner maximizes aggregate utility net of carbon damage,

$$\max \sum_{t=0}^{\infty} \beta^t \left\{ \sum_{n,j} \left[u(c_{njt}) + \phi_j v(F_{nt})\right] - \chi D(X_t) \right\},$$

subject to (8), the local forest laws of motion, and (10), where $D'(X) > 0$ captures global carbon damages. Let Ω_{nt} be the planner's shadow value of country n's forest stock and $\Psi_t \geq 0$ the shadow cost of global atmospheric carbon. These satisfy

$$\Omega_{nt} = \beta \left[(\phi_H + \phi_L)v'(F^G_{n,t+1}) + \Omega_{n,t+1}G'(F^G_{n,t+1})\right], \tag{11}$$

$$\Psi_t = \beta \left[\chi D'(X^G_{t+1}) + (1-\rho_X)\Psi_{t+1}\right]. \tag{12}$$

The planner's forest-management conditions are

$$A_{nj}q_d(d^G_{njt}, a^G_{njt}) = \frac{\Omega_{nt} + \xi_d \Psi_t}{u'(c^G_{njt})}, \tag{13}$$

$$\frac{m_n'(e_{njt}^G)}{\zeta} = \frac{\Omega_{nt} + \xi_e \Psi_t}{u'(c_{njt}^G)}. \tag{14}$$

Each condition has two components: the local forest amenity value Ω_{nt}/u' and the global carbon shadow cost Ψ_t/u'. The planner thus internalizes two externalities that laissez-faire ignores: the within-country amenity accruing to both groups (via $\phi_H + \phi_L$ in (11)) and the cross-country carbon externality (via Ψ_t in (12)). Appendix Proposition 5 collects the full characterization.

6.2.2 Laissez-faire in the global economy

Under laissez-faire, each group in each country maximizes its own utility taking the global carbon stock as given. The Euler equation follows (4) indexed by country, and the forest-management condition mirrors (5) with the private shadow value $\mu_{njt}^{LF} = \beta[\phi_j v'(F_{n,t+1}^{LF}) + \mu_{nj,t+1}^{LF} G'(F_{n,t+1}^{LF})]$.

Laissez-faire omits two terms: the within-country forest benefit to the other group (since μ_{njt}^{LF} uses ϕ_j rather than $\phi_H + \phi_L$) and the global carbon shadow cost Ψ_t. The productivity asymmetry $A_{Nj} = \gamma A_j > A_{Sj}$ further amplifies the distortion: country N has a stronger private return to deforestation.

6.2.3 Carbon-rights market

A global cap on net emissions is enforced by trading carbon rights at price p_t^C:

$$\sum_{n \in \{N,S\}} E_{nt} \leq \bar{E}_t, \qquad \sum_n B_{nt} = \bar{E}_t, \tag{15}$$

where B_{nt} is country n's initial allocation. The carbon price adjusts private incentives: deforestation now carries cost $p_t^C \xi_d$ and reforestation earns credit $p_t^C \xi_e \zeta$, giving the forest-management condition

$$A_{nj} q_d(d_{njt}^C, a_{njt}^C) - p_t^C \xi_d = \frac{1}{\zeta} m_n'(e_{njt}^C) - p_t^C \xi_e = \frac{\mu_{njt}^C}{u'(c_{njt}^C)}, \tag{16}$$

where the private forest shadow value μ_{njt}^C follows (7) indexed by country. Market clearing requires

$$\sum_{n,j} \left[\xi_d d_{njt}^C - \xi_e \zeta e_{njt}^C \right] = \bar{E}_t. \tag{17}$$

Proposition 1 *A competitive carbon market alone does not generally implement the global social optimum.*

The proof (Appendix C.6) shows that simultaneously implementing the planner's deforestation and reforestation conditions requires different carbon prices when $\xi_d \neq \xi_e$, because the

planner's within-country amenity wedge Ω_{nt} (internalizing $\phi_H + \phi_L$) diverges from the private shadow value μ^C_{njt} (internalizing only ϕ_j). The carbon market prices the carbon dimension of forest management but leaves the local amenity distortion uncorrected. In the quantitative exercise, we measure how much of the welfare gap it nevertheless closes.

6.3 Parameterization, and Calibration

Functional Forms For utility over consumption we adopt the constant relative risk aversion (CRRA) form,

$$u(c) = \frac{c^{1-\sigma}}{1-\sigma}, \qquad \sigma > 0. \tag{18}$$

Utility over the forest stock is also CRRA,

$$v(F) = \frac{F^{1-\eta}}{1-\eta}, \qquad \eta > 0, \tag{19}$$

following Acemoglu et al. (2012), where η governs the curvature of environmental preferences. The natural growth of the forest is logistic,

$$G(F) = F + rF\left(1 - \frac{F}{K}\right), \tag{20}$$

where $r > 0$ is the intrinsic growth rate and $K > 0$ the carrying capacity, so that absent net forest use the stock converges to $F = K$ with $G'(K) = 1 - r$.

Deforestation and physical capital combine in a constant-elasticity-of-substitution (CES) technology,

$$q(d, a) = \left(\omega\, d^{\psi} + (1-\omega)\, a^{\psi}\right)^{\alpha/\psi}, \qquad \alpha \in (0,1),\ \ \psi < 1,\ \psi \neq 0,\ \ \omega \in (0,1), \tag{21}$$

where ω is the deforestation share, α indexes the (decreasing) returns to scale, and $1/(1-\psi)$ is the elasticity of substitution between d and a. The reforestation cost function is convex,

$$m(e) = \kappa e^{\theta}, \qquad \theta > 1, \tag{22}$$

reflecting the convex costs of planting and maintaining trees at larger scales. Then $m'(e) = \theta\kappa e^{\theta-1}$ and the reforestation condition $m'(e_j) = \zeta s_j$ has the closed form $e_j = \left(\frac{\zeta s_j}{\theta\kappa}\right)^{1/(\theta-1)}$.

Table 1. Calibration

Parameter	Symbol	Value	Source / target
Discount factor	β	0.96	Cooley and Prescott (1995)
Capital depreciation	δ	0.08	Gomme and Rupert (2007)
Risk aversion (consumption)	σ	2	Standard CRRA
Curvature of forest value	η	2	Symmetric to σ
Forest carrying capacity	K	1	Normalisation (cover share)
Intrinsic forest growth	r	0.15	Logistic regrowth
Returns to scale, harvest	α	0.60	Span-of-control Lucas (1978)
CES substitution	ψ	0.5	Standard CES
CES land weight	ω	0.4	Assumption
Reforestation cost curvature	θ	2	Quadratic adjustment cost; Hayashi (1982)
Reforestation cost scale	$\kappa, \kappa_N, \kappa_S$	1.9, 2.8, 1.2	\$1,935 ha^{-1} implementation cost
Reforestation effectiveness	ζ	0.14	NDVI treatment effect estimates
TFP, low-valuation group	A_L	0.42	$A_L > A_H$: comparative advantage in extraction
TFP, high-valuation group	A_H	0.40	
Forest preference, high	ϕ_H	0.80	Normalised preference spread
Forest preference, low	ϕ_L	0.20	
Carbon depreciation	ρ_X	0.02	30-50 years atmospheric residence time
Carbon damage normalization	χ	0.01	Assumption
Forest-to-carbon rate	ξ_e, ξ_d	0.65, 0.8	Eggleston et al. (2006)
Supply of carbon rights	$\bar{E}_t$	$\bar{E}_t = \bar{\kappa}\, E_t^{LF}$	Varying $\bar{\kappa}$

For the extension model with carbon market, we assume

$$D(X) = \frac{1}{2}X^2 \tag{23}$$

to capture quadratic cost of environmental damages following the climate economics literature (Nordhaus, 2017).

Calibration Table 1 collects the baseline parameter values; the remainder of this subsection records the evidence behind each.

Preferences and depreciation ($\beta, \delta, \sigma, \eta$). The annual discount factor $\beta = 0.96$ corresponds to a pure rate of time preference of roughly four per cent, the standard value in the quantitative-macro tradition (Cooley and Prescott, 1995; King and Rebelo, 1999). We set the annual depreciation rate to $\delta = 0.08$; values in the 0.06–0.10 range are conventional and depend on the breadth of the capital concept (Cooley and Prescott, 1995; Gomme and Rupert, 2007). We adopt $\sigma = 2$ for the coefficient of relative risk aversion and impose the symmetric curvature $\eta = 2$ on the forest-amenity flow, both within the range routinely used for CRRA preferences.

Forest growth (r, K). For the logistic forest function, we normalise carrying capacity to $K = 1$, so that F is read as a forest-cover share. The intrinsic growth rate $r = 0.15$ implies a characteristic regrowth time of order $1/r \approx$ 6–7 years in the low-density regime, slowing as $F \to K$; this is consistent with the multi-decadal recovery horizons documented for secondary forest, for which a positive logistic intrinsic growth rate fits observed regeneration (Acevedo et al., 2012).

Harvest technology (α, ψ, ω). Output combines cleared land d and capital a through a CES aggregator raised to the power α, $q(d, a) = \left(\omega d^{\psi} + (1-\omega)a^{\psi}\right)^{\alpha/\psi}$. The exponent $\alpha = 0.60$ delivers decreasing returns to the modelled input bundle, as in span-of-control formulations (Lucas, 1978). The substitution parameter ψ governs the elasticity $\sigma_q = 1/(1-\psi)$ between cleared land and capital. We assume the CES case $\psi = 0.5$ ($\sigma_q = 2$). We also assume the distribution parameter $\omega = 0.4$ being set to the cost share of cleared land in extractive output, leaving the larger weight on reproducible capital.

Reforestation cost ($\theta, \kappa, \kappa_N, \kappa_S$). Per-period reforestation effort e carries the convex cost $m(e) = \kappa e^{\theta}$. We set $\theta = 2$, the canonical quadratic adjustment cost (Hayashi, 1982). The scale parameter κ is pinned by the implementation cost of reforestation programmes. We use a central estimate of \$1,935 ha^{-1}, which sits inside the range reported in the restoration-cost literature (Busch et al., 2024; Cook-Patton et al., 2020). Normalizing this cost to fit the scale of our benchmark model gives us $\kappa = 1.9$. In our extension model with multiple countries with carbon market, we assume the cost of reforestation is cheaper in the South country than the North country because of lower land costs and labor costs: $\kappa_N = 2.8 > \kappa_S = 1.2$.

Reforestation effectiveness (ζ). The forest law of motion credits reforestation through $\zeta \sum_j e_j$, where ζ converts reforestation effort into durable forest stock. We discipline ζ using the reduced-form treatment effect of the reforestation programme on remotely sensed greenness: our estimates of average treatment effect is around $\Delta\text{NDVI} = 0.0027$, and the aboveground biomass gain per hectare is $\Delta AGB = 52.9 \times 0.0027 = 0.14$ as in Section 5.1. We therefore calibrate $\zeta = 0.14$.

Carbon Emission and Storage ($\rho_X, \chi, \xi_d, \xi_e$,) For the annual depreciation rate of the atmospheric carbon stock ρ_X, we calibrate it to 0.02 to capture a 30–50 year effective atmospheric residence time. We normalize $\chi = 0.01$ to fit the scale of our model. Consistent with Section 5.1, we take $\xi_e = 0.65$ which is the coefficient translating aboveground biomass to carbon storage. We then take $\xi_d = 0.8 > \xi_e$ to reflect the fact that deforestation releases more than just the

aboveground carbon that planting accumulates, such as belowground biomass and other foregone emissions.

Initial Conditions (F_0, A_{j0}) Lastly, for our transition path in the benchmark model, we will assume $F_0 = 0.9$ and $a_{H0} = a_{L0} = 0.3$ to capture a capital-poor economy with near-pristine forest. We assume both groups in the same country to share equal wealth. For the extension model with carbon market, we consider a set of initial conditions in which the North is relatively capital-rich and forest-poor and the South relatively capital-poor and forest-rich ($F_{N0} = 0.70$, $F_{S0} = 0.95$, $a_{N0} = 0.90$, $a_{S0} = 0.30$). We will see how forest stock evolves along with capital accumulation on the development path, as well as how incentives of deforestation and reforestation change over time.

6.4 Quantitative Results

6.4.1 Benchmark Model

We begin our quantitative results with the equilibrium dynamics in the benchmark model. Figure 11 traces the dynamic path of forest stock in the social optimum and the decentralized equilibrium for an economy that begins forest-rich but capital-poor. Both scenarios follow a pronounced *U-shape*: the forest is first drawn down to a trough around years 10–12, and only thereafter recovers towards its long-run level. The mechanism is the dynamic interaction between the competing incentives to exploit and restore the forest at the heart of the model: when capital is scarce its marginal return is high, economic agents clear land to finance accumulation; deforestation is intense and reforestation muted in the early phase. As capital deepens, its return falls while the forest, now depleted, becomes scarce and its marginal amenity value rises. The sign of net forest use flips, and the economy transitions from a deforesting stage to a net-reforesting stage. The model thus generates a *forest Kuznets curve* endogenously: deforestation rises with early development, peaks, and gives way to recovery as income accumulates, with no exogenous change in policy or preferences driving the turning point.

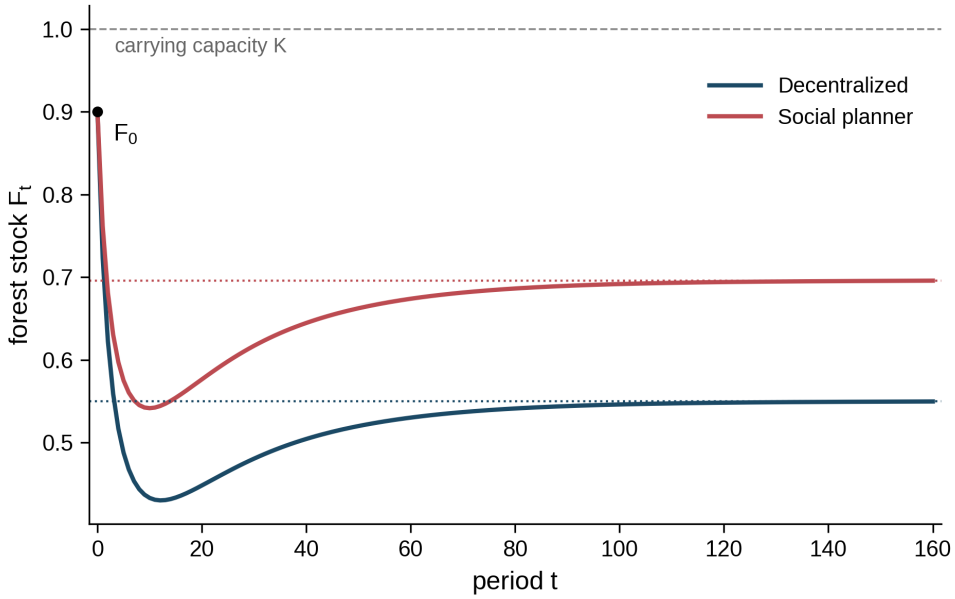


Figure 11. Forest stock along the equilibrium transition

Notes: The figure plots the forest-cover share F_t along the perfect-foresight transition under the decentralised Nash equilibrium (navy) and the constrained social planner (brick), starting from a forest-rich, capital-poor initial state ($F_0 = 0.90$, $a_{H0} = a_{L0} = 0.30$). Dotted horizontal lines mark the respective steady states ($F^*_{\text{DE}} = 0.55$, $F^*_{\text{SP}} = 0.70$); the dashed line is the carrying capacity $K = 1$. Parameter values are reported in Table 1.

The comparison between the two scenarios is one of this paper's central messages. The decentralized path overshoots far more violently: its forest trough falls to roughly 0.43, against 0.54 for the planner, and it converges to a permanently lower steady state. Privately optimizing economic agents act only on their own amenity valuation and ignore the forest benefit they confer on the other group; they therefore over-clear during the development phase and under-restore in the later stage. The constrained planner, internalizing the cross-group externality, both cushions the transitional overshoot and sustains a higher long-run stock: $F^*_{\text{SP}} = 0.70$ exceeds $F^*_{\text{DE}} = 0.55$ by about 0.15 in cover share, or 27 per cent. The vertical distance between the two curves is a direct visualization of the welfare cost of the externality, largest precisely during the deforestation phase when the forest is being drawn down most aggressively.

Appendix Figure B4 disaggregates these dynamics at the group level. The key finding is that the planner reverses the group conservation ranking relative to the decentralized outcome, loading the conservation burden onto the richer, more productive group L — the opposite of the laissez-faire allocation. The appendix also plots the time path of the optimal Pigouvian tax τ_{jt} (Proposition 4) that implements the social optimum.

We now turn to the two-country extension to explore how a global carbon market shapes the growth path of natural and physical capital.

6.4.2 The Role of Carbon Market

We discipline the carbon market by setting the cap relative to the decentralized laissez-faire path. In each period the cap on net global emissions is $\bar{E}_t = \bar{\kappa}\, E_t^{LF}$, where E_t^{LF} is the business-as-usual emission flow generated along the laissez-faire transition and $\bar{\kappa} \in (0, 1]$ indexes policy stringency. Tying the cap to the contemporaneous laissez-faire path, rather than to a fixed level, keeps the instrument economically meaningful throughout the transition, along which emissions are far higher in the early, capital-scarce development phase than at the steady state. Carbon rights are endowed equally across the four groups, and our baseline market sets $\bar{\kappa} = 0.70$.

We proceed in two steps. We first hold $\bar{\kappa}$ fixed and trace how the carbon market reshapes the transition paths of the forest stock and of productive capital in the two countries, against the decentralized and social-planner benchmarks. We then vary $\bar{\kappa}$ and ask how the steady-state forest stock, aggregate efficiency, and the distribution of welfare across countries and groups respond as the cap is tightened.

Forest dynamics. Figure 12 plots the forest transition in each country. Both economies reproduce the central dynamic of the benchmark model: a phase of aggregate deforestation followed by restoration, i.e., the *forest Kuznets curve*. The forest is drawn down while capital is scarce and its return high, and rebuilt as the economy matures and the marginal amenity value of the now-depleted forest rises. As in the benchmark, the social planner sustains a higher forest stock than laissez-faire in both countries, reflecting the failure of economic agents to internalize the positive externality that their conservation confers on others. The planner–laissez-faire gap is now somewhat wider than in the benchmark single economy case, because the decentralized economy fails to internalize two externalities rather than one: the local forest amenity and the global carbon externality.

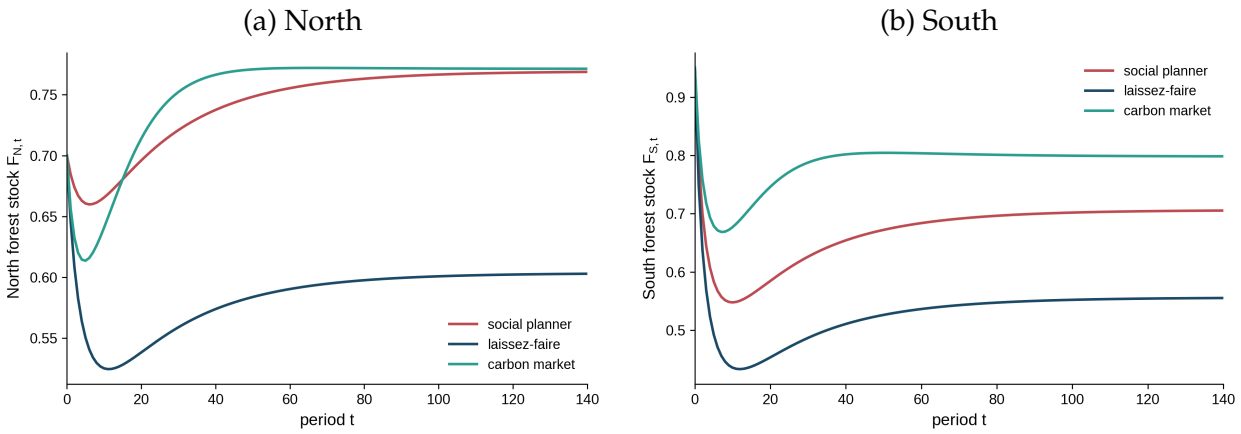


Figure 12. Forest Stock with Carbon Market

Notes: The figure plots the forest-cover share along the perfect-foresight transition under the global social planner (brick), laissez-faire (navy), and the decentralised economy with a global carbon market (teal). The carbon market caps net emissions at 70% of the contemporaneous laissez-faire (business-as-usual) path, with carbon rights endowed equally across groups and countries. The economy starts from a development-stage initial condition in which the North is relatively capital-rich and forest-poor and the South relatively capital-poor and forest-rich ($F_{N0} = 0.70$, $F_{S0} = 0.95$, $a_{N0} = 0.90$, $a_{S0} = 0.30$). Panel (a) reports the North; panel (b) the South. Parameters are as in Table 1.

The carbon market moves both economies substantially towards conservation. In the South the emission cap induces markedly more reforestation and less clearing than under laissez-faire; Southern cover rises above not only the laissez-faire path but the planner's as well, overshooting the social optimum throughout the transition. In the North, by contrast, forest cover lies below the planner during the early development phase, rises above it through the middle of the transition, and converges to approximately the planner's level in the long run. The asymmetry has a clean logic. A uniform carbon price equalizes the marginal cost of carbon abatement across countries and therefore concentrates sequestration where planting κ_n is cheapest, i.e., the South. The planner instead weighs each country's local amenity and converts the common carbon value into country-specific consumption units, and so does not wish abatement to be loaded so heavily onto the South. The market consequently over-conserves the low-cost South and roughly matches the high-cost North; strikingly, it *reverses* the planner's cross-country ordering of forest stock, holding the South more forested than the North when the planner would do the opposite.

Steady states and aggregate efficiency. Figure 13 varies the cap on carbon emission $\bar{\kappa}$ and reports steady-state outcomes; the horizontal axis runs from a slack cap on the left to a tight cap on the right. Panel (a) shows that tightening the cap monotonically raises the steady-state forest stock in both countries, and that the economy with carbon market always sustains more forest than laissez-faire. The response is also steeper for the South than for the North: as the low-

cost sequester, the South reforests more elastically as the carbon price rises, so its cover climbs faster and overtakes Northern cover at caps below roughly 80%. For sufficiently tight caps the market pushes forest cover above the planner's level in both countries, suggesting that tracking the planner's forest stock alone and maximizing aggregate welfare are distinct objectives.

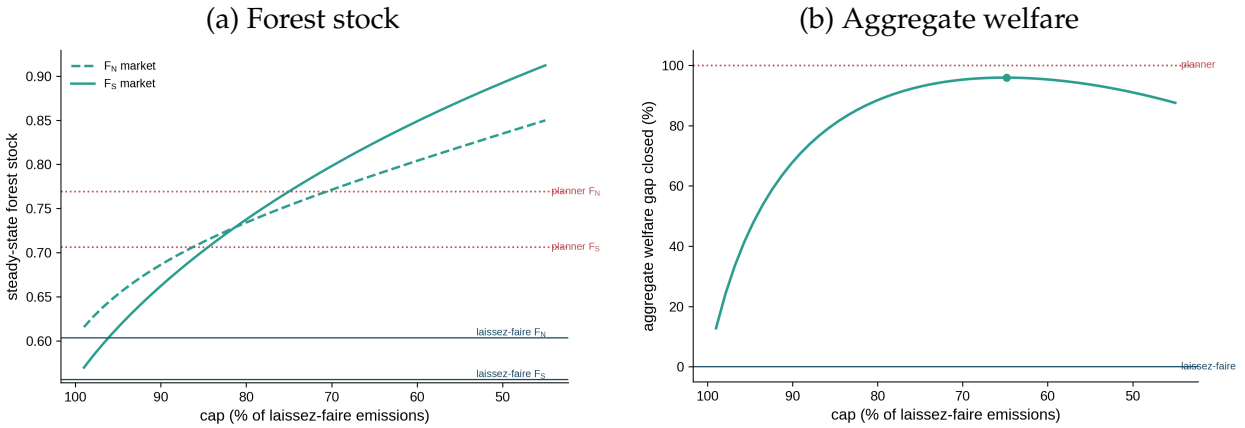


Figure 13. Steady States with Carbon Market

Notes: The figure traces steady-state outcomes as the emissions cap varies, expressed as a percentage of laissez-faire emissions (the horizontal axis runs from a slack cap on the left to a tight cap on the right). Panel (a) plots the steady-state forest stock in each country (teal: solid for the South F_S, dashed for the North F_N); brick dotted lines mark the social-planner levels and navy solid lines the laissez-faire levels. Panel (b) plots the share of the laissez-faire–to–planner aggregate welfare gap closed by the carbon market; 0% corresponds to laissez-faire and 100% to the global social optimum, and the marker indicates the welfare-maximising cap. Carbon rights are endowed equally across groups and countries. Parameters are as in Table 1.

Panel (b) plots the share of the laissez-faire–to–planner welfare gap closed by the market. The relationship is inverted-U-shaped. From a slack carbon emission cap $\bar{\kappa}$, tightening raises welfare by internalizing the carbon and amenity externalities, until the gap closed peaks at roughly 96% at a cap near 65% of laissez-faire emissions. However, tightening further lowers welfare, as the economy is forced to hold too much forest and to forgo the clearing, output and consumption that are optimal during development. It is interesting to note that the peak of closed welfare gap falls short of 100%: consistent with Proposition 1, a single carbon price cannot reproduce the planner's allocation, which embeds country- and group-specific amenity and consumption-equivalent wedges. Quantitatively, however, the shortfall is small. The best attainable carbon market closes most of the welfare gap, so that a well-designed market, though not first-best, is able to deliver the large majority of the available efficiency gains.

Distributional consequences. Figure 14 reports how each country's steady state aggregate welfare changes with the carbon emission cap $\bar{\kappa}$, relative to laissez-faire. Both countries gain at

every level of $\bar{\kappa}$, and the South gains somewhat more than the North. This speaks directly to the concern that forest carbon markets, and the reforestation programmes they finance, constitute a form of "carbon colonialism" in which the Global South surrenders its own development opportunities to manufacture credits for Northern emitters (Bachram, 2004). In our calibration the country accounts do not support that reading: the South is the net seller of carbon rights, and the revenue it earns leaves it strictly better off, indeed marginally more so than the North.

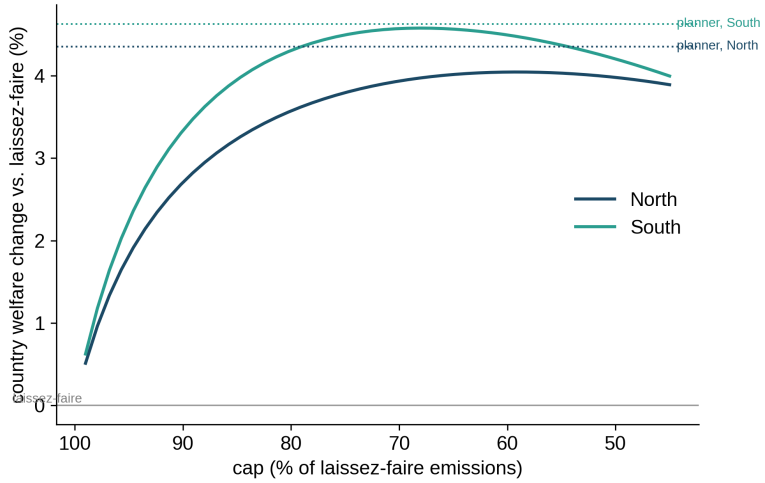
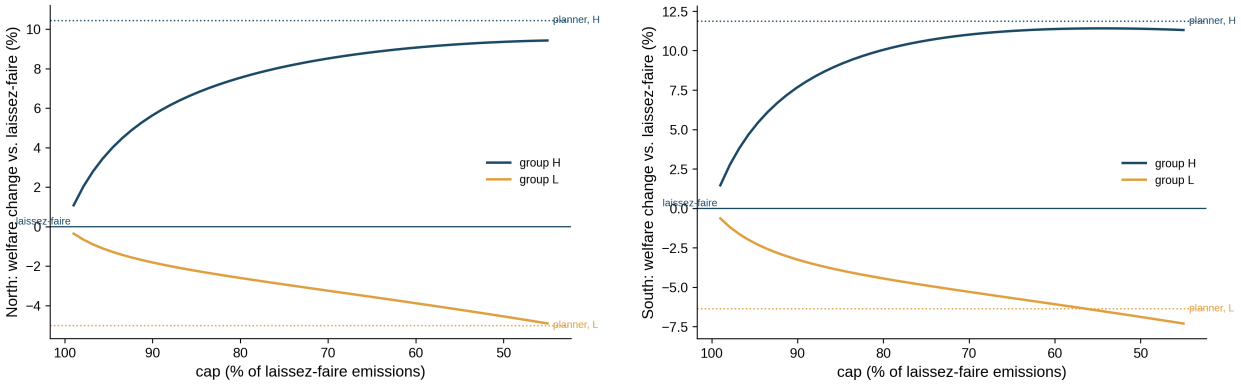


Figure 14. Country-Level Welfare Gains from the Carbon Market

Notes: The figure plots the steady-state change in each country's aggregate welfare relative to laissez-faire (in per cent) as the emissions cap varies; the horizontal axis runs from a slack cap on the left to a tight cap on the right. Dotted lines mark each country's welfare level under the global social planner. Carbon rights are endowed equally across groups and countries; the global carbon damage is split equally between the two countries in the country accounts. Parameters are as in Table 1.

The country aggregates, however, conceal a sharper pattern within each economy, shown in Figure 15, which plots each group's steady-state welfare change relative to laissez-faire as $\bar{\kappa}$ varies. The core message is that the carbon market is not Pareto-improving across groups. In both countries the high-valuation group H gains as the cap tightens, while the low-valuation, extractive group L is made worse off, its welfare falling below laissez-faire and converging towards the (also lower) planner level. The market, like the social optimum itself, thus acts as a redistribution from the high-deforesting group towards the conservationist group within each country. The mechanism is that pricing carbon and forest amenity taxes precisely the land-clearing in which group L has a comparative advantage; under an equal endowment of carbon rights, group L is the net buyer of permits and bears the cost, while group H is the net seller and, as Figure B5 indicates, channels part of the proceeds into faster capital accumulation, albeit modestly. The salient distributional consequence of introducing a global carbon market therefore runs not between the South

and the North, but between the extractive and the conservationist group within each hemisphere, orthogonal to the North-South divide emphasized by the carbon-colonialism critique.



(a) North (b) South

Figure 15. Distributional Impact of Carbon Market

Notes: The figure reports the steady-state welfare change of each group relative to laissez-faire (in per cent) as the emissions cap varies; positive values denote a welfare gain over laissez-faire and negative values a loss. In each panel the solid lines plot the high-valuation group H and the low-valuation group L, and the dotted lines mark each group's welfare level under the global social planner. Panel (a) reports the North; panel (b) the South. The horizontal axis runs from a slack cap on the left to a tight cap on the right; carbon rights are endowed equally across groups and countries. Parameters are as in Table 1.

7 Conclusion

Governments, NGOs, and private firms have committed hundreds of billions of dollars to planting trees as a climate solution. Yet until now, there has been no globally representative, causally identified evaluation of whether these programs actually work. This paper fills that gap.

We construct the Global Reforestation Atlas by combining georeferenced reforestation projects with satellite vegetation data, national forest policy records, cropland and mining maps, and household survey data. For each project, we compare vegetation outcomes to a synthetic control group of randomly drawn polygons in the same country, matching on pre-treatment vegetation trends.

The average program increases vegetation by 0.65 percent relative to controls, a modest gain that hides striking heterogeneity. Programs succeed in nine of the 25 Sub-Saharan African countries in our sample. Success is predicted by stronger national forest governance, lower pre-existing human activity within project sites, less surrounding agricultural pressure, and greater distance from mining operations. Where communities depend on livestock and farmland, planted trees face encroachment pressure that erodes any initial gains.

Programs are also not placed where they are most needed. Voluntary carbon market incentives direct investment toward wealthier, less agricultural locations where survival is easiest and carbon credits most likely to be verified. A welfare calculation finds that the average project generates only \$0.032 of climate benefit per dollar spent, far below the break-even threshold. Even the best-performing projects imply a cost of \$101.9 per tonne of CO_2 sequestered, well above engineering benchmarks that assume projects succeed.

A theoretical model shows that laissez-faire leads to 27 percent less forest than the social optimum due to two externalities: a within-country common-pool problem and a cross-country carbon externality. A global carbon market corrects the carbon channel but not the local amenity distortion, though international trade in carbon credits generates large efficiency gains by directing planting toward low-cost locations in the Global South.

The policy implications are clear. Reforestation works, but only where governance is strong, land pressure is low, and communities have a stake in forest survival. Current market mechanisms are directing billions of dollars to easy wins rather than where intervention is most needed. Independent satellite monitoring of the kind we implement here, linking verified vegetation outcomes to the projects themselves, offers a path toward evidence-based targeting and a better return on the world's reforestation investment.

References

Abadie, A. and Imbens, G. W. (2006). Large sample properties of matching estimators for average treatment effects. *econometrica*, 74(1):235–267.

Acemoglu, D., Aghion, P., Bursztyn, L., and Hemous, D. (2012). The environment and directed technical change. *American economic review*, 102(1):131–166.

Acevedo, M. A., Marcano, M., and Fletcher, R. J. (2012). A diffusive logistic growth model to describe forest recovery. *Ecological Modelling*, 244:13–19.

Aragón, F. M. and Rud, J. P. (2013). Natural resources and local communities: evidence from a peruvian gold mine. *American Economic Journal: Economic Policy*, 5(2):1–25.

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., and Wager, S. (2021). Synthetic difference-in-differences. *American economic review*, 111(12):4088–4118.

Assunção, J., Hansen, L. P., Munson, T., and Scheinkman, J. A. (2024). Carbon prices, forest conservation and reforestation in the Brazilian Amazon. *BIS Working Paper*, (1223).

Avitabile, V., Herold, M., Heuvelink, G. B., Lewis, S. L., Phillips, O. L., Asner, G. P., Armston, J., Ashton, P. S., Banin, L., Bayol, N., et al. (2016). An integrated pan-tropical biomass map using multiple reference datasets. *Global change biology*, 22(4):1406–1420.

Bachram, H. (2004). Climate fraud and carbon colonialism: The new trade in greenhouse gases. *Capitalism Nature Socialism*, 15(4):5–20.

Balboni, C., Berman, A., Burgess, R., and Olken, B. A. (2023). The economics of tropical deforestation. *Annual Review of Economics*, 15:723–754.

Barrage, L. and Nordhaus, W. (2024). Policies, projections, and the social cost of carbon: Results from the dice-2023 model. *Proceedings of the National Academy of Sciences*, 121(13):e2312030121.

Basri, M. C., Felix, M., Hanna, R., and Olken, B. A. (2021). Tax administration versus tax rates: Evidence from corporate taxation in indonesia. *American Economic Review*, 111(12):3827–3871.

Burgess, R., Hansen, M., Olken, B. A., Potapov, P., and Sieber, S. (2012). The political economy of deforestation in the tropics. *The Quarterly journal of economics*, 127(4):1707–1754.

Busch, J., Bukoski, J. J., Cook-Patton, S. C., Griscom, B., Kaczan, D., Potts, M. D., Yi, Y., and Vincent, J. R. (2024). Cost-effectiveness of natural forest regeneration and plantations for climate mitigation. *Nature Climate Change*, 14(9):996–1002.

Calel, R., Colmer, J., Dechezleprêtre, A., and Glachant, M. (2025). Do carbon offsets offset carbon? *American Economic Journal: Applied Economics*, 17(1):1–40.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of econometrics*, 225(2):200–230.

Cook-Patton, S. C. et al. (2020). Mapping carbon accumulation potential from global natural forest regrowth. *Nature*, 585:545–550.

Cooley, T. F. and Prescott, E. C. (1995). Economic growth and business cycles. In Cooley, T. F., editor, *Frontiers of Business Cycle Research*, pages 1–38. Princeton University Press, Princeton, NJ.

De Chaisemartin, C. and d'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American economic review*, 110(9):2964–2996.

Donaldson, D. and Storeygard, A. (2016). The view from above: Applications of satellite data in economics. *Journal of Economic Perspectives*, 30(4):171–198.

Ecosystem Marketplace (2025). State of the voluntary carbon market 2025.

Eggleston, H., Buendia, L., Miwa, K., Ngara, T., and Tanabe, K. (2006). 2006 ipcc guidelines for national greenhouse gas inventories.

Ferraro, P. J. and Pattanayak, S. K. (2006). Money for nothing? a call for empirical evaluation of biodiversity conservation investments. *PLoS biology*, 4(4):e105.

Foster, A. D. and Rosenzweig, M. R. (2003). Economic growth and the rise of forests. *The Quarterly Journal of Economics*, 118(2):601–637.

Giljum, S., Maus, V., Kuschnig, N., Luckeneder, S., Tost, M., Sonter, L. J., and Bebbington, A. J. (2022). A pantropical assessment of deforestation caused by industrial mining. *Proceedings of the National Academy of Sciences*, 119(38):e2118273119.

Goldstein, M. and Udry, C. (2008). The profits of power: Land rights and agricultural investment in ghana. *Journal of political Economy*, 116(6):981–1022.

Gomme, P. and Rupert, P. (2007). Theory, measurement and calibration of macroeconomic models. *Journal of Monetary Economics*, 54(2):460–497.

Hainmueller, J. (2012). Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political analysis*, 20(1):25–46.

Hansen, L. P. and Scheinkman, J. A. (2025). Emission prices, biomass, and biodiversity in tropical forests. *AEA Papers and Proceedings*, 115:391–396.

Hansen, M. C., Potapov, P. V., Moore, R., Hancher, M., Turubanova, S. A., Tyukavina, A., Thau, D., Stehman, S. V., Goetz, S. J., Loveland, T. R., et al. (2013). High-resolution global maps of 21st-century forest cover change. *science*, 342(6160):850–853.

Hayashi, F. (1982). Tobin's marginal q and average q: A neoclassical interpretation. *Econometrica*, 50(1):213–224.

Henderson, J. V., Storeygard, A., and Weil, D. N. (2012). Measuring economic growth from outer space. *American economic review*, 102(2):994–1028.

Imbens, G. W. (2015). Matching methods in practice: Three examples. *Journal of Human Resources*, 50(2):373–419.

Jayachandran, S. (2013). Liquidity constraints and deforestation: The limitations of payments for ecosystem services. *American Economic Review*, 103(3):309–313.

Jayachandran, S., De Laat, J., Lambin, E. F., Stanton, C. Y., Audy, R., and Thomas, N. E. (2017). Cash for carbon: A randomized trial of payments for ecosystem services to reduce deforestation. *Science*, 357(6348):267–273.

John, A., Allotey, S., Koebe, T., Tyukavina, A., and Weber, I. (2025). A global dataset of location data integrity-assessed reforestation efforts. *Scientific Data*, 12(1):1714.

King, R. G. and Rebelo, S. T. (1999). Resuscitating real business cycles. In Taylor, J. B. and Woodford, M., editors, *Handbook of Macroeconomics*, volume 1, chapter 14, pages 927–1007. Elsevier.

Kleven, H., Landais, C., and Leite-Mariante, G. (2025). The child penalty atlas. *Review of Economic Studies*, 92(5):3174–3207.

Lucas, R. E. (1978). On the size distribution of business firms. *Bell Journal of Economics*, 9(2):508–523.

Maus, V., Giljum, S., Gutschlhofer, J., da Silva, D. M., Probst, M., Gass, S. L., Luckeneder, S., Lieber, M., and McCallum, I. (2020). A global-scale data set of mining areas. *Scientific data*, 7(1):289.

Moore, F. C., Drupp, M. A., Rising, J., Dietz, S., Rudik, I., and Wagner, G. (2024). Synthesis of evidence yields high social cost of carbon due to structural model variation and uncertainties. *Proceedings of the National Academy of Sciences*, 121(52):e2410733121.

MSCI (2026). Carbon markets year in review 2025.

Nordhaus, W. D. (2017). Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences*, 114(7):1518–1523.

Potapov, P., Hansen, M. C., Pickens, A., Hernandez-Serna, A., Tyukavina, A., Turubanova, S., Zalles, V., Li, X., Khan, A., Stolle, F., et al. (2022). The global 2000-2020 land cover and land use change dataset derived from the landsat archive: first results. *Frontiers in Remote Sensing*, 3:856903.

Poudel, A., Shrestha, H. L., Mahat, N., Sharma, G., Aryal, S., Kalakheti, R., and Lamsal, B. (2023). Modeling and mapping of aboveground biomass and carbon stock using sentinel-2 imagery in chure region, nepal. *International Journal of Forestry Research*, 2023(1):5553957.

Probst, B. S., Toetzke, M., Kontoleon, A., Anádón, L. D., Minx, J. C., Haya, B. K., Schneider, L., Trotter, P. A., West, T. A. P., Gill-Wiehl, A., and Hoffmann, V. H. (2024). Systematic assessment of the achieved emission reductions of carbon crediting projects. *Nature Communications*, 15:9562.

Rennert, K., Errickson, F., Prest, B. C., Rennels, L., Newell, R. G., Pizer, W., Kingdon, C., Wingenroth, J., Cooke, R., Parthum, B., et al. (2022). Comprehensive evidence implies a higher social cost of co2. *Nature*, 610(7933):687–692.

Rodemeier, M. (2026). Willingness to pay for carbon mitigation: Field evidence from the market for carbon offsets. *The Review of Financial Studies*, 39(1):1–29.

Spawn, S. A., Sullivan, C. C., Lark, T. J., and Gibbs, H. K. (2020). Harmonized global maps of above and belowground biomass carbon density in the year 2010. *Scientific data*, 7(1):112.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of econometrics*, 225(2):175–199.

West, T. A., Börner, J., Sills, E. O., and Kontoleon, A. (2020). Overstated carbon emission reductions from voluntary redd+ projects in the brazilian amazon. *Proceedings of the National Academy of Sciences*, 117(39):24188–24194.

West, T. A. P., Wunder, S., Sills, E. O., Börner, J., Rifai, S. W., Neidermeier, A. N., Frey, G. P., and Kontoleon, A. (2023). Action needed to make carbon offsets from tropical forest conservation work for climate change mitigation. *Science*, 381(6660):873–877.

World Bank (2025). *State and Trends of Carbon Pricing 2025*. World Bank Group, Washington, DC.

A Data Appendix

A.1 John et al. (2025) Database

Our analysis builds on the global reforestation dataset compiled by John et al. (2025), which consolidates information on 1,289,068 planting sites from 45,628 projects across more than 50 sources. The full dataset contains 91 variables for each site, including identifiers, geographic boundaries, reported planting information, satellite-derived vegetation indices, land-cover and infrastructure indicators, and a set of ten location data integrity flags. We use the Sub-Saharan African subset of this dataset, which yields 1,964 project sites across 25 countries (see Section 2.1).

Variable selection. From the full set of attributes, we extract the following for each site: a unique polygon identifier (`poly_id`, corresponding to `site_id_created` in the original dataset), the reported planting year (`plant_yr`, derived from `planting_date_reported`), the country name (`country`), and administrative and project identifiers (`ctry_id`, `proj_id`, `site_id`, `site_rpt`). Table A1 lists the key original variables and their definitions.

Table A1. Key variables from John et al. (2025) used in our analysis

Variable name	Description
`site_id_created`	Unique numeric identifier for each planting site (our `poly_id`)
`project_id_reported`	Original project identifier
`host_name`	Organization hosting the project information
`country`	Country of the site
`planting_date_reported`	Year of planting (reported)
`site_sqkm_derived`	Site area in km^2 (derived)
`total_road_length_km`	Road length within site (km)
`built_area_2018`	Built-up area within site (km^2)
`treecover_atplanting`	Tree cover at planting time (fraction)
`forest_at_planting_glad`	GLAD forest cover at planting (fraction)
`circularity`	Circularity of site boundary
`polygon_acircle_oval_95`	Indicator for (almost) perfect circle
`project_geometries_invalid`	Indicator for invalid geometry
`exact_admin_area`	Indicator for exact match with administrative boundary
`Nesting Polygon`	Degree of nesting with other polygons
`Intersecting Polygon`	Degree of intersection with other polygons
`other_landcover_score`	Other land cover score
`stable_cropland_score`	Stable cropland score
`Road_presence`	Road presence indicator (LDIS component)
`Built_area_presence`	Built-up area presence indicator (LDIS component)

Handling of missing start years. The dataset contains `planting_date_reported` for 92.4% of all sites; the remaining 7.6% have a missing entry. Because the definition of the reported date varies across hosts (e.g., project registration, intervention year, or actual planting), we supplement this field with the LLM-extracted `planting_date_derived` where available. Sites for which no reliable planting year can be assigned are excluded from our sample. In the Sub-Saharan Africa subsample, a small number of sites are dropped for this reason.

LDIS filter. The Location Data Integrity Score (LDIS) is a composite index ranging from 0 to 10, constructed as the sum of ten binary indicators that flag potential concerns with a site's geographic

boundary. The ten components are: road presence inside the polygon, built-up area presence, forest cover at planting, other land cover, nesting with other polygons, intersection with other polygons, exact match with an administrative boundary, perfect circularity, invalid geometry, and stable cropland presence. In the full dataset, 21% of sites with polygon geometries attain a perfect LDIS of 10. In our main specification we do not filter on LDIS, preserving sample size and external validity. However, we conduct robustness checks restricting the sample to sites with LDIS ≤ 2 (i.e., those failing at most two of the ten checks), and results are quantitatively similar.

Spatial deduplication. The original dataset provides several fields to identify potential duplicates: `Nested_in`, `Intersecting_with`, and the continuous scores `Nesting Polygon` and `Intersecting Polygon`. These are based on a 95% area-overlap criterion: two polygons are considered duplicates if they share more than 95% of their respective areas. We adopt the same rule to remove duplicates: when two polygons overlap by more than 95%, we retain the one with the earliest reported planting date, or, if dates are identical, the one with the most complete set of attributes. Polygons that are nested within larger ones are retained only if the larger polygon is not already in our sample. After deduplication, we verify that no treated polygons overlap within the same country. The raw, un-deduplicated dataset is retained for robustness checks.

A.2 Google Earth Engine Extraction Pipeline

A.2.1 Data: Satellite-Based Forest Cover and NDVI Measures

We construct our outcome and control variables using satellite imagery processed in Google Earth Engine (GEE). The main variables are annual Normalized Difference Vegetation Index (NDVI) and several forest cover indicators derived from the Hansen Global Forest Change dataset. All variables are extracted at the polygon level—both for the actual treated project polygons and for a set of randomly assigned control polygons within the same country.

1. Polygon Data and Control Assignment

Our analysis uses polygons that delineate the boundaries of specific forestation projects across multiple African countries. Each polygon represents a reforestation site which means it is treated. These treated polygons were uploaded to GEE as vector assets after splitting large files into manageable batches. For each treated polygon, we generate a set of 100 control polygons to serve as a counterfactual. The control polygons are constructed as follows:

i. The area of the treated polygon is computed.

ii. A buffer radius is derived so that a circle of that radius has the same area.

iii. Using the country boundary (from the Large Scale International Boundaries dataset), 100 random points are generated within the country but excluding a 1 km buffer around the treated polygon. This buffer avoids spatial spillover.

iv. A circular polygon of the same area is created around each random point, producing a control unit with the same size as the treated one.

Each control polygon inherits the project attributes (country, planting year, etc.) and is linked back to its specific treated polygon via a treated_polygon_id. This design ensures that control polygons are drawn from the same national context and have a comparable scale, but are not directly adjacent to the treatment site.

2. Satellite-Derived Variables

For every treated and controlled polygon, we extract satellite-based variables for all years from 2000 to 2024.

Variable	Description	Temporal cove
landsat_ndvi_mean	Mean Landsat NDVI within polygon	2000–2024, anr
landsat_ndvi_median	Median Landsat NDVI	2000–2024
landsat_ndvi_sd	Within-polygon standard deviation of Landsat NDVI	2000–2024
modis_ndvi_mean	Mean MODIS NDVI	2000–2024
modis_ndvi_median	Median MODIS NDVI	2000–2024
modis_ndvi_sd	Within-polygon standard deviation of MODIS NDVI	2000–2024
hansen_treecover_mean	Mean tree cover in 2000 (%)	Cross-section (
hansen_treecover_median	Median tree cover in 2000	Cross-section
hansen_lossyear_mean	Mean of encoded loss year	Cross-section
hansen_gain_mean	Mean forest gain indicator (0/1)	Cross-section (

Hansen Forest Cover Variables. The Hansen Global Forest Change dataset (version 1.12) provides 30 m resolution maps from which we extract three metrics: tree cover in 2000 (treecover2000), the year of detected forest loss (lossyear, encoded as 0 for no loss and 1–24 for loss in 2001–2024), and a binary indicator of forest gain between 2000 and 2012 (gain). We compute the mean, median, and standard deviation of each metric within each polygon. Because treecover2000 and gain are static, their year-to-year variation comes only from changes in the set of polygons considered; lossyear is later converted into annual binary loss indicators in the panel construction.

NDVI Measures. We use two independent sources of NDVI to ensure robustness to different spatial and temporal resolutions. The Landsat annual NDVI composites (LANDSAT/COMPOSITES/C02/T1_ provide gap-filled, 30 m resolution annual NDVI maps, from which we extract the mean, median, and standard deviation for each polygon and year. The MODIS MOD13A1 product delivers NDVI at 500 m resolution every 16 days; we compute pixel-wise annual means, apply the official scale factor (0.0001), and then extract the same summary statistics. Together, these two sources provide complementary measures of vegetation greenness at different spatial and temporal scales.

3. Batch Processing and Export

Extracting these variables for the entire set of polygons in one operation is computationally expensive and often exceeds GEE memory limits. We therefore split the extraction into inner batches. For each country, the full list of treated polygons is divided into subsets of 15 polygons. For each subset, we generate the corresponding 1,500 control polygons and immediately run the variable extraction. This batched approach prevents timeouts and allows easy local reassembly of the final panel dataset. The exported data files are in long format: each row corresponds to a polygon–year–variable combination, with columns for the mean, median, and standard deviation of the satellite measure, plus all identifying information (polygon ID, country, planting year, etc.). The data are subsequently merged, cleaned, and reshaped to a wide panel format for econometric analysis.

4. Resulting Variables

The final panel includes the satellite-based variables summarized in Table **??**. This dataset forms the basis for our difference-in-differences and event-study specifications, enabling us to track changes in vegetation greenness and forest cover around the project planting year.

A.3 Cropland and Mining Data Sources

The cropland data are from Potapov et al. (2022), available at https://zenodo.org/records/12505548. The mining data are from Maus et al. (2020), available at https://www.nature.com/articles/s41597-020-00624-w. For each project polygon, we compute the share of cropland and an indi-

cator for any active mine within a 50 km buffer, using standard spatial joins in Python with the `geopandas` library.

B Additional Results

B.1 Identification Methods

B.1.1 Synthetic Difference-in-Differences

Let Y_{it} denote the annual average NDVI for polygon i in year t, where i indexes both treated project sites and their matched control polygons, and T_i denotes the year in which project i was launched. We estimate the average treatment effect on the treated (ATT) using the synthetic difference-in-differences (SDID) estimator of Arkhangelsky et al. (2021).

The SDID estimator proceeds in two steps. First, it constructs unit weights $\hat{\omega}_i$ for the control polygons by solving:

$$\min_{\omega} \sum_{t<T_i} \left(Y_{it}^{\text{treated}} - \sum_{j \in \mathcal{C}} \omega_j Y_{jt} \right)^2 \quad \text{s.t.} \quad \omega_j \geq 0, \quad \sum_j \omega_j = 1, \tag{24}$$

where $\mathcal{C}$ denotes the set of control polygons. This reweights the control group to match the pre-treatment NDVI trajectory of the treated polygon, reducing confounding from differential pre-trends. Second, it applies a two-way fixed effects DiD estimator to the reweighted sample, yielding:

$$\hat{\tau}^{\text{SDID}} = \arg\min_{\tau, \alpha_i, \beta_t} \sum_{i,t} \hat{\omega}_i \left(Y_{it} - \alpha_i - \beta_t - \tau D_{it} \right)^2, \tag{25}$$

where $D_{it} = \mathbf{1}[i \text{ treated}, t \geq T_i]$ is the treatment indicator, α_i are polygon fixed effects, and β_t are year fixed effects. The polygon fixed effects absorb any time-invariant differences between treated and control sites; the year fixed effects absorb common shocks such as ENSO-driven rainfall variation.

The identifying assumption underlying SDID is that, conditional on the pre-treatment NDVI trajectory and the synthetic control weights, the counterfactual vegetation trend for treated polygons would have paralleled the (reweighted) control group in the absence of treatment. This is a weaker assumption than standard DiD parallel trends, because the reweighting step explicitly adjusts for differences in pre-treatment trajectories. It is stronger than synthetic control in that it also uses within-polygon variation across time.

B.1.2 Event Study Specification

To document the dynamic treatment effects and test for pre-trends, we estimate the event study specification:

$$Y_{it} = \alpha_i + \beta_t + \sum_{k \neq -1} \delta_k \cdot \mathbf{1}[t - T_i = k] + \varepsilon_{it}, \tag{26}$$

where k indexes years relative to treatment, and $k = -1$ (the year before treatment) is the omitted reference period. The pre-treatment coefficients $\{\delta_k\}_{k<0}$ provide a direct test of the parallel trends assumption: if treated and control sites were on the same NDVI trajectory before the project launched, these coefficients should be statistically indistinguishable from zero. The post-treatment coefficients $\{\delta_k\}_{k \geq 0}$ trace the dynamic treatment effect.

We estimate equation (26) separately for each country in our sample, producing the country-level event study figures that constitute the Reforestation Atlas. We also pool all countries and estimate a single average event study profile, weighting by the number of treated polygons per country.

B.2 Forest and Structural Transformation

(a) Agricultural Share of GDP (b) GDP Growth Rate

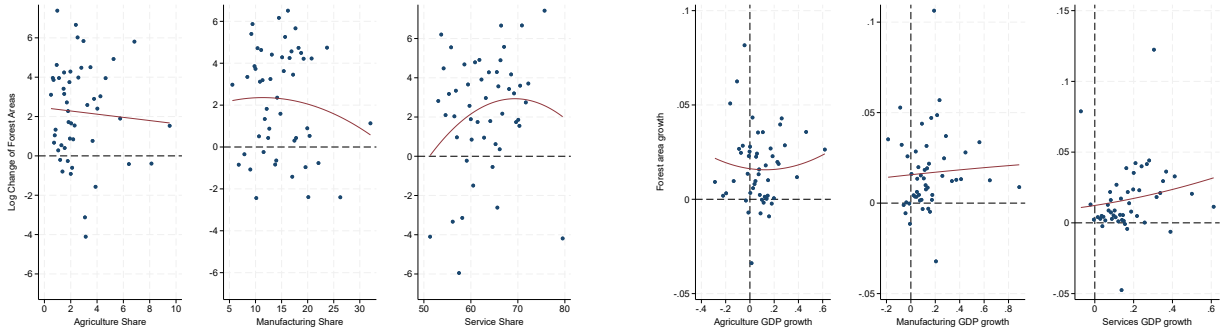


Figure B1. Forest Cover and Economic Development: OECD Countries

Notes: Each panel is a combined three-panel binscatter plot (50 quantile bins) with a quadratic fit, for OECD countries. Panel (a) plots the log change in forest area (inverse hyperbolic sine of the absolute change) against the agriculture, manufacturing, and services shares of GDP (%) respectively. Panel (b) plots forest area growth (log difference) against sector-specific GDP growth rates (log difference) for agriculture, manufacturing, and services; growth rates are winsorized at the 1st and 99th percentiles. The sample covers benchmark years 1990, 2000, 2010, 2015, 2020, and 2024. Forest area data are from the FAO Global Forest Resources Assessment 2025. GDP data are from the World Bank World Development Indicators.

(a) Agricultural Share of GDP (b) GDP Growth Rate

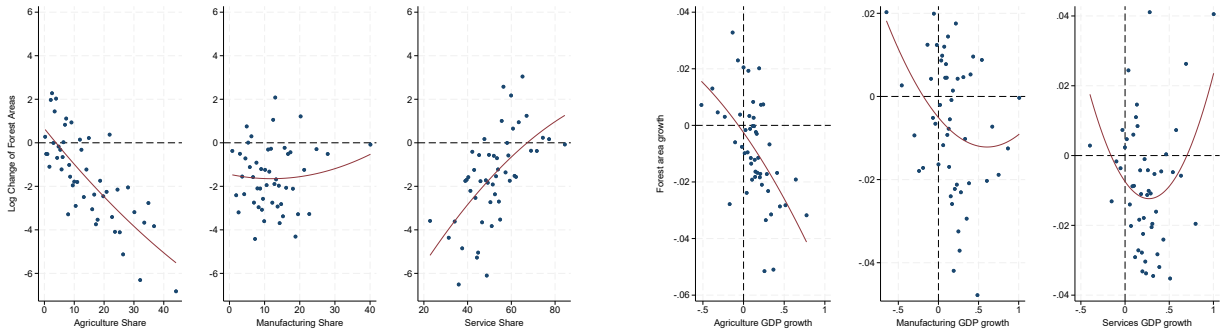


Figure B2. Forest Cover and Economic Development: Non-OECD Countries

Notes: Each panel is a combined three-panel binscatter plot (50 quantile bins) with a quadratic fit, for Non-OECD countries. Panel (a) plots the log change in forest area (inverse hyperbolic sine of the absolute change) against the agriculture, manufacturing, and services shares of GDP (%) respectively. Panel (b) plots forest area growth (log difference) against sector-specific GDP growth rates (log difference) for agriculture, manufacturing, and services; growth rates are winsorized at the 1st and 99th percentiles. The sample covers benchmark years 1990, 2000, 2010, 2015, 2020, and 2024. Forest area data are from the FAO Global Forest Resources Assessment 2025. GDP data are from the World Bank World Development Indicators.

B.3 Additional Figures for Reforestation Atlas

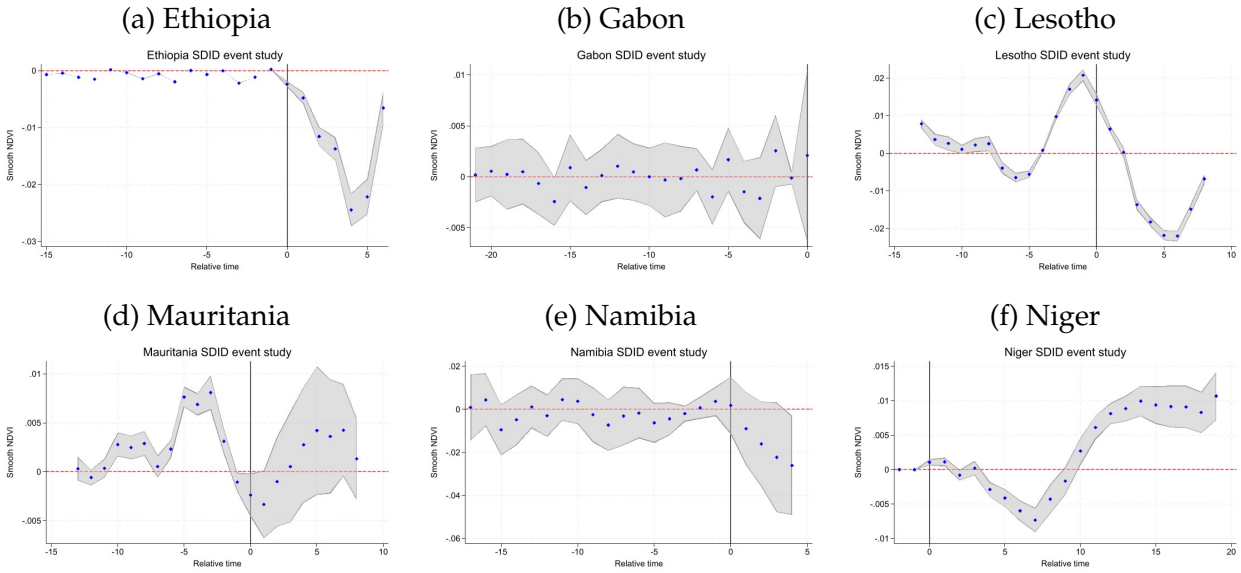


Figure B3. SDID Event Study: NDVI Effects — Failure/Uninformative

Notes: Each panel plots the SDID event study coefficients for one country, estimated using `sdid_event` with placebo-based inference. The outcome variable is smoothed mean NDVI (`smooth_mean`). The unit of observation is a site (`unique_id`), defined as one treated polygon matched to approximately 100 control units drawn from GEE satellite extracts. The country-level estimate pools all sites within the country; Ethiopia is excluded from this analysis. The x-axis shows years relative to project launch (treatment year); the reference period is $t = 0$ (the first treated year, with the vertical line at $x = 0$). The outer shaded band represents 95% confidence intervals; the inner band represents 90% confidence intervals (computed as $\hat{\beta} \pm 1.645 \times \hat{\sigma}$). Pre-treatment coefficients (years < 0) test the parallel trends assumption; post-treatment coefficients (years ≥ 0) trace the dynamic treatment effect. The sample is restricted to years after 2002. NDVI data are from Google Earth Engine satellite extracts; project treatment years are from the ATLAS afforestation project database.

B.4 Additional Figures for Model Quantifications

Capital accumulation. Figure B5 reports the dynamic path of physical capital for each group in each country. Two features stand out. First, the differences across regimes are modest in level and operate group by group, because the steady state condition $A_{nj} q_a(d_{nj}, a_{nj}) = 1/\beta - 1 + \delta$ pins the long-run marginal product of capital; the regimes can move capital only through the optimal forest and capital input mix. Second, the regime ordering differs across groups. For the high-valuation group H the planner accumulates slightly more capital than laissez-faire, whereas for the low-valuation, extractive group L the ordering reverses, mirroring the path of capital accumulation in the benchmark model. More importantly, we see that for high-valuation group, carbon market speeds up capital accumulation, whereas for low-valuation group, it slows down capital accumulation especially in the initial periods, with the effects being more pronounced for South country. Intuitively, under an equal endowment of carbon rights, the conservationist high-valuation group H is a net seller of credits in both countries, while the extractive group L is a net buyer, so credit revenue relaxes H's budget and finances faster accumulation while the cost of buying permits slows L's. This is further reinforced by factor reallocation: pricing carbon raises the cost of clearing, which group L curtails sharply, and since the marginal product of capital is now lower, the smaller land input drags down capital.

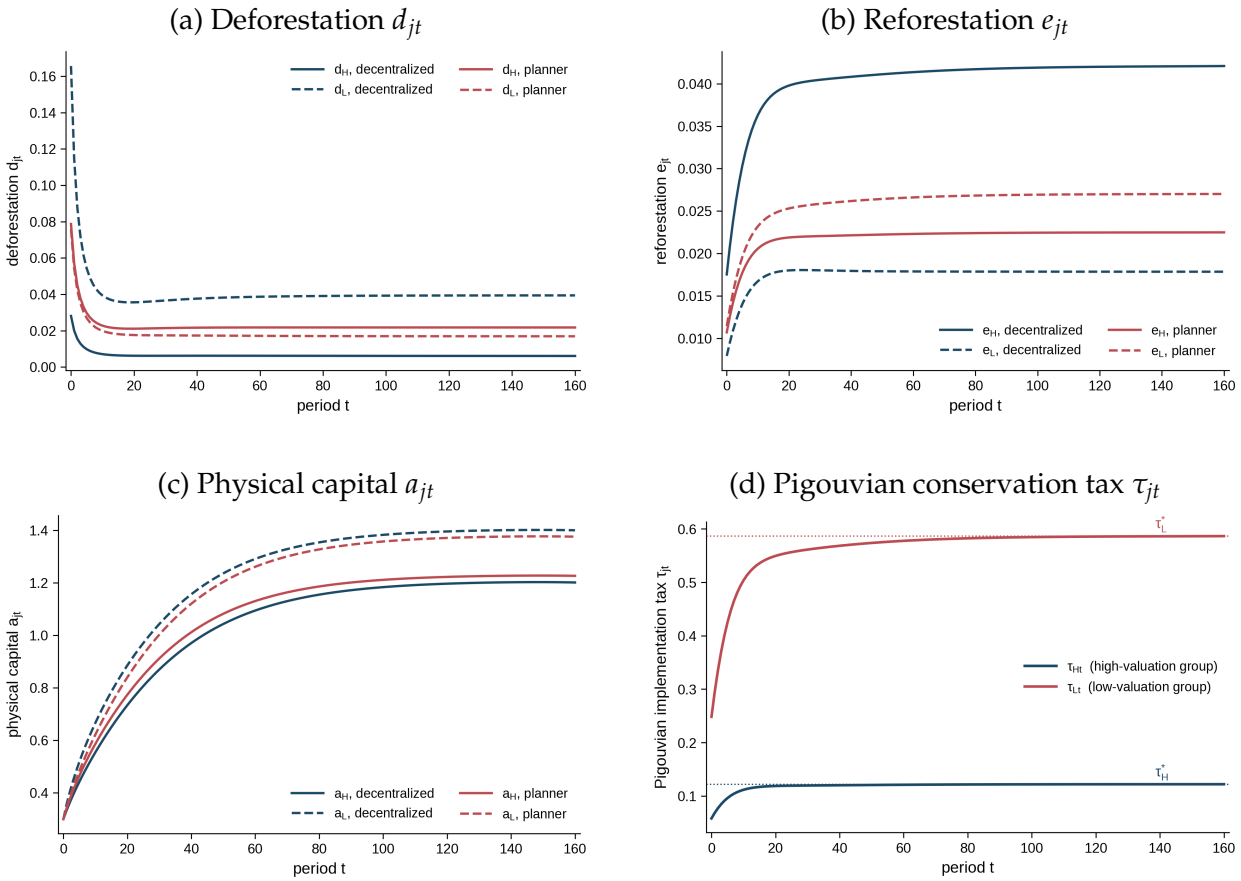


Figure B4. Group-level dynamics and the corrective tax

Notes: Panels (A)–(C) plot group-level allocations along the transition; navy denotes the decentralized equilibrium and brick the social planner, with solid lines for the high-valuation group H and dashed lines for the low-valuation group L. Panel (D) plots the time path of the optimal Pigouvian implementation tax τ_{jt} that decentralizes the planner's allocation, with navy for τ_{Ht} and brick for τ_{Lt}; dotted lines mark steady-state values. The economy starts from $F_0 = 0.90$, $a_{H0} = a_{L0} = 0.30$; parameters are as in Table 1. The two groups share the common forest but differ in their amenity valuation ($\phi_H > \phi_L$) and in extractive productivity ($A_L > A_H$).

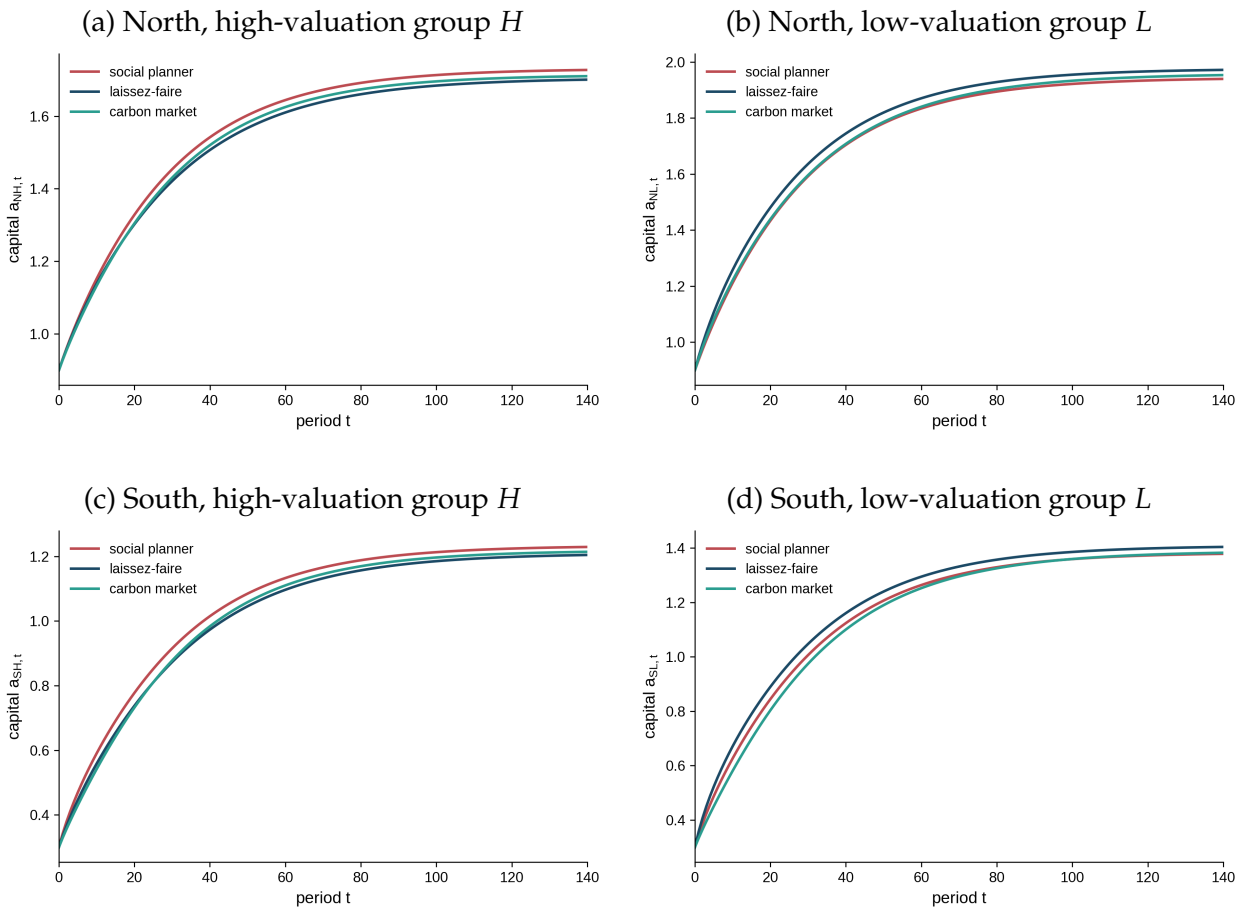


Figure B5. Capital Accumulation with Carbon Market

Notes: The figure plots physical capital along the perfect-foresight transition for each group in each country under the global social planner (brick), laissez-faire (navy), and the decentralized economy with a global carbon market (teal). Panels (a) and (b) report the high- and low-valuation groups in the North; panels (c) and (d) the corresponding groups in the South. The initial condition and carbon-market cap are as in Figure 12. Parameters are as in Table 1.

C Model Appendix

C.1 Single-Economy Benchmark

This subsection formally characterizes the single-economy case introduced in Section 6, isolating the within-country common-pool externality that motivates the two-country extension. Technical derivations are in Sections C.2 and C.3.

C.1.1 Decentralized Equilibrium

Proposition 2 (Dynamic decentralized equilibrium and steady state) *Given initial conditions* (F_0, a_{H0}, a_{L0}), *an interior decentralized equilibrium consists of sequences* $\{c_{jt}, d_{jt}, e_{jt}, a_{j,t+1}, \lambda_{jt}, \mu_{jt}\}_{t=0}^{\infty}$ *for* $j \in \{H, L\}$, *and a forest-stock path* $\{F_{t+1}\}_{t=0}^{\infty}$, *such that each group maximizes its objective taking the other group's choices as given. The equilibrium is characterized by the Euler equation* (4), *the forest-management condition* (5), *the shadow-price equations* (6) *and* (7), *the asset-accumulation equation* (2), *the forest law of motion* (3), *and the transversality condition*

$$\lim_{T\to\infty} \beta^T \lambda_{jT} a_{j,T+1} = 0. \tag{27}$$

The interior steady state $\left(c_j^*, d_j^*, e_j^*, a_j^*, \lambda_j^*, \mu_j^*, F^*\right)$ *for* $j \in \{H, L\}$ *is characterized by*

$$A_j\, q_a\left(d_j^*, a_j^*\right) = \frac{1}{\beta} - (1-\delta),$$

together with

$$\mu_j^* = \frac{\beta\phi_j v'(F^*)}{1-\beta G'(F^*)}, \qquad A_j\, q_d\left(d_j^*, a_j^*\right) = \frac{1}{\zeta} m'(e_j^*) = \frac{\mu_j^*}{u'(c_j^*)},$$

the forest law of motion $F^* = G(F^*) + \sum_{j\in\{H,L\}} \left(\zeta e_j^* - d_j^*\right)$, *and the stationary budget constraint* $c_j^* = A_j q\left(d_j^*, a_j^*\right) - m(e_j^*) - \delta a_j^*$.

Proposition 2 summarizes two margins. The first is the intertemporal saving margin: the steady-state capital stock satisfies the golden rule $A_j q_a(d_j^*, a_j^*) = 1/\beta - 1 + \delta$, so the long-run asset position is pinned down by fundamentals and initial wealth a_{j0} shapes only the transition. The second is the forest-management margin: each group equates the private return to deforestation and the private cost of reforestation with its own shadow value of the forest. The key object is the consumption-equivalent shadow value in steady state,

$$s_j^* \equiv \frac{\mu_j^*}{\lambda_j^*} = \frac{1}{u'(c_j^*)} \frac{\beta\phi_j v'(F^*)}{1-\beta G'(F^*)}.$$

The high-ϕ group tends to be more conservation-oriented, but the exact ranking of deforestation also depends on long-run capital and extraction productivity, since q_d depends on a_j^*.

C.1.2 Social Planner's Problem

The constrained social planner internalizes the cross-group forest externality, solving

$$\max_{\{c_{jt},d_{jt},e_{jt},a_{j,t+1},F_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \sum_{j\in\{H,L\}} \left[u(c_{jt}) + \phi_j v(F_t)\right], \tag{28}$$

subject to the group-level budget constraints (2) and the forest law of motion (3). The optimality conditions (derived in Section C.3) introduce a planner shadow value Ω_t that replaces each group's private shadow value:

$$A_j q_d(d_{jt}, a_{jt}) = \frac{1}{\zeta} m'(e_{jt}) = \frac{\Omega_t}{u'(c_{jt})} \equiv s_{jt}^{SP}, \tag{29}$$

$$\Omega_t = \beta\left[(\phi_H + \phi_L)v'(F_{t+1}) + \Omega_{t+1} G'(F_{t+1})\right]. \tag{30}$$

Comparing (30) with (7), group j's private shadow value internalizes only ϕ_j, while the planner internalizes $\phi_H + \phi_L$. This is the common-pool externality.

Proposition 3 (Dynamic social optimum and steady state) *Given initial conditions (F_0, a_{H0}, a_{L0}), an interior social optimum consists of sequences $\left\{c_{jt}^{SP}, d_{jt}^{SP}, e_{jt}^{SP}, a_{j,t+1}^{SP}, \lambda_{jt}^{SP}\right\}_{t=0}^{\infty}$ for $j \in \{H, L\}$, a forest-stock path $\{F_{t+1}^{SP}\}_{t=0}^{\infty}$, and planner shadow values $\{\Omega_t\}_{t=0}^{\infty}$, solving (28) subject to (2) and (3). The social optimum is characterized by (4), (29), (30), (2), (3), and the transversality condition*

$$\lim_{T\to\infty} \beta^T u'(c_{jT}^{SP}) a_{j,T+1}^{SP} = 0, \qquad j \in \{H, L\}. \tag{31}$$

The interior steady state $\left(c_j^{SP,*}, d_j^{SP,*}, e_j^{SP,*}, a_j^{SP,*}, \Omega^*, F^{SP,*}\right)$ satisfies the same golden rule $A_j q_a(d_j^{SP,*}, a_j^{SP,*}) = 1/\beta - (1-\delta)$, together with

$$\Omega^* = \frac{\beta(\phi_H + \phi_L)v'(F^{SP,*})}{1 - \beta G'(F^{SP,*})}, \qquad A_j q_d(d_j^{SP,*}, a_j^{SP,*}) = \frac{1}{\zeta} m'(e_j^{SP,*}) = \frac{\Omega^*}{u'(c_j^{SP,*})},$$

and the forest law of motion and stationary budget constraints from Proposition 2, evaluated at the social optimum.

The saving margin is unchanged: both allocations satisfy the same golden rule. The distortion is confined to the forest-management margin, where the planner internalizes the benefit each group's reforestation creates for the other. This leads to a higher steady-state forest stock, $F^{SP,*} \geq F^*$.

Conservation wedge and Pigouvian implementation. Define the gap between planner and private shadow values as $\Delta_{jt} \equiv \Omega_t - \mu_{jt}$. Subtracting the private recursion (7) from the planner recursion (30) gives

$$\Delta_{jt} = \beta\left[\phi_{-j}\, v'(F_{t+1}) + \Delta_{j,t+1} G'(F_{t+1})\right], \tag{32}$$

so the wedge is driven by the *other* group's forest valuation ϕ_{-j}. In steady state, $\Delta_j^* = \beta\,\phi_{-j}\, v'(F^{SP,*})/[1 - \beta G'(F^{SP,*})] > 0$.

Proposition 4 (Pigouvian implementation) *Let $\tau_{jt} \equiv \Delta_{jt}/u'(c_{jt})$ denote the group-j conservation wedge in consumption-equivalent units. The socially optimal forest path is decentralized by a per-unit tax τ_{jt} on deforestation and a per-unit subsidy $\zeta\tau_{jt}$ on reforestation for each group j. In steady state,*

$$\tau_j^* = \frac{1}{u'(c_j^{SP,*})} \frac{\beta\, \phi_{-j}\, v'(F^{SP,*})}{1 - \beta G'(F^{SP,*})} > 0.$$

The wedge is strictly positive for both groups and proportional to ϕ_{-j}: it is precisely the marginal forest benefit that group j ignores when acting unilaterally. The type-L group therefore receives the larger corrective subsidy when $\phi_H > \phi_L$, because the uninternalized benefit it imposes on type H is the larger of the two. This structure carries directly into the quantitative results, where Panel (D) of Figure B4 plots the time path of τ_{jt} along the transition.

Figure B4 disaggregates the transition path from Figure 11. Panel (A) shows the path of deforestation. In the decentralized economy the low-valuation group L (dashed navy) clears far more land than the high-valuation group H (solid navy), i.e., the group that values the forest more conserves it, exactly as the private first-order conditions predict. Clearing spikes at the capital-poor start and decays monotonically as accumulation proceeds. Panel (B) shows the mirror image for reforestation: the high-valuation group H replants substantially more than L throughout, and both groups' replanting effort rises over the transition as the forest becomes scarce.

The planner allocations (brick) reverse the group ranking, and the reversal carries important intuition of the benchmark model. Once the externality is internalized, the planner applies a *common* social valuation of the forest to both groups: the preference wedge $\phi_H > \phi_L$ that drives the decentralized outcome is switched off, and the two groups no longer differ in how much the forest is worth in welfare terms (their forest shadow values coincide). The only remaining asymmetry is productivity. Because L is the more productive extractor ($A_L > A_H$), it is richer and consumes more, so its marginal utility of consumption is lower. The common social forest value therefore translates into a *higher consumption-equivalent* conservation value for L, since handing a unit of output to the richer group buys little marginal welfare while the forest it would cost is valued identically across groups. Through the reforestation condition $\zeta s_j = m'(e_j)$ and the clearing condition $A_j q_{d,j} = s_j$, this makes the planner ask the richer, more productive group L to clear *less* and replant *more*. The planner thus loads the conservation burden onto the group for which the consumption sacrifice is cheapest at the margin, an allocation that runs opposite to the pure production-efficiency motive and opposite to the decentralized ranking.

Panel (C) reports the dynamic path of capital accumulation. Both groups accumulate from the common capital-poor start towards their steady states, with the more productive group L settling at a slightly higher stock. The decentralized and planner capital paths are close, because the steady-state capital condition pins down the long-run capital stock largely independently of the forest margin; the externality operates principally through the forest and conservation channels rather than through aggregate accumulation.

Panel (D) plots the optimal Pigouvian tax characterized in Proposition 4 that implements the planner's allocation as a decentralized equilibrium. The tax is low early and rises over the transition: when the forest is still abundant its marginal scarcity value is small, so little correction is required; as the forest is drawn down through the development phase it becomes scarce and the optimal tax climbs to its steady-state level. Correction is therefore steepest exactly when the decentralized economy is over-deforesting, reinforcing the reading of Figure 11. Across groups the tax inherits the structure $\tau_j \propto \phi_{-j}$: the low-valuation group faces the far larger levy ($\tau_L^* \approx 0.59$ versus $\tau_H^* \approx 0.12$), because the externality it imposes falls on the high-valuation group, whose forgone intrinsic value from forest is large.

C.2 Decentralized Equilibrium in Benchmark Model

This appendix derives the optimality conditions for the decentralized equilibrium. Throughout, group $j \in \{H, L\}$ takes $\{d_{-j,t}, e_{-j,t}\}_{t=0}^{\infty}$ as given. Let λ_{jt} be the multiplier on group j's budget constraint and μ_{jt} the multiplier on the forest law of motion. The Lagrangian is

$$\begin{aligned}\mathcal{L}_j = \sum_{t=0}^{\infty} \beta^t \Bigg\{ & u(c_{jt}) + \phi_j v(F_t) \\ & + \lambda_{jt} \left[(1-\delta) a_{jt} + A_j q(d_{jt}, a_{jt}) - c_{jt} - m(e_{jt}) - a_{j,t+1} \right] \\ & + \mu_{jt} \left[G(F_t) + \sum_{k \in \{H,L\}} (\zeta e_{kt} - d_{kt}) - F_{t+1} \right] \Bigg\}. \end{aligned} \tag{33}$$

Consumption and saving. The first-order condition for consumption is

$$u'(c_{jt}) = \lambda_{jt},$$

which is (6). The first-order condition with respect to $a_{j,t+1}$ is

$$\lambda_{jt} = \beta \lambda_{j,t+1} \left[1 - \delta + A_j q_a(d_{j,t+1}, a_{j,t+1}) \right],$$

since $a_{j,t+1}$ earns its net marginal product in production next period. Combining with (6) gives the Euler equation (4), with $R_{j,t+1} \equiv 1 - \delta + A_j q_a(d_{j,t+1}, a_{j,t+1})$.

Deforestation and reforestation. The first-order conditions with respect to d_{jt} and e_{jt} are

$$\lambda_{jt} A_j q_d(d_{jt}, a_{jt}) = \mu_{jt}, \qquad \lambda_{jt} m'(e_{jt}) = \mu_{jt} \zeta. \tag{34}$$

Dividing through by λ_{jt} gives the forest-management condition (5):

$$A_j q_d(d_{jt}, a_{jt}) = \frac{1}{\zeta} m'(e_{jt}) = \frac{\mu_{jt}}{\lambda_{jt}}.$$

Forest shadow value. The envelope condition with respect to F_{t+1} gives

$$\mu_{jt} = \beta \left[\phi_j v'(F_{t+1}) + \mu_{j,t+1} G'(F_{t+1}) \right],$$

which is (7).

Transversality and capital accumulation. The optimal path must satisfy the transversality condition (27), which prevents the asymptotic over-accumulation of physical capital relative to its shadow value. Because a_{jt} is a productive input rather than a claim paying a market return, the accumulation equation (2) is nonlinear in a_{jt} through $q(d_{jt}, a_{jt})$ and cannot be solved forward into a present-value budget constraint with a common discount factor. The relevant restrictions on the path are therefore (2) itself, together with (27), given a_{j0}.

Steady state. A decentralized steady state is a time-invariant allocation $\left(c_j^*, d_j^*, e_j^*, a_j^*, \lambda_j^*, \mu_j^*, F^* \right)$, $j \in \{H, L\}$, to which the dynamic path converges. The Euler equation (4) implies the modified

golden rule

$$1-\delta+A_j q_a(d_j^*,a_j^*)=\frac{1}{\beta},\qquad \text{equivalently}\qquad A_j q_a(d_j^*,a_j^*)=\rho+\delta,\quad \rho\equiv\tfrac{1}{\beta}-1.$$

The consumption first-order condition gives $\lambda_j^*=u'(c_j^*)$. Setting $\mu_{j,t+1}=\mu_{jt}=\mu_j^*$ and $F_{t+1}=F_t=F^*$ in (7) gives

$$\mu_j^*=\frac{\beta\phi_j v'(F^*)}{1-\beta G'(F^*)},$$

provided $1-\beta G'(F^*)>0$. The stationary forest-management condition is

$$A_j q_d(d_j^*,a_j^*)=\frac{m'(e_j^*)}{\zeta}=\frac{\mu_j^*}{u'(c_j^*)},$$

the forest law of motion gives

$$F^*=G(F^*)+\sum_{j\in\{H,L\}}\left(e_j^*-d_j^*\right),$$

and the stationary budget constraint (2) gives

$$c_j^*=A_j q(d_j^*,a_j^*)-m(e_j^*)-\delta a_j^*.$$

These are the conditions in Proposition 2. The golden rule ties a_j^* to d_j^*, so the long-run asset positions (a_H^*,a_L^*) are determined by fundamentals. The initial conditions (a_{H0},a_{L0}) instead govern the transition path toward this steady state.

Forest-management ranking. Let $s_j^*\equiv\mu_j^*/u'(c_j^*)$ denote the steady-state consumption-equivalent shadow value of the forest. The reforestation condition $m'(e_j^*)=s_j^*$ with m' increasing implies that e_j^* is strictly increasing in s_j^*, so $s_H^*>s_L^*$ implies $e_H^*>e_L^*$. Deforestation d_j^* and a_j^* are jointly determined by

$$A_j q_a(d_j^*,a_j^*)=\rho+\delta,\qquad A_j q_d(d_j^*,a_j^*)=s_j^*.$$

Totally differentiating this block in s_j^* (holding A_j fixed), and assuming q is strictly jointly concave so that $q_{dd}<0$, $q_{aa}<0$, and $q_{dd}q_{aa}-q_{da}^2>0$, yields

$$\frac{\partial d_j^*}{\partial s_j^*}=\frac{q_{aa}}{A_j\left(q_{dd}q_{aa}-q_{da}^2\right)}<0,\qquad \frac{\partial a_j^*}{\partial s_j^*}=\frac{-q_{da}}{A_j\left(q_{dd}q_{aa}-q_{da}^2\right)}.$$

Thus a higher forest shadow value lowers deforestation within a group, while its effect on capital has the sign of $-q_{da}$ (negative when d and a are complements, $q_{da}>0$). The cross-group ranking of d_H^* versus d_L^* now depends jointly on s_j^*, the productivity gap A_H versus A_L, and the induced capital stocks. These conditions are collected in Proposition 2 in Section C.1.

C.3 Social Planner's Problem in Benchmark Model

This appendix derives the social planner's optimality conditions. To make the comparison with the decentralized equilibrium transparent, the planner respects the same group-level budget constraints (2) but internalizes the fact that the aggregate forest stock enters the utility of both groups. Let λ_{jt}^{SP} denote the multiplier on group j's budget constraint and let Ω_t denote the planner's

shadow value of the aggregate forest stock. The planner's Lagrangian is

$$\mathcal{L}^{SP} = \sum_{t=0}^{\infty} \beta^t \Bigg\{ \sum_{j\in\{H,L\}} \left[u(c_{jt}) + \phi_j v(F_t)\right] + \sum_{j\in\{H,L\}} \lambda_{jt}^{SP} \left[(1-\delta)a_{jt} + A_j q(d_{jt}, a_{jt}) - c_{jt} - m(e_{jt}) - a_{j,t+1}\right] + \Omega_t \left[G(F_t) + \sum_{j\in\{H,L\}} (\zeta e_{jt} - d_{jt}) - F_{t+1} \right] \Bigg\}. \tag{35}$$

For notational simplicity, this derivation suppresses the SP superscript on choice variables. The superscript is restored in the statement of Proposition 3.

Consumption and saving. The first-order condition for consumption is

$$u'(c_{jt}) = \lambda_{jt}^{SP}. \tag{36}$$

The first-order condition with respect to $a_{j,t+1}$ is

$$\lambda_{jt}^{SP} = \beta \lambda_{j,t+1}^{SP} \left[1 - \delta + A_j q_a(d_{j,t+1}, a_{j,t+1})\right], \tag{37}$$

since each unit of $a_{j,t+1}$ earns its net marginal product in group j's own production next period. Combining (36) and (37) gives the same Euler equation as in the decentralized equilibrium, namely (4), with $R_{j,t+1} \equiv 1 - \delta + A_j q_a(d_{j,t+1}, a_{j,t+1})$.

Deforestation and reforestation. The first-order condition with respect to deforestation by group j is

$$\lambda_{jt}^{SP} A_j q_d(d_{jt}, a_{jt}) = \Omega_t. \tag{38}$$

The first-order condition with respect to reforestation by group j is

$$\lambda_{jt}^{SP} m'(e_{jt}) = \Omega_t \zeta. \tag{39}$$

Using (36), equations (38) and (39) imply the planner's forest-management condition (29):

$$A_j q_d(d_{jt}, a_{jt}) = \frac{1}{\zeta} m'(e_{jt}) = \frac{\Omega_t}{u'(c_{jt})}.$$

Planner's forest shadow value. The envelope condition with respect to F_{t+1} gives

$$\Omega_t = \beta \left[(\phi_H + \phi_L) v'(F_{t+1}) + \Omega_{t+1} G'(F_{t+1})\right],$$

which is equation (30) in the main text.

The key difference from the decentralized equilibrium is therefore the shadow value of the forest. In the decentralized equilibrium, group j's private shadow value in (7) internalizes only its own forest valuation ϕ_j. The planner's shadow value in (30) internalizes the total forest valuation $\phi_H + \phi_L$.

Transversality and capital accumulation. The planner's allocation must also satisfy the group-level budget constraints (2), the forest law of motion (3), and the transversality conditions (31). Because each a_{jt} is a productive input rather than a claim paying a common market return, the accumulation equation (2) is nonlinear in a_{jt} through $q(d_{jt}, a_{jt})$ and cannot be solved forward into a present-value budget constraint with a common discount factor. The relevant restrictions on each group's path are therefore (2) itself, together with (31), given a_{j0}.

Planner's steady state. A social optimum steady state is a time-invariant allocation

$$\left(c_j^{SP,*}, d_j^{SP,*}, e_j^{SP,*}, a_j^{SP,*}, \lambda_j^{SP,*}, \Omega^*, F^{SP,*}\right), \qquad j \in \{H, L\},$$

to which the planner's allocation converges. The Euler equation (4) implies the modified golden rule

$$1 - \delta + A_j q_a(d_j^{SP,*}, a_j^{SP,*}) = \frac{1}{\beta}, \qquad \text{equivalently} \qquad A_j q_a(d_j^{SP,*}, a_j^{SP,*}) = \rho + \delta, \quad \rho \equiv \tfrac{1}{\beta} - 1.$$

The consumption first-order condition (36) implies

$$\lambda_j^{SP,*} = u'(c_j^{SP,*}).$$

Setting $\Omega_{t+1} = \Omega_t = \Omega^*$ and $F_{t+1} = F_t = F^{SP,*}$ in (30) gives

$$\Omega^* = \frac{\beta(\phi_H + \phi_L)v'(F^{SP,*})}{1 - \beta G'(F^{SP,*})},$$

provided $1 - \beta G'(F^{SP,*}) > 0$. The steady-state version of (29) is

$$A_j q_d(d_j^{SP,*}, a_j^{SP,*}) = \frac{m'(e_j^{SP,*})}{\zeta} = \frac{\Omega^*}{u'(c_j^{SP,*})}.$$

The forest law of motion (3) implies

$$F^{SP,*} = G(F^{SP,*}) + \sum_{j \in \{H,L\}} \left(e_j^{SP,*} - d_j^{SP,*}\right),$$

and the group-level budget constraint (2) implies

$$c_j^{SP,*} = A_j q(d_j^{SP,*}, a_j^{SP,*}) - m(e_j^{SP,*}) - \delta a_j^{SP,*}.$$

These are the steady-state conditions collected in Proposition 3 in Section C.1. The modified golden rule ties $a_j^{SP,*}$ to $d_j^{SP,*}$, so the long-run asset positions are determined by fundamentals and the initial conditions govern only the transition path.

C.4 Two-Country Extension: Full Derivations

This section provides the complete derivations for the two-country model, supplementing the condensed presentation in Section 6. Throughout, $n \in \{N, S\}$ indexes countries and $j \in \{H, L\}$ indexes groups within each country.

C.4.1 Global Social Planner

A global social planner chooses $\left\{c^G_{njt}, d^G_{njt}, e^G_{njt}, a^G_{nj,t+1}, F^G_{n,t+1}, X^G_{t+1}\right\}_{t=0}^{\infty}$ for all $n \in \{N, S\}$ and $j \in \{H, L\}$, taking initial conditions $\{F_{n0}\}_{n\in\{N,S\}}$, X_0, and $\{a_{nj0}\}_{n,j}$ as given. The planner maximizes

$$\max \sum_{t=0}^{\infty} \beta^t \left\{ \sum_{n\in\{N,S\}} \sum_{j\in\{H,L\}} \left[u(c_{njt}) + \phi_j v(F_{nt})\right] - \chi D(X_t) \right\}, \tag{40}$$

subject to the country-level budget constraints (8), the local forest laws of motion, the emissions equation (9), and the global carbon law of motion (10). The function $D(X)$ captures global damage from atmospheric carbon, with $D'(X) > 0$ and $D''(X) \geq 0$.

Let λ^G_{njt} denote the multiplier on group j's budget constraint in country n, Ω_{nt} the planner's shadow value of the local forest stock in country n, and $\Psi_t \geq 0$ the shadow cost of the global atmospheric carbon stock. The Euler equation takes the same form as (4), now indexed by country and group:

$$u'(c^G_{njt}) = \beta R_{nj,t+1}\, u'(c^G_{nj,t+1}), \qquad R_{nj,t+1} \equiv 1 - \delta + A_{nj} q_a(d^G_{nj,t+1}, a^G_{nj,t+1}),$$

where the gross return to wealth is each group's marginal product of capital net of depreciation, scaled by the country-specific productivity A_{nj}.

The planner's local forest shadow value satisfies (11) and the global carbon shadow cost satisfies (12). The forest-management conditions are (13) and (14). The allocation must also satisfy the transversality condition

$$\lim_{T\to\infty} \beta^T u'(c^G_{njT}) a^G_{nj,T+1} = 0, \qquad n \in \{N, S\}, \quad j \in \{H, L\}. \tag{41}$$

The full characterization is collected in Proposition 5. The planner internalizes two externalities absent from private decisions. First, within each country, the planner's shadow value Ω_{nt} (equation (11)) reflects the total valuation $\phi_H + \phi_L$ of both groups, whereas each group's private shadow value reflects only its own ϕ_j. Second, across countries, the planner's carbon shadow cost Ψ_t (equation (12)) internalizes the global damage from net emissions, whereas each group takes the atmospheric carbon stock as given.

C.4.2 Laissez-Faire in the Global Economy

Under laissez-faire, each group j in country n chooses $\{c_{njt}, d_{njt}, e_{njt}, a_{nj,t+1}\}_{t=0}^{\infty}$ to maximize its own discounted utility, taking other groups' choices and the global carbon stock path as given. This allocation is the direct country-level analogue of the decentralized equilibrium in Proposition 2. Let μ^{LF}_{njt} denote the private shadow value of country n's local forest for group j.

The Euler equation takes the same form as (4) with the country-and-group-specific return $R_{nj,t+1} \equiv 1 - \delta + A_{nj} q_a(d^{LF}_{nj,t+1}, a^{LF}_{nj,t+1})$. The private forest shadow value is the country-level analogue of (7):

$$\mu^{LF}_{njt} = \beta \left[\phi_j v'(F^{LF}_{n,t+1}) + \mu^{LF}_{nj,t+1} G'(F^{LF}_{n,t+1})\right]. \tag{42}$$

At an interior optimum, forest management satisfies the country-level analogue of (5):

$$A_{nj} q_d(d^{LF}_{njt}, a^{LF}_{njt}) = \frac{1}{\zeta} m'_n(e^{LF}_{njt}) = \frac{\mu^{LF}_{njt}}{u'(c^{LF}_{njt})}. \tag{43}$$

Relative to the global planner, laissez-faire omits two terms. It omits the within-country forest benefit accruing to the other group, because μ^{LF}_{njt} depends only on ϕ_j rather than $\phi_H + \phi_L$. It also omits the global carbon shadow cost Ψ_t, because each group takes the global carbon stock as given. Hence laissez-faire generally features excessive deforestation and insufficient reforestation relative to the global social optimum.

The productivity asymmetry further amplifies this distortion. Since $A_{Nj} = \gamma A_j > A_j = A_{Sj}$, the high-productivity country N has a higher private return to deforestation. Holding shadow values fixed, country N therefore has a stronger incentive to deforest than country S.

C.4.3 Carbon-Rights Market: Full Setup and Discussion

Country n receives an initial allocation of carbon rights B_{nt}, with $\sum_{n\in\{N,S\}} B_{nt} = \bar{E}_t$. Carbon rights are traded in a competitive global market at price p^C_t. Production from deforestation requires carbon rights, while reforestation generates carbon credits. The group-level budget constraint becomes

$$a_{nj,t+1} = (1-\delta)a_{njt} + A_{nj}q(d_{njt}, a_{njt}) - c_{njt} - m_n(e_{njt}) - p^C_t \xi_d d_{njt} + p^C_t \xi_e \zeta e_{njt} + T_{njt}, \tag{44}$$

where T_{njt} is the value of the initial carbon-rights allocation rebated to group j in country n. The carbon market therefore changes private incentives along the carbon dimension: deforestation now carries the additional cost $p^C_t \xi_d$, while reforestation receives the credit value $p^C_t \zeta \xi_e$.

Let μ^C_{njt} denote the private shadow value of the local forest stock under the carbon market. The Euler equation again takes the same form as (4). The local forest shadow value is the country-level analogue of (7):

$$\mu^C_{njt} = \beta \left[\phi_j v'(F^C_{n,t+1}) + \mu^C_{nj,t+1} G'(F^C_{n,t+1}) \right]. \tag{45}$$

The forest-management condition (16) equates the carbon-adjusted marginal return to deforestation, the carbon-adjusted marginal cost of reforestation, and the private forest shadow value in consumption units.

Proposition 1 shows that a competitive carbon market alone does not generally implement the global social optimum. The carbon price is asked to do two jobs and is a generally incomplete instrument. It can price the carbon content of forest management, captured by ξ_d and ξ_e, but cannot simultaneously correct the local forest amenity wedge. The planner's efficiency wedge contains both a global carbon component (via Ψ_t) and a local forest amenity component (via Ω_{nt} versus μ^C_{njt}), whereas the carbon market prices only the former. Therefore, unless local forest amenities are separately internalized — for example, through a complementary domestic policy instrument — the carbon-rights equilibrium remains distorted relative to the global social optimum.

In fact, even abstracting from the local amenity value of forests, a single carbon price faces a further obstacle. The planner's carbon shadow cost is converted into consumption units using the group-specific marginal utility of income: the planner's carbon wedge for group j in country n is proportional to $\Psi_t / u'(c_{njt})$. When marginal utilities differ across groups or countries — as they do along any transition path with heterogeneous wealth — a common carbon price cannot reproduce these heterogeneous consumption-equivalent wedges, unless additional instruments equalize marginal utilities across agents.

This distinction is central for the quantitative analysis. The relevant question is not whether a carbon market fully implements the first best, but how much of the welfare loss it can close. In the quantitative exercise, we evaluate the welfare gains from introducing a global carbon market and decompose the remaining gap relative to the social optimum into the carbon component and the residual local-amenity component.

C.5 Global Social Optimum

Proposition 5 (Global social optimum) *Given initial conditions* $\{F_{n0}\}_{n\in\{N,S\}}$, X_0, *and* $\{a_{nj0}\}_{n,j}$, *the global social optimum is characterized by the country-level Euler equation* (4) *(indexed by country and group), the planner objective* (40), *the local forest shadow-value equation* (11), *the global carbon shadow-cost equation* (12), *the forest-management conditions* (13) *and* (14), *the state equations* (10) *and the local forest laws of motion, the budget constraints* (8), *and the transversality condition* (41).

C.6 Proof of Proposition 1

The carbon price required to implement the planner's deforestation choice is:

$$p_t^C = \frac{\Psi_t}{u'(c_{njt}^G)} + \frac{1}{\xi_d}\left[\frac{\Omega_{nt}}{u'(c_{njt}^G)} - \frac{\mu_{njt}^C}{u'(c_{njt}^C)}\right].$$

The carbon price required to implement the planner's reforestation choice is:

$$p_t^C = \frac{\Psi_t}{u'(c_{njt}^G)} + \frac{1}{\xi_e}\left[\frac{\Omega_{nt}}{u'(c_{njt}^G)} - \frac{\mu_{njt}^C}{u'(c_{njt}^C)}\right].$$

A single price satisfies both only if $\xi_d = \xi_e$ or if

$$\frac{\Omega_{nt}}{u'(c_{njt}^G)} - \frac{\mu_{njt}^C}{u'(c_{njt}^C)} = 0.$$

However, Ω_{nt} is governed by (11), where the planner internalizes $\phi_H + \phi_L$. The private shadow value satisfies $\mu_{njt}^C = \beta[\phi_j v'(F_{n,t+1}^C) + \mu_{nj,t+1}^C G'(F_{n,t+1}^C)]$, which internalizes only ϕ_j. Therefore the condition does not generally hold. □

C.7 Static Model: Proof of Gains from Trade

Setup. Two countries, North (N) and South (S). Northern firms collectively emit E units of carbon, fixed by technology, and must hold one carbon credit per unit of emission. Reforestation in country i sequesters γ_i units of carbon per hectare, where $\gamma_S > \gamma_N$, reflecting the South's superior natural conditions. The cost of reforesting R_i hectares is

$$c_i(R_i) = \frac{w_i}{2\gamma_i}R_i^2,$$

where $w_N > w_S$ is the wage in each region. The South has a double cost advantage: higher sequestration per hectare *and* lower labour cost. The marginal cost per tonne of carbon sequestered is $w_i R_i/\gamma_i^2$, which is unambiguously lower in the South for any given output level.

Global carbon market. A competitive intermediary reforests in the South and sells credits to Northern firms at price p. The intermediary's optimal scale is $R_S^* = p\gamma_S/w_S$. Market clearing requires

$$\gamma_S R_S^* = E \implies p^* = \frac{w_S E}{\gamma_S^2}.$$

Total land used: $R_S^* = E/\gamma_S$. Total compliance cost:

$$C_{\text{global}} = c_S(R_S^*) = \frac{w_S E^2}{2\gamma_S^2}.$$

Autarky. Northern firms source credits domestically. Market clearing gives

$$p^{**} = \frac{w_N E}{\gamma_N^2} > p^*.$$

Since $w_N > w_S$ and $\gamma_N < \gamma_S$, the autarky price strictly exceeds the global price. Total compliance cost:

$$C_{\text{autarky}} = \frac{w_N E^2}{2\gamma_N^2} > C_{\text{global}}.$$

Gains from trade. The cost saving from international trade in carbon credits is

$$\Delta C = C_{\text{autarky}} - C_{\text{global}} = \frac{E^2}{2}\left(\frac{w_N}{\gamma_N^2} - \frac{w_S}{\gamma_S^2}\right) > 0. \tag{46}$$

The gains are increasing in the productivity gap γ_S/γ_N and the wage gap w_N/w_S. The global market also concentrates reforestation in the South ($R_S^* = E/\gamma_S < E/\gamma_N = R_N^{**}$), using less land to achieve the same sequestration target. □

D Additional Background

D.1 Satellite Images: Reforestation and Subsequent Land Use Change

Figure D2 presents Google Earth satellite imagery for a representative reforestation site in Ghana, planted in 2018. The sequence illustrates the general trajectory of a successful project: the site is largely bare in the pre-treatment period (2012–2014), greens up noticeably after planting (2018 onward), and shows sustained vegetation growth through 2025. This is what a well-functioning reforestation program looks like from space.

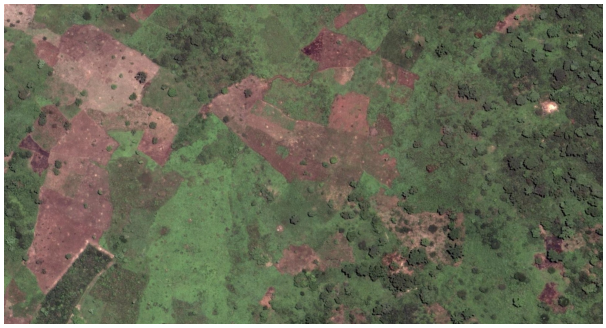
The Ghana case is instructive precisely because it is a success. In the countries where we find negative treatment effects, Ethiopia, Namibia, and Niger, we observe the early stages of this sequence — initial site preparation and planting — but the subsequent greening is either absent or reversed. Satellite inspection of sites in these countries reveals farming activity, livestock grazing tracks, and cleared patches emerging within project boundaries in the years following planting, consistent with the in-migration and encroachment mechanism we propose. The contrast with Ghana suggests that the difference between a successful and a failed program is not the initial planting effort, which is visible in both cases, but whether the planted trees survive the subsequent pressure from surrounding land users.



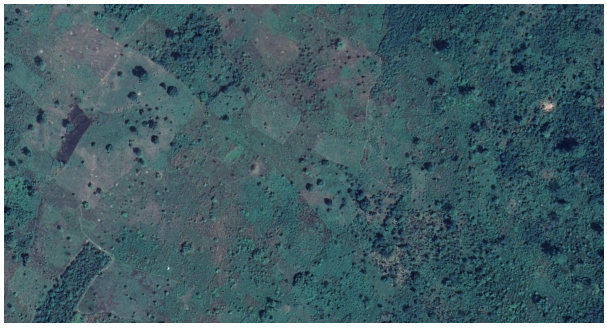
Figure D1. Reforestation Project Site in Ghana: Project Overview

Notes: Project boundary overlaid on satellite imagery of the reforestation site in Ghana, planted in 2018.

(a) May 2012 (pre-treatment)



(b) March 2018 (planting year)



(c) November 2021



(d) December 2023



(e) November 2024



(f) April 2025



Figure D2. Satellite Imagery of a Reforestation Site in Ghana (Planted 2018)

Notes: Google Earth satellite imagery for a reforestation project site in Ghana, planted in 2018. The panels show the site at six points in time spanning 2012 to 2025. The site transitions from bare land prior to planting to sustained dense vegetation by 2021–2025.

D.2 Institutional Background: The Voluntary Carbon Market

The voluntary carbon market (VCM) is a buyer-driven system in which corporations, organisations, and individuals purchase carbon credits on a non-mandatory basis to offset their greenhouse gas emissions. Each credit represents the reduction, avoidance, or removal of one metric ton of CO_2 equivalent. In contrast to compliance markets established by regulation, such as the European Union Emissions Trading System or California's cap-and-trade programme, VCM participation is motivated by corporate climate strategies, net-zero pledges, and environmental, social, and governance objectives. As of 2025, approximately 28 percent of global greenhouse gas emissions are covered by direct carbon pricing, with 80 carbon pricing instruments in operation across economies that collectively account for roughly two-thirds of world GDP (World Bank, 2025).

After a decade of rapid expansion, the VCM entered a transition from scale to quality beginning in 2023. According to Ecosystem Marketplace (2025), global VCM transaction volume fell by 25 percent year-on-year to 84 million metric tons in 2024, the lowest level since 2018. Transaction value declined by 29 percent to $715 million, returning to 2020 levels. Despite this contraction, the average credit price declined by only 5.5 percent, reflecting resilient demand for high-quality credits. Retirements, the most direct measure of end-user demand, stabilised at 182 million metric tons across the ten largest standards. MSCI (2026) report that new credit issuance dropped 13 percent in 2025 to 276 million metric tons, while retirements reached a record 202 million metric tons, and corporate capital commitments to carbon projects totalled $16 billion, led by technology firms such as Microsoft, Google, and Meta. Long-term growth projections vary widely: estimates of market size in 2034 range from $15.7 billion to $451.8 billion, reflecting deep uncertainty about regulatory developments and integrity reforms.

Carbon credits fall into two broad categories with sharply diverging price trajectories. Avoidance and reduction credits, generated by projects that prevent emissions that would otherwise occur (e.g., REDD+ forest conservation, improved cookstoves, and landfill methane capture), have historically dominated issuance but face persistent concerns over additionality and permanence. Removal credits, which actively withdraw CO_2 from the atmosphere, are subdivided into nature-based removals (afforestation, reforestation, soil carbon, blue carbon) and engineered removals (direct air capture, bioenergy with carbon capture and storage, biochar, enhanced weathering). Nature-based removals currently account for 95 percent of all carbon dioxide removal credit issuance, but engineered removals are the fastest-growing segment, projected to expand at a compound annual rate of nearly 56 percent. In 2024, the average price of removal credits was 381 percent higher than that of avoidance credits, and credits rated BBB or above by independent agencies traded at a premium exceeding $7 per ton relative to low-integrity credits (MSCI, 2026).

The institutional architecture of the VCM comprises several distinct groups of participants. The supply side is anchored by standards and registries, most notably Verra (Verified Carbon Standard), Gold Standard, the American Carbon Registry, and the Climate Action Reserve, which establish methodologies, certify projects, and issue credits. Project developers such as South Pole, Anew, and Everland design and operate carbon projects on the ground. On the demand side, the principal buyers are multinational corporations with net-zero commitments: over 4,200 firms have validated or committed to science-based targets, and more than 28,000 corporate buyer profiles are tracked in market databases. Intermediation occurs through brokers, digital platforms (e.g., Pachama, Salesforce Net Zero Marketplace), carbon investment vehicles, and rating agencies such as Sylvera and Calyx Global. In parallel, integrity governance bodies, namely the Integrity Council for the Voluntary Carbon Market (ICVCM) on the supply side and the Voluntary Carbon Markets Integrity Initiative (VCMI) on the demand side, have developed quality benchmarks and corporate claims guidance. The ICVCM's Core Carbon Principles, launched in 2023, have become a reference standard for high-integrity credits, with CCP-approved projects commanding significant price and volume premia.

A growing body of academic evidence has documented systematic weaknesses in credit quality. Probst et al. (2024), synthesising 65 studies covering 972 million metric tons of credits, find that fewer than 16 percent represent genuine emission reductions. Avoidance projects, including Chinese wind power and improved forest management in the United States, show no statistically significant climate benefits, while nature-based removal projects overstate sequestration. West et al. (2023) demonstrate that voluntary REDD+ projects in the Brazilian Amazon systematically overestimate avoided deforestation relative to synthetic controls, and Calel et al. (2025) report similar findings across multiple project types. Fragmentation across more than 60 certification standards compounds the information asymmetry between buyers and sellers. Nonetheless, reforms are advancing: the ICVCM has approved seven project types and 36 methodologies under

the CCP label, the European Union has permitted the use of international carbon credits toward its 2040 climate targets, and disclosure mandates under the ISSB standards and the EU Corporate Sustainability Reporting Directive are increasing transparency. The VCM thus remains a critical, if contested, element of global climate finance, with the potential to channel significant resources to reforestation and other nature-based solutions if integrity challenges can be resolved.