The Global Allocative Efficiency of Deforestation

Prakash Mishra*

January 22, 2026

Draft: please click here for the latest version of this paper.

Abstract

This paper quantifies global inefficient and spatially misallocated agricultural deforestation: carbon emissions-intensive deforestation on land with low agricultural yields. I estimate a global, plot-specific distribution of the elasticity of deforestation to changes in agricultural returns. Landowners in high-deforestation regions are less price elastic. I nest this distribution in a multi-sector equilibrium trade model to solve for a welfare-maximizing global carbon tax, assuming a $190 per ton social cost of carbon and no international revenue transfers. Against this benchmark, 63.1% of carbon emissions from deforestation since 1982 are inefficient. Food production rises under the optimal tax, driven by reallocating agriculture from high-emissions to high-yield land across countries. The US, Canada, and Europe benefit from the status quo allocation of emissions, providing a political economy explanation for the lack of coordinated deforestation policy. When emissions cannot reallocate across countries, optimal carbon taxes achieve 24% fewer emissions reductions and food production declines.

Keywords: deforestation, carbon tax, misallocation, trade, land use
JEL codes: Q56, Q57, Q15, D61, H23

*Dartmouth College, Email: mishrap@wharton.upenn.edu. I thank my dissertation committee, Arthur van Benthem, Susanna Berkouwer, Gilles Duranton, and Ulrich Doraszelski for continuous advising and feedback. This work was part of my dissertation while at Wharton. I am grateful to participants from Wharton, Dartmouth, Yale SOM, Caltech, the University of Georgia, Columbia, UW-Madison, AERE, and Toulouse. I learned immensely from discussions with Jessie Handbury, Sagar Saxena, Shresth Garg, Ben Lockwood, Allan Hsiao, Anna Russo, Teevrat Garg, Mathias Reynaert, Bob Staiger, and Torben Mideksa. This paper benefited from discussions by Heitor Pellegrina, Seleni Cruz, and Sylvain Chabé-Ferret. I gratefully acknowledge financial support from FORMAS grant number 2020-00371, the Wharton Risk Management Center's Russell Ackoff Doctoral Student Fellowship, Wharton's Global Initiative, and the Wharton Climate Center.

1 Introduction

Carbon dioxide emissions have risen precipitously over the last century; without intervention, global temperatures are projected to exceed 1.5°C above pre-industrial levels before 2100. Climate policy has slowed emissions growth, primarily by targeting the manufacturing and transportation sectors (Gillingham and Stock 2018). However, policies targeting deforestation remain underdeveloped relative to other sectors (Balboni et al. 2023). The majority of deforestation since 1980 has been driven by agricultural expansion (Curtis et al. 2018). The resulting emissions from this agricultural deforestation drive 15% of annual global carbon dioxide emissions, more than the entire ground transportation (9%) and airline (4%) fleets (IPCC 2023).[1] Yet, current climate policies rarely target agricultural deforestation, often out of concern for the tradeoff between climate and food security or economic development (Clausing and Wolfram 2023).[2] How can policymakers retarget deforestation to lower emissions while minimizing the loss of necessary food production?

Conceptually, a policymaker should allocate deforestation to land with high agricultural yields and low carbon emissions to maximize efficiency. Low-yield, high-emissions deforestation is misallocated. In theory, efficient climate policy, a global Pigouvian tax at the social cost of carbon, corrects this misallocation.

The extent to which a Pigouvian tax can reduce deforestation operates through three mechanisms. First, consumers must be able to substitute away from food produced in high-deforestation regions like Brazil, potentially consuming more food produced abroad. Trade barriers hamper this substitution (Tombe 2015). Second, decreasing deforestation in a given country reduces agricultural labor demand, pushing labor into non-agricultural sectors. Non-agricultural labor productivity also varies across countries, affecting comparative advantage in agricultural production (Lagakos and Waugh 2013; Nath 2025). Finally, other, lower emissions land must have sufficient agricultural yields to meet excess food demand: this requires knowing the global, joint distribution of agricultural returns and emissions.

I assemble global data on agricultural yields, deforestation emissions, land use, and transportation costs. The data consists of 1,591,999 plots of land at 0.1 arc-degree (approximately 10 kilometers at the equator) resolution. Deforestation since 1982 has released 184 tons of CO_2 per ton of maximum potential agricultural output (i.e., yields under optimal agronomic conditions) on deforested land. At a social cost of carbon of \$190/ton, breaking even with the social value of carbon requires a present value of \$3,487 per ton of potential yield. This is around 10 times the global average.

To account for the role of prices and to simulate the effect of deforestation policies on food production, deforestation emissions, and welfare, I develop a global computable general equilibrium

1. Carbon dioxide, or CO_2, emissions from deforestation come from the burning or processing of felled trees. Trees, in the process of respiration, take in carbon dioxide and release oxygen, removing carbon from the atmosphere and storing it as woody, carbon-rich, biomass.

2. Bolsonaro, controversially, on why the Amazon continues to be developed: "Indigenous people want to work, they want to produce and they can't." (Phillips 2019) A Polish official on recent deforestation: "In the West, first they built their infrastructure, and then laws to protect nature began to be introduced. We — through no fault of our own — have been developing for only 20-odd years and we are forced to (protect the environment) now, taking into account the restrictive environmental protection law." (Skiba 2023)

trade model of deforestation with a granular plot-level supply side. I employ a two-sector, two-factor framework: agriculture (using land and labor with a deforestation externality) and manufacturing (constant-returns, competitive producers using labor). Two agents enter the model: landowners, who make a discrete choice regarding land use (forest/natural use, agriculture, or other) and workers, who supply labor either to agricultural landowners or the outside manufacturing sector.[3] Both agents consume a mix of agriculture and manufacturing goods, buying varieties of each across countries (an Armington demand system). Global social welfare accounts for the carbon externality and depends on three model features: (1) plot-level heterogeneity in yields, deforestation emissions, and transportation costs, which determine producer surplus; (2) cross-country trade costs, which affect consumer surplus and terms of trade; and (3) general equilibrium interactions from labor reallocation across sectors.

As a key input, I estimate the elasticity of land use decisions with respect to agricultural revenue. In a multinomial logit, land use shares in the data map to realized land use returns in dollars. This mapping recovers a key parameter: the elasticity of deforestation to a $1 increase in agricultural returns. Cross-sectional differences of $1 in agricultural returns will be correlated with unobservable land profitability for agriculture. Identification requires variation plausibly unrelated to field-level unobservables. I test three approaches to identification: using tariff shocks abroad to instrument domestic farm-gate prices (a shift-share), using distance to cities and steep slopes to instrument freight costs, and a spatial differences design. The average global elasticity is between 1.2-1.8 across approaches. I estimate a global distribution of unobserved heterogeneity in this elasticity through an expectation-maximization approach (Train 2008). Deforestation in carbon dense regions is less elastic to prices, which hampers gains from carbon pricing.[4]

For the demand side, estimation follows from Costinot, Donaldson, and Smith (2016) and Carleton, Crews, and Nath (2023). The cross-country agricultural trade elasticity is 4.8, meaning a 1% increase in the price of agriculture in France lowers the share of French agriculture in consumption abroad by 4.8%. With both demand and supply estimated, the model rationalizes deforestation given a starting land use state. Intuitively, the general equilibrium wage in the model shifts the level of land returns, while the empirical model above is identified off of relative returns. The level of wages are identified from macroeconomic data. Consistent with prior work focused on agricultural labor supply, manufacturing wages are higher than agricultural wages (Tombe 2015; Nath 2025). Then, if labor can shift across sectors, then a carbon tax results in some efficiency gains. Given a quantified model, I produce six findings.

First, I solve for a global welfare-maximizing tax on carbon emissions at a $190/ton social cost of carbon. I impose two political economy constraints: countries cannot transfer carbon tax revenues, and countries are bound to a common carbon tax. The optimal tax would have abated 63.1%

3. The static land use discrete choice problem captures the breakeven agricultural profit which generates the "first" (and most carbon intensive) deforestation event for primary forest, abstracting from secondary forest or land use dynamics.

4. These differences in elasticities are often driven by institutional differences in the tropics (Angelsen 1999; Hidalgo et al. 2010)

of realized deforestation-related emissions since 1982. Under this policy, the largest reductions in deforestation occur in the interior of major forest basins: Amazonia, the Congo, and Borneo. Remaining deforestation occurs on land with high yields and low emissions.

Second, in equilibrium, agricultural land reallocation results in total food production rising. In partial equilibrium (holding prices and wages fixed), food production would fall by 40%; in general equilibrium, prices adjust and food production rises slightly. Rising farm-gate prices and falling wages create increasing margins for inframarginal food producers, lowering the costs of deforestation prevention in equilibrium. The average landowner profit rises by \$1.91/ha/year, while emissions-weighted profits fall by \$34.88/ha/year. Windfall profits accrue to landowners who produce on low carbon-intensity land.

Third, 24% of optimally avoided emissions come from reallocating deforestation across countries. Furthermore, shutting down this reallocation would result in food losses rather than a gain in food production. To demonstrate this, I solve for the country-level carbon tax which maximizes welfare subject to a constraint that countries must maintain a constant share of global deforestation emissions. Constraining emissions shares dramatically lowers the optimal carbon tax on forested countries, reflecting that they deforest more laxly in business-as-usual. The US and Canada also face lower taxes. In constrast, Mexico and other sub-Saharan African countries are *over*-taxed: they should deforest more but for the constraint. Geographically, emissions which cannot be avoided but for reallocation exist on the periphery of major forest basins: emissions increase in the Brazilian biodiverse cerrado region and nearly all of Indonesia save the interior of Borneo. In terms of welfare, the Congo River Basin, the US, Canada, and Europe all benefit from country-level constraints. Thus, globally coordinated deforestation policy would put these countries at a relative disadvantage, reducing their incentive to cooperate. Nearly all remaining avoided deforestation comes from cross-biome reallocation within countries. Absent the ability to move agriculture from forest interiors to slack grasslands, tropical countries cannot achieve emissions reductions.

Fourth, distributionally, the global carbon tax is clearly regressive. Tropical, high-forest countries with low GDP per capita lose the most as a consequence of the tax. On average, landowning agriculturalists win as a consequence of the policy because most landowners are not located on densely forested land. They merely benefit from rising prices and downward pressure on wages induced by excess labor supply from lost agricultural work. The tax thus amounts to a transfer from marginal, high-forest landowners and workers, who pay higher food costs, to these incumbent landowners and landowners in wealthy countries.

Fifth, I quantify the importance of granular biomass measurement. Tropical monitoring systems for deforestation have become increasingly granular (Assunção, Gandour, and Rocha 2023). However, offset markets and large-scale international investments through the UNFCCC rely on country-biome specific carbon density (UNFCCC 2013). To understand the gains from granular carbon pricing at scale, I implement the optimal carbon tax rate but assume the planner can only levy taxes based on country-specific carbon densities. One-quarter of optimal emissions reductions are still emitted with imperfect monitoring. The carbon densest – and most emissions-costly –

regions emit differentially more, as they are by construction more dense than country-level averages. Because the bulk of the deforestation externality is driven by 20% of the most highly emissive landowners, granularity matters.

Sixth, there are large co-benefits from taxing carbon for biodiversity conservation. While there is no single measure of biodiversity value, I leverage the International Union for the Conservation of Nature (IUCN) Red List habitat maps to measure the potential prevalence of threatened species on every pixel on earth. The previously cited global carbon-only tax decreases loss of threatened species habitats by 17 percent relative to business-as-usual. When the planner can also penalize species loss and values habitat at $1 per threated animal-hectare, they decrease habitat loss by a further 20 percentage points while increaseing emissions by less than 2 percentage points. Additional conservation of threatened species results in more emissions in much of the Global North – Europe and Canada in particular – as well as Southeast Asia and Indonesia where these goals are most spatially distinct.

Related Literature. My approach brings together research on agricultural trade (Costinot, Donaldson, and Smith 2016; Sotelo 2020; Gouel and Laborde 2021; Farrokhi and Pellegrina 2023; Farrokhi et al. 2023; Dominguez-Iino 2023) with research which applies discrete choice methods to land use and deforestation decisions (Scott 2013; Souza-Rodrigues 2019; Araujo, Costa, and Sant'Anna 2025; Hsiao 2025). This latter empirical work highlights relatively elastic deforestation decisions. Carbon tax incentives should thus sharply reduce deforestation. The present paper explicitly estimates a global distribution of land use elasticities to understand the counterfactual geography of agricultural production under carbon taxes, highlighting the role of reallocation in preventing deforestation. This approach opens up a link to research on agricultural misallocation (Adamopoulos and Restuccia 2014; Tombe 2015; Nath 2025).

On the other hand, prior research in trade primarily highlights misallocated agricultural production (Tombe 2015; Foster and Rosenzweig 2021; Nath 2025) or evaluates effects of trade policy on forest cover (Abman and Lundberg 2020), while I link agricultural misallocation and deforestation. This link has been a major gap in prior literature: Alix-Garcia et al. (2025) emphasizes that there is limited micro- or global-scale evidence regarding "win-win" policies for limiting deforestation while encouraging development. This paper suggests that globally coordinated environmental policy is a win-win in part because of prior documented misallocation in agriculture.

A significant theoretical literature motivates partial regulation-induced spillovers across forested countries (Harstad and Mideksa 2017; Harstad 2020; 2023). Such emissions leakage has been evaluated in several environmental contexts (Fowlie 2009; Fowlie, Holland, and Mansur 2012; Perino, Ritz, and Benthem 2025; Restrepo and Mariante 2024; Hsiao 2025). I account for leakage by modeling both equilibrium price effects and changes in factor rewards in general equilibrium.

I highlight two related papers closest to this one. First, Farrokhi et al. (2023) highlight the impact of trade liberalization on deforestation dynamics. In contrast, I focus on carbon tax policy and allocative efficiency. Second, Dominguez-Iino (2023) studies how market power in intermediary firms can affect pass-through of South American carbon taxes onto deforesting farms. I instead focus

on a granular and global approach, highlighting the role of the rest of the world in deforestation prevention.

2 Data and motivating facts

2.1 Data

An observation is a 0.1 arc-degree by 0.1 arc-degree square grid. These plots are approximately 10×10 kilometers at the equator. All statistics are area-weighted to account for differences in area further from the equator. The globe contains 5,040,000 such cells, of which 1,591,999 cover non-polar terrestrial land. In this section, I provide a brief overview of data sources; I reserve other details for the data appendix (Appendix B).

Land use. I use European Space Agency Climate Change Initiative, hereafter ESACCI, maps of landcover as an annual panel from 1995 to 2019 for estimation. I reclassify the base ESACCI data to assign pixels to the categories: forest, cropland, grassland, bare areas, and urban areas. For the purposes of this paper, I consolidate grassland, bare, and urban areas into "other" land use. Critically, this data allows for direct observation of *both* forest and cropland land use shares. In descriptives, I also use Song et al. (2018) as a continuous measurement of tree canopy cover change between 1982 and 2016 for a slightly longer panel.

Deforestation emissions To calculate emissions from deforestation, I use a biomass map from NASA's Oak Ridge National Laboratory DAAC. I assume a constant biomass density within plots. For a given plot ω with land use share $s^F(\omega)$, I calculate deforestation emissions as:

$$\text{emissions}(\omega; s^F) = \frac{44}{12}\frac{1}{2}[1 - s^F(\omega)]\text{biomass}_{1980}(\omega) \tag{1}$$

That is, I convert biomass into carbon using a factor of $\frac{1}{2}$, and carbon into carbon dioxide using a ratio of molecular weights, 44/12 (UNFCCC 2013). Then, I assume deforestation proportionally destroys biomass to the original stock of biomass on a parcel – this amounts to assuming deforestation is a complete burn.

I briefly discuss why I treat biomass as an exogenous feature of land. Emissions depend on two variables in general: a measure of tree volume – in regional studies this is usually a count or area, while in situ studies often use diameter at breast height – and a regional or species-specific coefficient which converts these measurements to carbon (UNFCCC 2013). This latter biomass coefficient is effectively a land-specific productivity measure: conditional on the same level of treecover, land in the Amazon tends to sequester more carbon per unit forest.[5] To verify this, I regress biomass on treecover percentage. Even with flexible cubic spline terms in tree cover, the model explains only

5. IPCC methodologies can be strikingly broad in their use of biomass coefficients. For REDD+ (avoided deforestation offset) carbon accounting, biomass coefficients tend to focus on multiple countries in nature, and within-region variation only comes from rainfall intensity. Country fixed effects and forest share collectively capture 73% of variation.

58% of variation in biomass per hectare—indicating substantial residual variation in carbon density conditional on forest extent.

Potential agricultural productivity Agricultural productivity is measured using the FAO GAEZ maps, which have a native 9-by-9 km resolution. This dataset, used extensively in prior economic research (Costinot, Donaldson, and Smith 2016), provides crop-specific measures of biophysical suitability based on climate, soil quality, and terrain. I use high-input, rainfed yields for the climate epoch 1980-2010. My focal crops are rice (the maximum of dryland and wetland yields), wheat, maize, soybean, potatoes (maximum of white and sweet potato yields), sugar cane, and oilpalm. Oilpalm is included due to its relevance for Southeast Asian deforestation (Hsiao 2025). For each pixel, I construct an agricultural productivity index as a weighted average of crop-specific potential yields. Following Scott (2013), I use country-level crop land use shares as weights. Throughout this analysis, I maintain the assumption that pixel-level potential yields are reasonable proxies for across-pixel, within country differences in agricultural productivity.

Biodiversity data In counterfactual simulations examining co-benefits, I incorporate data on the biodiversity value of forests, measured as potential habitat area for threatened species from the IUCN Red List; see Appendix **??** for details.

Production and trade flows International production and trade flow data come from the International Trade and Production Database for Estimation (ITPD-E; Borchert et al. 2021), which provides HS 2-digit sector-level data by country pair and year. I adopt their classification of HS-2 codes into agriculture, and group all other sectors—manufacturing, mining, and services—into a single non-agricultural sector. For country-crop-level prices and production, I use data from the Food and Agriculture Organization.

2.2 Describing global forest cover

Global forests vary substantially in carbon density.

The top panel of Figure 1 maps potential emissions from deforestation, derived from the carbon content of trees as in Equation (1). The bottom panel maps the maximum potential agricultural yield across analyzed crops. Both panels depict land-specific productivity measures driven by climate and soil characteristics. While tropical regions support high biomass accumulation due to warmth and moisture, many tropical soils are nutrient-poor without sustained fertilizer application. Dense, carbon-rich woody biomass is better adapted to grow in this tropical band: the Amazon River Basin (Brazil and surrounding countries), the Congo River Basin (primarily in the Democratic Republic of the Congo), and Southeast Asia (particularly Indonesia). Temperate forests in Europe and the US are not as dense.[6]

6. True long-run evidence on carbon density prior to the satellite record is scarce. The longest run temporally consistent records on forest cover come from pollen records, but these records are of varying sources and quality with large confidence bands on implied carbon. Generally, sources agree that the Amazon has always been the most carbon rich. However, this may be due to area or density at various points in time depending on trends in long-run precipitation (Malhi et al. 2012; Knight et al. 2022).

In Appendix Figure A.6, correlates agricultural yields with an alternate measure of forest value – species richness. The left panel shows only the within country relationship. On average, higher yield land is also more species rich.

Deforestation is spatially concentrated.

Since 1982, 3.34% of global land area has lost more than 10% of its vegetation cover. The average parcel experienced a 1.26 percentage point increase in vegetation cover. Aggregate deforestation from 1982-2016 has removed forests from 1,299,865 sq. km. of land. Appendix Table A.15 breaks down vegetation levels and changes from 1982-2016 in the top 50 countries by deforestation rate. Many countries with high deforestation rates have nonetheless experienced net increases in vegetation cover over this period. In most countries, the majority of land area experienced stable or increasing forest cover: exceptions are Paraguay, Mozambique, Belize, and Zambia.

Some deforestation occurs on land with high social costs and low private returns.

As shown in the first panel of Table 1 1, carbon emissions per hectare of deforestation vary dramatically across regions. The top 25% of global land has at least 29 tons of biomass per ha, which if burned would generate 53 tons of CO_2 emissions following equation (1). The bottom 25% of land generates at most 3% as much CO_2 per hectare as these carbon-rich parcels.

Second, of global land area, deforested regions (with losses of at least 10% of vegetation cover) are on average more carbon dense, emitting 17 tons of CO_2 more per hectare than the average global parcel. This land does have better best-case yields and lower travel time to nearby cities than the average parcel of land.

Third, tropical deforested regions comprise only 0.2% of the global surface. Tropical regions are more biomass rich than other deforested regions: the average deforested hectare emits 73 tons of CO_2. Differences in potential yields are small on average (though composition of these yields is different, as indicated by markedly lower wheat yields). However, tropical deforestation tends to occur much further from market.

Taken together, deforestation since 1982 has contributed an estimated 7 billion tons of carbon dioxide emissions to the atmosphere by removing forests on 1,308,000 sq. km. of land. The social cost of these emissions depends on the choice of social cost of carbon. In terms of maximum potential yields, this deforestation contributed a best-case 0.38 billion tons of food production. This new production accounts for 17.4% of the initial maximum potential yields implied by 1980 land use. To break-even against the emissions cost of deforestation at a \$190 social cost of carbon, the average ton of food which was produced must be valued at \$3,494/ton. If this food was produced at the same level every year, a discount rate of 0.98 (the same used to produce the \$190 SCC) would suggest a value per ton of \$70.

In Appendix Table A.16, I summarize country's individual contributions to global emissions over this time period. Brazil contributes 34.4% of global emissions since 1982. Boreal deforestation, driven largely by rotating timber production patterns in Canada and Russia, accounts for 11%

of global emissions in the last 40 years. Appendix Table A.14 then summarizes yield in these countries. Some countries appear to be better on average at deforesting high yield land, including Argentina. Countries whose deforestation is largely driven by timber production – the USA, Russia, and Canada – deforest very poor quality agricultural land.

3 Model

3.1 Economic Environment

There are N plots of land of equal area and L workers, partitioned among J countries. Countries are indexed by i. Land and workers are immobile across countries. Workers are mobile across sectors within countries.Each plot contains a unit measure of land parcels, with each parcel owned by a distinct landowner. Thus, there are two distinct agents: workers and landowners.

There are two private goods—agriculture (produced with land and labor) and manufacturing (produced with labor only)—and one public good: carbon sequestration from forests. International trade in private goods is subject to iceberg transport costs. Consequently, prices for the two traded goods (agriculture and manufacturing) differ across countries.

Each plot ω is characterized by its endowments and chooses land use $h \in \{F, A, O\}$: forest, agriculture, or other. Each plot is endowed with characteristics: potential yields, initial forest cover, and transport costs.

Agriculture and the landowners' problem. Each landowner must decide among land uses. Each plot is endowed an observable (to the landowner and the econometrician via potential yield data) productivity $\eta^A(\omega)$. Landowners require their parcel and a worker to produce $a_i\eta^A(\omega)$ units of output, where a_i is a country-level productivity shifter. Landowners sell output at price p_i per unit and hire labor at wage w_i.

$$\pi^A_\epsilon(\omega) = a_i[p_i - c(\omega)]\eta^A(\omega) - w_i^A + \epsilon^A(\omega). \tag{2}$$

This landowner's agricultural return is micro-founded by a Leontief production function in land parcels and labor, with total factor productivity $a_i\eta^A(\omega)$. The cost term $c(\omega)$ captures plot-specific transport costs to market.

Landowners can alternatively leave their land in a natural (forested) use and earn returns π^F. Forest land yields non-market amenity value $\xi^F(\omega)$ to the landowner, which may reflect private use value, bequest motives, or unmeasured returns. The returns to natural forest are then:

$$\pi^F_\epsilon(\omega) = \xi^F(\omega) + \epsilon^F(\omega).$$

I allow landowners to choose a third land use: other. Other corresponds to urban, grassland, and bare areas. Returns to other land uses, $O(\omega)$, are taken as exogenous to the model.

$$\pi^O_\epsilon(\omega) = \xi^O(\omega) + \epsilon^O(\omega).$$

The following discrete choice problem results:

$$\pi(\omega) \coloneqq \mathbb{E}_\epsilon \max_{h \in \{F,A,O\}} \pi^h_\epsilon(\omega) = \mathbb{E}_\epsilon \max_{h \in \{F,A,O\}} \pi^h(\omega) + \epsilon^h(\omega). \tag{3}$$

The shocks ϵ^h are observed by the landowner but unobserved to the econometrician, representing shocks to the individual landowners' parcels within plot ω. Each landowner allocates their entire parcel to a single use, so individual land use is a corner solution: $\mu^h_\epsilon(\omega) \in \{0, 1\}$ for each h. Then, a representative landowners' returns $\pi(\omega)$ on plot ω is the expected return across realizations of parcel-specific shocks ϵ^h, with corresponding conditional choice probability $\mu^h(\omega)$.

Agricultural land use reflects selection on several margins. Within a country, I assume agricultural land has no amenity value. Bar other land uses, selection within a country is driven by transportation costs, yields, and the unobserved revenue shock ϵ. Incorporating other land uses within the same country, urban areas compete for land with an empirical distribution of values $\xi^O(\omega)$; the model is agnostic on their source. Across countries, selection depends on differences in the marginal cost of agriculture (p_i), including the the opportunity cost of manufacturing labor (w_i). Importantly, the marginal agricultural producer in a given country is a convex combination of plot-level producers. Then, aggregate agricultural supply curvature and producer surplus represent plot-level variation in revealed returns from deforestation rather than the pure distributional assumption on landowner-level shocks ϵ (Eaton and Kortum 2002).

Cattle production is not explicitly modeled in this approach. This reflects a data limitation. Country-level data provides pastoral land stock, prices, and outputs. However, I lack global plot-level satellite data on pasture that would identify pasture separately from other grassland. Existing data sources cannot reliably distinguish managed pasture from natural grassland.[7] In counterfactuals, I treat pasture as part of the 'other' category.

Labor supply. Total labor supply in country i is perfectly inelastic at L_i. If a worker is employed in agriculture, they are paid wage w^A_i by their landowner. All non-agricultural workers, of mass $L^M_i \leq L_i$, earn a wage w^M_i in the manufacturing sector, which I return to below. The wedge between wages in each sector is fixed at a proportionality constant l_i, $w^M_i = l_i w^A_i$ (Tombe 2015).

Trade and transportation. For the agricultural industry, I adopt a hub-and-spoke transportation network assumption (Farrokhi and Pellegrina 2023). Agricultural goods are first domestically ferried to a major city or port market. From these "hubs," goods are sold on the international market. Within the country, the transportation cost enters landowners' marginal cost in Equation (2). Internationally, transportation costs are good-specific iceberg costs, T^h_{ij} (delivering a unit from country i to country j requires producing $T^h_{ij} \geq 1$ units in i).

I assume manufacturing is located directly in hubs. As a result, manufacturing goods only face international trade costs T^M_{ij}. Trade costs T^h_{ij} capture both natural barriers (distance, geography) and policy barriers (tariffs).

7. At the time of writing, the Global Pasture Watch collates data discerning pastoral grassland and natural grasses. However, the authors warn that this data cannot distinguish grass from cropland consistently.

The manufacturing industry. Manufacturing is perfectly competitive with constant returns to scale in labor. Define manufacturing technological productivity in country i as $\bar{\eta}_i^M$. This gives a profit maximization problem, with price of output p_i^M in country i and wage payment w_i^M to L_i^M total workers:

$$\pi_i^M(M_i) = (p_i^M \bar{\eta}_i^M - w_i^<)L_i^M$$

Producers price at marginal cost. Since labor is the only input, firms in country i set prices according to:

$$w_i^M = \bar{\eta}_i^M p_i^M \tag{4}$$

3.2 Consumption

The two agents, landowners and workers, consume agriculture and manufacturing. I adopt homothetic preferences which will imply that all agents spend a constant share of their income on each good. Thus, to ease exposition, I focus on a representative consumer in this section. In country i, the representative consumer earns income Y_i.

Preferences. Preferences follow Costinot, Donaldson, and Smith (2016), augmented to include an emissions externality. All countries share common price elasticities σ and σ^h, discussed below. Across countries, consumers can differ in their tastes for individual varieties based on preference shifters $\vec{\theta}$. In each country $j = 1, 2 \ldots, J$, the representative consumer has a nested, constant elasticity of substitution utility. In the innermost nest, consumers choose between country-specific varieties of the agricultural A and manufacturing M goods. Preference shifters θ_{ij}^h capture quality differences across origin countries, and σ^h governs the substitution elasticity across origins within each sector. Higher values of σ^h indicate greater substitutability between countries of origin.

$$U_j^h(X) = \left(\sum_{i=1}^{N} \theta_{ij}^h X_{ij}^h \right)^{\frac{\sigma^h - 1}{\sigma^h}} \quad \text{for } h \in \{1, 2, \ldots, K, M\} \tag{5}$$

Total demand for good h is $X_j^h = \sum_{i=1}^N X_{ij}^h$. In the final, outermost nest, the consumer chooses across sectoral goods:

$$U_j(x) = \left(\sum_{h \in \{A,M\}} (\theta_j^h X_j^h)^{\frac{(\sigma-1)}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \tag{6}$$

The outer elasticity $\sigma > 0$ determines complementarity ($\sigma < 1$) or substitutability ($\sigma > 1$) between agriculture and manufacturing.

Finally, consumers experience disutility from global carbon emissions D, modeled as additively separable with marginal damage κ. Because emissions are a global externality, individual consumers do not internalize the damage D from their consumption choices. The global welfare criterion is W:

$$W = \sum_{j=1}^{N} \psi_j U_j(X) - \kappa D$$

where ψ_j are welfare weights.

Budget constraint. To complete the utility maximization problem, I define an aggregate budget constraint. Expenditure on a given variety (country) i of agriculture is $E_{ij}^A = (1 + T_{ij}^A) p_i^A X_{ij}^A$; in manufacturing, it is $E_{ij}^M = (1+T_{ij}) p_i^M X_{ij}^M$. The representative consumer faces an aggregate budget constraint on expenditures:

$$\sum_{i=1}^{J} (E_{ij}^A + E_{ij}^M) = Y_j.$$

With a utility function and budget constraint defined, I next derive expenditure shares for the representative consumer in each nest of the utility maximization problem.

Inner nest expenditures. I begin with results for inner nest expenditures, which hold fixed the outer nest expenditures E_j^h in a given sector h. In the manufacturing industry, the representative consumer spends share $\lambda_j^{i|M} = \frac{E_{ij}^M}{E_j^M}$ of total manufacturing expenditures E_j^M on goods from i,

$$\lambda_j^{i|M} = \frac{(\theta_{ij}^M \bar{\eta}_i^M)^{\sigma^M - 1} (T_{ij}^M w_i)^{1-\sigma^M}}{\sum_{n=1}^{J} (\theta_{nj}^M \bar{\eta}_n^M)^{\sigma^M - 1} (T_{nj}^M w_n)^{1-\sigma^M}}. \tag{7}$$

Similarly, the representative consumer spends share $\lambda_j^{i|A} = \frac{E_{ij}^A}{E_j^A}$ of total expenditures on agriculture E_j^A on agricultural goods from i:

$$\lambda_j^{i|A} = \frac{(\theta_{ij}^A)^{\sigma^A - 1} (T_{ij}^A p_i)^{1-\sigma^A}}{\sum_{n=1}^{J} (\theta_{nj}^A)^{\sigma^A - 1} (T_{nj}^A p_n^A)^{1-\sigma^A}}.$$

Inner nest price indices. Using these inner nest expenditure shares, I derive the Dixit-Stiglitz price index of each sector P_j^h. In manufacturing, the price index is

$$P_j^M = \left(\sum_{n=1}^{J} (\theta_{nj}^M \bar{\eta}_n^M)^{\sigma^M - 1} (T_{nj}^M w_n)^{1-\sigma^M} \right)^{\frac{1}{1-\sigma^M}}.$$

Similarly, in agriculture, the price index is

$$P_j^A = \left(\sum_{n=1}^{J} (\theta_{nj}^A)^{\sigma^A - 1} (T_{nj}^A p_n^A)^{1-\sigma^A} \right)^{\frac{1}{1-\sigma^A}}.$$

Outer nest expenditures. Finally, moving to the outer nest, I derive expenditure shares on each sector λ_j^h as a function of these price indices for each sector.

$$\lambda_i^h = \begin{cases} \frac{(\theta_i^h)^{\sigma-1}(P_i^h)^{1-\sigma}}{\sum_{h \in \{A,M\}} (\theta_i^h)^{\sigma-1}(P_i^h)^{1-\sigma}} & \text{if } \sigma \in (0,1) \cap (1,\infty) \\ \theta_i^h & \text{if } \sigma = 1 \end{cases} \tag{8}$$

where I single out the Cobb-Douglas case when $\sigma = 1$. Thus, as a share of aggregate income in i, Y_i (defined below), the representative consumer spends $E_{ji}^A = \lambda_i^A \lambda_j^{i|A} Y_i$ on agricultural goods from j. I will use the notation $\lambda_{ji}^h = \lambda_i^h \lambda_j^{i|h}$ to denote the unconditional share of income that the representative consumer in i spends on good h from country j.

3.3 Equilibrium: Business-as-usual

I describe here aggregation of supply and demand curves and the market equilibrium concept. I refer to this equilibrium as the business-as-usual (BAU) equilibrium.

Aggregation. For plot ω, the expected plot-level supply curve is

$$q(\omega) = \eta^A(\omega)\mu^A(\omega),$$

where $\mu^A(\omega)$ represents the average share of land across individual parcels within plot ω devoted to agriculture. Aggregating over the set Ω_i of N_i plots ω in i gives an aggregate supply curve:

$$Q_i^A = \sum_{\omega=1}^{N_i} q(\omega)$$

Aggregate rents depend on the functional form of the distribution of parcel-level shocks per plot ω, $\epsilon^h(\omega)$. For Type-I extreme value shocks with a scale parameter $\frac{1}{\gamma(\omega)}$, the aggregate rent takes the following closed-form solution (see Section 3.5 of Train 2009):

$$\Pi_i = \sum_{\omega=1}^{N_i} \frac{1}{\gamma(\omega)} \log \sum_{h \in \{F,A,O\}} \exp[\gamma(\omega)\pi^h(\omega)],$$

Define total income Y_i in the economy of i as the sum of wage income and landowner rents.

$$Y_i = \sum_{h \in \{A,M\}} w_i^h L_i^h + \Pi_i.$$

Finally, define aggregate manufacturing labor demand L_i^M as the residual labor supply after agricultural landowners make their land use decisions:

$$L_i^M = \left(L_i - \sum_{\omega=1}^{N_i} \mu^A(\omega) \right).$$

Market clearing. Equilibrium pins down $2 \times J - 1$ endogenous prices: J country-level agricultural prices p_i, $J-1$ country-level wages w_i with one wage normalized. Goods market clearing in agriculture equates supply (from the landowners' problem) to international expenditures (from the consumer choice problem), giving J unique equilibrium conditions:

$$\sum_{j=1}^{J} E_{ij}^A = p_i Q_i^A. \tag{9}$$

Next, I construct equilibrium manufacturing wages. The manufacturing industry absorbs residual labor supply net of labor used by agriculture. Additionally, all manufacturing surplus is paid out to manufacturing laborers, providing a balance condition:

$$\sum_{j=1}^{J} E_{ij}^M = w_i^M L_i^M \tag{10}$$

I hold the wedge between agricultural and manufacturing wages proportionally fixed, so that:

$$w_i^A = l_i w_i^M \tag{11}$$

Existence of the equilibrium follows from the proof of Theorem 2 of Alvarez and Lucas (2007).[8]

3.4 Internalizing carbon damages from deforestation

Damages from deforestation. Production of agriculture has an unpriced secondary product: carbon emissions. Emissions from deforesting plot ω are $d(\omega)$, denominated in tons CO_2 per hectare. Total global emissions globally sum over plot-level emissions,

8. Numerically, my model converges to a unique equilibrium across a wide range of starting guesses.

$$D = \sum_{\omega=1}^{N} d(\omega)\mu^A(\omega). \tag{12}$$

Counterfactual: carbon taxes. Consider a carbon tax t on deforestation denominated in dollars per ton carbon. Landowners adjust land use decisions as:

$$\pi_\epsilon^A(\omega, t_i) = \pi_\epsilon^A(\omega, 0) - td(\omega)$$

I do not tax conversion to 'other' land uses, as I lack data on returns from these categories. Consequently, counterfactuals allow for variation in manufacturing returns only through the wage w_i^M.

Tax revenue $\mathcal{T}_i = \sum_{\omega=1}^{N_i} td(\omega)\mu^A(\omega)$ is re-distributed within-country as a lump-sum. Economy-wide income, including tax revenues, is $Y_i = \Pi_i + w_i L_i + \mathcal{T}_i$. Appendix A.1 writes out a global planners' problem and optimality conditions for a single global carbon tax.

I demonstrate identification of this tax counterfactual in changes in Appendix Section A.2.

Measuring welfare implications of counterfactuals Total consumption value in business-as-usual is measured by real income, $\frac{Y_i}{P_i^C}$ (Arkolakis, Costinot, and Rodríguez-Clare 2012). I measure changes in dollar-denominated welfare using equivalent variation (Shapiro 2016):

$$EV_i = \frac{Y_i}{P_i^C}\left(\frac{\hat{Y}_i}{\hat{P}_i^C} - 1\right)$$

4 Empirical strategy

The model identifies the key parameters governing deforestation decisions. I estimate two sets of parameters: (i) the distribution of land use elasticities $\gamma(\omega)$, identified from land use shares via Equation (3), and (ii) the cross-country trade elasticity σ^A, identified from trade shares via Equation (7). The land use elasticity determines how responsive deforestation is to changes in agricultural returns. The trade elasticity σ^A determines how readily consumers substitute across countries of origin in response to relative price changes.

4.1 Landowner's problem

Deriving empirical land use shares. The focal estimation on the supply-side concerns the landowner's problem, Equation (3). To derive estimating equations, I assume the error terms $\epsilon^h(\omega)$ are distributed Type-I Extreme Value, where $h \in \{F, A, O\}$.[9] The shocks have location parameter

9. This assumption limits estimation to land with interior solutions and extrapolates to the full sample. Appendix Table A.17 suggests interior-share land does tend to be more populated, higher yield, and closer to market. Conceiving

0 with scale parameter $\frac{1}{\gamma(\omega)}$.

The Type-I Extreme value assumption gives rise to the following closed-form land use shares for plot ω:

$$\mu^h(\omega) = \frac{\exp[\pi^h(\omega)]^{\gamma(\omega)}}{\sum_l \exp[\pi^l(\omega)]^{\gamma(\omega)}}$$

Next, I derive a linear regression from nonlinear choice probabilities (Hotz and Miller 1993). I take the example of forest-versus-agriculture: derivation of other land use comparisons follow suit. Define $Y^{hl}(\omega) = \log \mu^h(\omega) - \log \mu^l(\omega)$. Taking logs of the above choice probability, and differencing across land uses gives:

$$Y^{AF}(\omega) \coloneqq \log \mu^A(\omega) - \log \mu^F(\omega) = \gamma(\omega)[\pi^A(\omega) - \pi^F(\omega)]$$

Cropland-versus-forest decision. I parameterize the regression for the crop-forest tradeoff as:

$$Y^{AF}(\omega) = \gamma(\omega)[p_i - \gamma_\tau \tau(\omega)]\eta^A(\omega) + \upsilon_i - \gamma(\omega)\xi^F(\omega) \tag{13}$$

where υ_i is a country fixed effect absorbing wages, and $\gamma_\tau \tau(\omega)$ represents transport costs as a function of travel time $\tau(\omega)$. I measure the variable $\tau(\omega)$ in two ways. First, I use travel time in hours to cities using Google Maps query data from Nelson et al. (2019). I dollarize travel times for a subset of countries using freight rates in Appendix C.1. In this case, the regressor $\tau(\omega)$ is in dollar terms. Second, I use an interaction of straight-line distance to city and yields (yielding a variable denominated in t-km per ha). In this case, γ_τ is unrestricted and is dollarized using the response of deforestation to the price-yield interaction (which is in \$/ha/year).

Identification and instruments. Estimation of Equation (13) is biased if, within countries, agricultural productivity or distance to market correlates with unobserved forest amenities $\xi^F(\omega)$.[10] Distance to market raises transport costs, reducing agricultural profitability. But causation may also run in reverse: land with poor agricultural fundamentals (or high forest amenity value) attracts less infrastructure investment, resulting in greater distance to markets. This channel attenuates $\gamma(\omega)$. I require an instrument which shifts the relative rental rate of natural land and non-natural (agricultural) land.

I employ two instruments based on physical geography. First, I calculate distance to major cities and ports whose locations have not changed in the last 20 years (Souza-Rodrigues 2019). Second, I calculate the share of grid cells with slope above 15 degrees, a novel approach within forest land use studies. Globally, at high resolution (90 m), 18% of forest occurs on land above

of many shocks within a 10 × 10 km grid cells is reasonable: the average farm is smaller than 5 ha, or 0.005% of grid cell area (Foster and Rosenzweig 2021). I consider a sample selection correction in Appendix C.1.1.

10. Classical price endogeneity is resolved using a country-level fixed effect which absorbs aggregate unobserved supply. As long as farms are price takers, price is as-good-as-randomly-assigned within-country.

15 degrees in slope (Lundbäck et al. 2020), while prior work has demonstrated that housing and agriculture tend to drop off dramatically at 15 degrees (see Appendix Figure A.1 and Saiz 2010; Harari 2020). The exclusion restriction assumes slope affects land use only through agricultural costs, not through forest amenities directly. In Appendix C.1, I report the first stage with stepwise controls. I demonstrate robustness to hyper-local variation using a spatial differences design to alleviate concerns about remaining unobservables.

As a robustness check, I also alternatively identify explicitly off of differences in price. I describe the instrument in Appendix C.4. Price variation domestically comes from a tariff variation abroad. I weight tariffs abroad by historical trade shares, producing a shift-share-type instrument. Exclusion requires that tariffs abroad do not reflect supply shocks domestically. I provide suggestive evidence of this supply channel in the Appendix.

Other land uses versus forest. Here, I calibrate returns to match the data. Recall that the returns to other land use are exactly an amenity:

$$Y^{OF}(\omega) = \gamma(\omega)\xi^{O}(\omega) - \gamma(\omega)\xi^{F}(\omega)$$

Given estimates of $\gamma(\omega)$ from the forest-agriculture margin, I recover $\xi^{O}(\omega) - \xi^{F}(\omega)$ as the residual that rationalizes observed other-vs-forest shares.

4.2 Demand parameters and trade barriers

I estimate σ^A from agricultural trade flows. I calibrate the manufacturing cross-country elasticity of substitution to $\sigma^M = 4$ (Bernard et al. 2003; Simonovska and Waugh 2014), and the outer nest elasticity of substitution between agriculture and manufacturing to $\sigma = 0.5$ (a homothetic approximation of Comin, Lashkari, and Mestieri 2021).

Cross-country elasticity of substitution. Following 8, the logged expenditure share of country j consumers on the country i variety of agriculture A is given by:

$$\log \lambda_j^{i|A} = (1 - \sigma^A) \log p_i^A + \delta_j^A + \epsilon_{ij}^A$$

where trade barriers and idiosyncratic preferences $(1 - \sigma^A) \log \frac{T_{ij}^A}{\theta_{ij}^A}$ are grouped into the residual, and the importer-by-crop fixed effect δ_j^A is proportional to the aggregate agricultural price index P_j^A of country j.

OLS estimation of the above demand is biased by unobserved taste shocks θ_{ij}^A. For example, a sudden increase in taste for the agricultural goods from i in j would raise $\lambda_j^{i|A}$, but would simultaneously drive up prices p_i in the origin country by bidding up prices on the margin. I leverage country-level average potential yields η_i^A as a shifter of agricultural production which is

plausibly exogenous relative to demand. Identification requires that 1980 climate and soil-derived productivity in i is not driven by taste shocks in j as measured in data from 2000 onwards.

In practice, I estimate this specification pooling across panel years in my data. To be consistent with the cross-sectional model, I use importer-by-year fixed effects. Within any one cross-section, I recover similar point estimates.

5 Empirical Results

5.1 Land use model

Global deforestation elasticity. Table 2, top panel, presents my preferred estimates of the average deforestation elasticity. Standard errors are clustered at the tile level (100 grid cell rectangular grids, following Araujo, Costa, and Sant'Anna 2025). Appendix C.1 shows the first stage of the distance-to-city instrument. Interpreting the central elasticity, raising agricultural returns by \$3.54/ha raises agricultural land use shares by 1.6%. Comparable estimates are similar (Souza-Rodrigues 2019; Araujo, Costa, and Sant'Anna 2025). Appendix Table A.4 shows estimates are robust to alternative measures of transportation costs. Appendix C.1.1 addresses potential sample selection bias from estimating only on plots with interior land use shares.

Heterogeneous elasticities. A single deforestation elasticity is useful for capturing aggregate average magnitudes, but heterogeneity in the elasticity can greatly affect welfare costs of a carbon tax and the ability to reallocate food production across space. Table 2 , bottom panel, presents results from an EM estimator allowing for three latent landowner types with different elasticities. Type probabilities vary at the country-by-ecoregion level. Details of the estimation procedure are provided in Appendix Section C.2. Revenue elasticities have a coefficient of variation of 66%: standard deviations are large relative to the mean.

Ignoring this heterogeneity and projecting counterfactuals using only the average elasticity will result in biased conclusions. The optimal carbon tax and its welfare implications depend explicitly on the covariance between elasticities and emissions. Appendix Table C.2 shows that the distribution of elasticities is correlated with observables. For example, the Brazilian Amazon has more tenuous property rights institutions which make price incentives less effective (visible in Appendix Figure A.4). This finding echoes a similar result from Dominguez-Iino (2023). These covariances affect the incidence and, under some cases, the level of the optimal tax (see Appendix A.1).

Low deforestation elasticities also capture heterogeneity stemming from timber production: Northern Canadian Boreal forest and Russian forest both tend to be less elastic regions, as they tend to produce timber. From an identification standpoint, this result suggests I am not inaccurately loading on timber profitability when computing agricultural returns; it is purely heterogeneity. For counterfactuals, low elasticities in timber producing regions imply lower reallocation away from such regions.

Reconciling different local average treatment effects. The tariff shifter instrument

(Appendix C.4) results in a lower elasticity on average than spatial first differences and the cross-sectional estimation. With heterogeneous elasticities across landowners, each instrument produces a different local average treatment effect. To reconcile these elasticities, I compare the tariff IV estimates with the results for each heterogeneous landowner type in Table C.2. The tariff IV estimate corresponds to the lowest-elasticity type (column 1), which is also the most prevalent. Tariff variation is most relevant to landowners with higher transportation costs on the margin – \$1 per hectare in cost relative to \$0.42/ha of the next type. In this way, landowners with lower cropland share and higher transport costs – arguably, those further from agricultural clusters – are first on the margin when tariffs rise.

5.2 Demand estimation

In Table 3, I report my main findings for the values of the agricultural elasticity of substitution across country varieties – σ^A respectively. Columns (3) and (4) shows that log-linear IV and PPML-IV estimates of σ^A fall between 3-5. Log-linear IV estimates may be biased due to dropped zero trade flows (Fally 2015). First stages are in the expected direction: increasing agricultural productivity reduces prices. Going forward, I use the PPML-IV estimate.

5.3 Model inversion

I invert the model to obtain values of the preference parameters $\vec{\theta}^M$, price indices, and wages. I require data on: (1) trade and consumption shares, (2) agricultural prices. Because I estimate a global distribution of the revenue elasticity $\gamma(\omega)$ and demand parameters σ^A, I invert the model at 100 global bootstrapped draws of each parameter and resolve counterfactuals across each instance of the inverted model. This produces standard errors for all counterfactual results while respecting global spatial correlation in land use elasticities.

Given estimates of $\gamma(\omega)$ as inputs, I invert the land use decisions $\mu^h(\omega)$ to get amenities $\xi^F(\omega), \xi^O(\omega)$. For 0 and 1 land shares, I use a tolerance of 1^{-6}. In Appendix C.1.1, I demonstrate that there is no evidence of statistical bias resulting from sample selection for interior land shares, though I slightly overstate land reallocation on land which is at 0 or 1.

Next, I guess a value of country-level GDP which will rationalize the data. Given a vector of GDP, the quantity of agricultural output which rationalizes observed agricultural prices is:

$$Q_i^A = \frac{\lambda_i^A Y_i}{p_i}$$

With quantities in hand, I calibrate productivity shifters a_i as:

$$a_i = \frac{Q_i^A}{\sum_{\omega=1}^{N_i} \eta^A(\omega)\mu^A(\omega)}$$

I calibrate agricultural wages so that the agricultural wage bill equals a share e_i of agricultural value added, where e_i is the labor share reported by FAO.

$$w_i^A L_i^A = e_i p_i Q_i^A$$

Given an agricultural wage and profit vector, manufacturing wages are the solution to the linear system in Equation (10):

$$w_i^M = [I_n - \lambda^M]^{-1} \lambda^M \overbrace{\left(\frac{w^A L^A + \Pi}{L^M}\right)}^{\substack{\text{vector of } A \text{ value added per} \\ M \text{ worker}}}$$

positive wage vector exists when diagonal elements of λ^M (domestic manufacturing shares) are less than one.[11] This inversion finds the manufacturing wage which rationalizes cross-country manufacturing spending patterns λ^M with a fixed, known income per worker in each country. The wage gap in equation (11) is thus set implicitly as $l_i = \frac{w_i^M}{w_i^A}$.

Finally, I verify that the implied GDP from solving this system and the guessed GDP vector are identical. Otherwise, I resolve the system using the implied GDP from solving this system as the updated guess for GDP. This algorithm is quite stable and converges after a second iteration.

Preference parameters are obtained by combining the equilibrium prices, wages, with data on expenditure shares as in equations (5) and (6). For countries with 0 bilateral trade flows, I fix the associated preference parameters to be 0. Price indices are identified up to scale. θ_{ij}^A are set to match agricultural price indices to agricultural prices. $\theta_{ij}^M, \theta_j^h$ match manufacturing and consumption price indices to 1.

Finally, I set the marginal emissions cost κ of deforestation-based CO_2 to \$3.80, an annualized value of a \$190/ton social cost of carbon at a discount rate of 2%. Because yields are per-annum, and GDP is measured per annum, this harmonizes units across a per ton SCC and measurements in dollars per year. This assumption is a simplification: climate damages are not necessarily a perpetuity nor are they constant over time.

In this inversion, country-level GDP and actual crop production in tons remain untargeted moments. I visualize model fit in Appendix Figure A.7. In the left papnel A.7a, I achieve a 74.5% correlation between model GDP and actual GDP data from the Penn World Table. The

11. This equation also provides an analytic rationale for needing a separate agricultural wage to rationalize the data. When agricultural wages equal manufacturing wages, this equation need not have a positive solution, as the equilibrium condition will then solve

$$w_i^M = [I_n \vec{L}^{M,-1} 1_{\{n \times n\}} - \lambda^M]^{-1} \lambda^M \Pi$$

For the data values of non-agricultural land use shares and manufacturing flows, I cannot recover a positive wage. Intuitively, this means that the per-hectare productivity of manufacturing is much higher than that of agriculture.

model overstates GDP variance among low-income countries. In the right papnel A.7a, I compare model-implied output to the total output in tons across crops. I achieve an R^2 of 0.59 without calibrating the residual a_i output shifters to match model-implied quantities. Incorporating this calibration step, the R^2 stays nearly constant at 0.58 but the slope of the fit of model on data goes from 0.56 to 0.84.

6 Results

6.1 Welfare-maximizing carbon tax

Optimal tax rate. Using a linear grid search over taxes, I find a unique, global welfare-maximizing single tax rate of \$8/ton. Table 4, Panel A reports consolidated results of the carbon tax. The reported 95% confidence intervals reflect uncertainty in the level and spatial distribution of land use and demand elasticities. The 95% confidence interval for the optimal tax rate is always higher than the marginal social cost of emissions because (1) manufacturing productivity (and thus wages) exceed agricultural productivity and (2) there are gains in agricultural productivity from reallocating food out of deforestation-intensive regions.

Aggregate emissions fall markedly. Emissions from deforestation fall by 63.1% (51.9, 76). I classify fully half of actual global deforestation emissions since 1985 with 95% confidence as inefficient. This result extends prior insights that Brazilian deforestation is inefficient and cheap to prevent (Souza-Rodrigues 2019; Araujo, Costa, and Sant'Anna 2025).

Aggregate crop production rises, driven by cross-country reallocation. Reallocating land through the tax increases global food production by 0.6% (0.1, 0.7). The US, Canada, and the EU experience the largest level increases in production; the Congo River Basin, China, Australia and Indonesia all reduce food production the most. Without any price adjustment, food production would have fallen instead by nearly 10% globally, driven by an 8% decrease agricultural land use.

Agricultural prices modulate partial equilibrium losses and reflect rising marginal costs of food in the tropics. For example, prices increase by \$21/ton or 8.5% of baseline prices in the Congo River Basin region, the largest price increase of any region I study. Indonesia is second in the magnitude of its price increase, experiencing only a 4.5% or \$12 increase. These large increases in price drive reallocation. Within countries, prices incentivize additional agricultural production on less dense but potentially worse (due to high transport costs, low yields, or high amenity) forested land. They also reflect the need to consume food from abroad, which is more expensive even if produced at lower cost due to trade barriers.

Thus, prices and quantities reflect two compositional changes: across- and within-country substitution in addition to the pure level effect. I produce a partial equilibrium but exact decomposition of food production changes into three categories to explore these mechanisms. Define $\bar{x}$ as the global average of a variable x and $\bar{x}_i$ as the within-country i average of the same:

$$dQ^A = \sum_{\omega=1}^{N} \overbrace{\Delta\mu^A(\omega)\bar{a}\bar{\eta}^A}^{\text{Level effect}} + \underbrace{\Delta\mu^A(\omega)(a_i\bar{\eta}_i^A - \bar{a}\bar{\eta}^A)}_{\text{Cross-country composition}} + \overbrace{\Delta\mu^A(\omega)(a_i\eta^A(\omega) - a_i\bar{\eta}_i^A)}^{\text{within-country composition}} \tag{14}$$

Panel A of Table 5 reports this partial equilibrium decomposition across bootstrapped land use and demand elasticities. The level component reflects a simple intuition: under a tax on deforestation emissions, agricultural land use falls under most reasonable spatial distributions of the land use elasticity. Thus, reallocation towards more productive land is required to produce the positive food production result.

On net, land use reallocates towards higher-yield land both across and within countries. Notably, the confidence intervals provide a strict ordering of the different sources of reallocation. Cross-country reallocation – where for example, carbon taxes induce the US grows a larger share of food than it did under business as usual – is a lesser share of total reallocation than within-country reallocation. When deforestation is taxed efficiently, higher prices induce further agricultural entry on higher yield land. Because I adopt a revealed preference framework, this high-yield land is rationally not used to produce in business-as-usual due to higher transportation costs or amenities from forest cover. However, by reducing the supply of lower-yield but cheaper land on the forest frontier, it becomes more attractive to produce elsewhere.[12]

Panel A applies the same decomposition to emissions. Consistent with food production, emissions reductions are primarily driven by reallocating towards low emissions land – compositional changes in where production happens – rather than merely reducing agricultural land supply (a level effect).

Both decompositions occur in partial equilibrium. Partial equilibrium decompositions understate the role of cross-country reallocation, as absent the opportunity to reallocate across countries prices will change on the margin. In Subsection 6.2, I provide a general equilibrium decomposition of emissions which builds on this insight.

Finally, trade plays an important role in supporting cross-country reallocation. Tropical basins are pushed towards net importing of food products. Indonesia and the Congo River Basin increase their imports by 25-30% from top trading partners (the base rate of agricultural trade is relatively small, but this is still a large increase). Brazil is less impacted, increasing imports by 1.8-3% with its top trading partners. Reliance on food production from the US, Canada, the EU, and China increase. Countries on the periphery of tropical basins tend to export more – Venezuela, Bolivia, and Myanmar are particular examples.

Cropland reallocates from forest basin interiors to forest peripheries. Figure 2a illustrates changes in the cropland share at pixel scale. Cropland reallocates a fair amount. There are very few countries where cropland uniformly declines everywhere, which is the partial equilibrium prediction of the model. For example, within Brazil, deforestation falls in the Amazonian interior, but stays stable or even intensifies in the cerrado and Atlantic Forest regions. These regions are on average

12. Endogenizing the transportation network should further strengthen this result, as transport costs in these alternate regions should fall as demand for them rises.

less carbon dense per unit agricultural output.[13] Similarly in other major forest basins, the interior forests of the Congo River Basin and Borneo are both too expensive to deforest. Substitution in Sub-Saharan Africa moves to East African forests, while in SE Asia deforestation increases in Myanmar, Malaysia, and Vietnam.

Land classified as other does not experience any reallocation on average. Two choices affect this finding. First, the returns to other land use are held fixed at their amenity value. In practice, land returns to the other land use are also affected by equilibrium wages changing. Wages changing would favor substitution into agriculture, putting some pressure on urban areas. Second, other land use purposely captures cattle ranching which I leave unmodeled and untaxed in this framework due to a lack of data. In principle, such land also creates externalities.

At the pixel level, most landowners benefit from a carbon tax on deforestation. Figure 2b maps out the welfare implications for agriculturalists. Globally, 20% of landowners see falling inclusive values; the remaining 80% benefit from the policy. The average landowners' inclusive value rises by \$1.91/ha, with a 95% confidence interval of (\$0.45, \$7.08), while emissions-weighted profits fall by \$34.88/ha (\$110, \$17). Comparing these changes in returns to the geography of deforestation elasticities from Appendix Figure A.4, it is clear that dispersion in the scale parameter is critical in predicting regional variation in inclusive values.

Landowners in India, Northern China, and the Sahel benefit. These regions tend to have lower price elasticities and low-carbon natural vegetation. They experience the tax largely through increases in price. The variance of their unobservable shock ϵ^h is much larger (equivalently, they have a smaller elasticity of deforestation). Then, small increases in price are associated with larger welfare gains through selection on the unobservable shocks $\epsilon^h(\omega)$. In contrast, Midwestern farmers, Northern Mexican farmers, and Eastern European farmers are also beneficiaries of the policy to a lesser extent. Because this latter group of producers is more elastic, welfare changes are closer to the absolute level of expected profits.

Welfare benefits are large, but costs are regressive. Table A.19 summarizes production and welfare consequences. Annual real income rises slightly as a consequence of the Pigouvian tax, driven by the wedge in sectoral wages.[14] Canada, Europe, and the US are major beneficiaries of the policy.

Agricultural landowners face on average a positive welfare benefit of the Pigouvian tax. Recall from Section 2.2 that few landowners in land share terms carry out the bulk of deforestation. These few landowners are indeed marginal to the tax and face much higher costs of production. However, most global agricultural landowners face no downside from a Pigouvian tax because they are not on marginal land. Prices respond on the margin, and rise. Wages do not rise by as much, as the tax

13. Currently, the cerrado and Amazon biomes produce different mixes of agricultural output. Crops are not perfectly substitutable to consumers (Costinot, Donaldson, and Smith 2016). However, the gap between biome-level returns is \$36/ha. To surmount this gap would require a relative price increase for goods grown in the Amazon of \$27.70 per ton or 16% of business-as-usual prices.

14. This is verifiable by running the model with only the price clearing condition and setting both wages to 0. In this case, the optimal tax rate is exactly the rate of marginal damages and real income falls under taxation.

is targeted at land, generating a new wedge in labor and land returns. Thus, the average farmer earns higher prices, pays lower *relative* wages, and pays little or no tax.

Aggregate income growth is positive within each group due to the wage distortion between manufacturing and agriculture. However, positive net gains hide heterogeneity across countries. In Appendix Table A.18, the tax amounts to a transfer from laborers to inframarginal landowners. Lower agricultural labor demand puts downward pressure on aggregate wages. Rising agricultural prices erode the real value of these wages. Combining these two forces, workers in the hardest-hit tropical countries are worse off.

6.2 How important is cross-country reallocation of deforestation?

I next consider how much of the welfare benefit of the Pigouvian tax comes from reallocating agricultural production as opposed to levels of emissions reductions. I fix the share of emissions in a region, either a country or a country-by-biome pair, to remain at their business-as-usual level. I then solve for a vector of country, or country-biome specific, tax rates which maximizes global welfare. The full planning problem is explained in Appendix Section A.1.3. In that appendix, I describe my solution algorithm for country and country-biome level tax rates. I find the optimal solution through a numerically accurate, gradient-free approach.[15]

Without cross-country reallocation of deforestation, the average country-level tax is \$19/ton. Panel B of Table 5 reports results. 24% of emissions reductions from the optimal global carbon tax require cross-border reallocation of land use. As evidenced by the higher tax rate, this lower level of emissions reductions also comes at higher marginal cost per unit emissions. With 95% confidence, further, preventing cross-country reallocation of emissions implies a strict loss of food production. Prior partial equilibrium estimates understate cross-country reallocation because prices rising in one country will induce greater reallocation in equilibrium.

Distributionally, the quota provides insight into who winners and losers from the business-as-usual allocation of deforestation, taking as given a global value of the carbon externality. With the country-level emissions constraint in place, Mexico and regions in sub-Saharan Africa – Zimbabwe, Nigeria, and Mozambique as the clearest examples – deforest less than they would under the global optimum. On the other hand, Malaysia, the Congo River Basin, and Indonesia deforest far more than is optimal.

Figure 3 maps the country-level changes in emissions. The densest forest cover in the Amazonian interior, Borneo interior, and Congo Forest interiors is always avoided. Deforestation on the periphery of these major forest basins is harder to prevent absent cross-country intervention. In Brazil, the entire cerrado region and parts of major navigable waterways in the Amazon Basin rely 60-100% on reallocation of emissions abroad. Similarly, nearly all of Indonesia remains deforested save the interior of Borneo when other countries cannot pick up emissions slack. In emissions

15. I hold fixed the tax rate of Mali, Namibia, and Mauritania at the average tax rate. A small number of land parcels govern agricultural output in these countries, causing fixed point methods to fail. Per Appendix Table A.16, the share of global deforestation emissions in each country is less than 0.1%.

terms, the Congo River Basin reduces deforestation the least when constrained relative to the unconstrained optimum. Deforestation reductions in the Congo are nearly entirely reliant on substitution to East African producers.

I next constrain reallocation within countries, across biome regions. Introducing a constraint at the country-by-biome level instead of the country level leads to an average country-level tax across country-biomes of $7 per ton. The constrained-optimal tax reduces global emissions by a mere 1.1% relative to business-as-usual, or 1.8% of the optimal global unconstrained tax. A global tax of $7 per ton would reduce far more emissions than this country-biome-specific tax: the limited emissions reductions come entirely from the constraint, not the lower level of achievable tax rate. The limited remaining emissions reductions come from the deepest interior of the three tropical forest basins: the Amazon, Indonesia, and the Congo River Basin. Each biome can only reduce emissions by at most 5% relative to the optimal global carbon tax benchmark. All remaining emissions require reallocation (e.g., from the Amazon to the cerrado as described above). As shown in Panel C of Table 5, these few emissions come at very little (less than 0.1%) of global food and real income cost.

The country-biome level tax rates reveal where current deforestation is the most constrained relative to the optimum. Corroborating the story above, Brazilian Tropical & Subtropical Grasslands, Savannas & Shrublands – the cerrado region – have a 33% higher shadow value of carbon emissions than Brazilian Tropical & Subtropical Moist Broadleaf Forests. Both shadow values are low relative to constraints in the Chinese temperature and tropical forests – both double the constraint in the cerrado – or Peruvian forest – ten times the cerrado constraint. European forest is also similarly more taxed.

Existing economic literature, consistent with my reduced form evidence, suggests preventing emissions in dense forested regions can be achieved at low cost. For example, Souza-Rodrigues (2019) suggests that $1 per ton in the Amazon is sufficient to prevent a majority of realized emissions. I expand on this point: major reallocation is required to support food production.[16] Even focusing within Brazil alone, reallocation of agriculture across biomes to the cerrado is a necessary component of preventing further degradation of the Amazon. This result also suggests that investments in agricultural production in less deforestation-intensive countries has an environmental dividend when it competes directly with deforestation-intensive production in the tropics.

6.3 Does granular biomass measurement matter?

Thus far, I have demonstrated the central role of granularity in the welfare incidence of a carbon tax on producers and the importance of within-country reallocation for deforestation prevention in equilibrium. Granularity also matters in a simpler, public finance sense. Better accounting of an externality means fewer distortions from taxation.

16. In general, the role of reallocation would be even larger with the need to reallocate cattle production (untaxed here). Local complementarities between cattle production and production of feedcrops could make reallocation harder, and thus reduce equilibrium reallocation.

Researchers and policymakers cannot consistently observe individual landowners' emissions. Offset markets, for example, require detailed measurements of trees' diameter at breast height to award carbon credits, and even these measurements are able to be gamed. In the IPCC framework, used to assign offset project baseline credits prior to field testing, projects are assigned a carbon stock based on a country-by-biome average parameter taken from major ecological studies.

In this section, I take these measurement constraints as given and binding. I assume the planner taxes based only an aggregated regional measure of carbon intensity. I apply the same tax rate as the global optimum, but assume the planner cannot verify plot-level emissions.

When the planner taxes based on country-biome average emissions, 8.5% of baseline emissions which should be avoided at the optimal tax are not avoided. This result is reported in Panel B of Table 4. Geographically speaking, drawing on Figure 4, increases in emissions relative to the optimum are concentrated in two regions: first, temperate forests in Europe and the US; and second, subtropical forests and grasslands in Brazil and Indonesia. These regions have low average biomass but have high within-biome biomass variance. In contrast, tropical forests in Brazil and the Congo reduce emissions by more than is optimal. In these forests, the variance cuts in the opposite direction: highly dense clusters of tropical forest raise the average tax rate on remaining forest. To maintain higher stringency on tropical forests, leakage to nearby grassland regions also increases deforestation in these regions.

6.4 Calculating biodiversity co-benefits from carbon taxes

In this section, I modify the global welfare criterion to allow for biodiversity to impose an externality at the rate κ^B. Modeling a global value of biodiversity is stylized and should not be interpreted as an economic primitive. Indeed, biodiversity values vary across species and geography. While biodiversity may offer certain global public goods (medical innovation), it is also certainly local. However, biodiversity objectives have been coordinated at global scale since the Ramsar Convention through land protection efforts (Grupp et al. 2023). Thus, quantifying co-benefits of carbon-based deforestation policy for biodiversity is a valuable exercise for thinking about joint policies for land protection and payments for ecosystem services.

To quantify co-benefits of forest-related carbon policy on biodiversity, I consider a perturbation approach. I hold the carbon tax rate fixed at the optimal tax rate under the original objective function. I introduce a second instrument, a biodiversity price t^b. The tax is denominated in dollars per threatened animal habitat. For example, a potential hectare of agricultural land with three threatened species faces an additional fixed cost of $3t^b$. I re-optimize global welfare to determine the single welfare-maximizing price on biodiversity t^b, which enters the landowners' problem akin to the original carbon tax:

$$\pi_\epsilon^A(\omega, t_i^{SCC}, t^b) = \pi_\epsilon^A(\omega, t_i^{SCC}, 0) - t^b b(\omega)$$

In this framework, co-benefits of carbon for biodiversity protection arise from the spatial endowments of carbon and threatened species. If highly valuable threatened species are largely found in carbon dense regions, then marginal pricing of biodiversity loss will have limited effects on land use after carbon is already optimally priced. Thus, deviations from the baseline land use allocation will quantify how much substitution there is between carbon-optimal and jointly carbon-species optimal land allocations.

Panel C of Table 4 summarizes the impacts of optimal biodiversity taxation with a social value of $1 per animal habitat. The optimal policy reduces 36.8% of threatened species loss relative to business as usual. Notably, the policy trades off 12% of global emissions reductions in the carbon-only benchmark to achieve these biodiversity gains.

Figure 5 illustrates why this trade-off between biodiversity and carbon emissions must occur. Marginal pricing of biodiversity dramatically increases habitat loss prevention relative to a carbon-only tax policy. A $1 per animal-hectare biodiversity price reduces species loss from 16% of business-as-usual losses to 37% – more than doubling species protections.

The species-carbon frontier in the first panel of this figure is nonlinear. Initially, moving from a carbon-only world to protecting both carbon and biodiversity goals sacrifices some carbon. However, after an inflection point, both carbon and biodiversity goals appear complementary and more stringent protection of species also leads to stronger protection of forest carbon.

Some countries decrease emissions to meet biodiversity conservation goals; other increase carbon emissions. Animal conservation and carbon objectives are most complementary in the Amazon. At the most complementary, for every 1 ha of animal habitat saved, Peru avoids 0.57 t of carbon emissions. Brazil and Colombia also net save carbon emissions in the Amazon; the Congo River Basin also experiences complementarities. On the other hand, Malaysia and Indonesia give up $3 and $8 in carbon costs per 1 ha increase rare species habitat, respectively. Europe and the US also face tradeoffs in forest carbon sequestration and emissions.

7 Discussion

I summarize three takeaways from the quantitative exercises in this paper. First, deforestation is a granular problem driven by a tail of landowners. Measurement of this granular heterogeneity is important even for macro-scale programs targeting deforestation: when planners observe only country-biome average biomass, half of optimal emissions reductions are lost. Second, correcting deforestation can be a win-win policy. Agricultural production is misallocated across countries—high-productivity countries produce too little agriculture as a share of global production. A carbon tax on deforestation emissions gains efficiency both from pricing an unpriced carbon externality and through reallocation. When labor reallocates across sectors and out of agriculture, deforestation prevention becomes a pro-development policy. Third, while deforestation is a granular problem, unlocking these win-win benefits requires significant cross-country reallocation: 40% of optimal emissions reductions depend on shifting agricultural production across borders.

This reallocation disadvantages key actors in global deforestation policy. The US, Canada, and Europe benefit from the status quo allocation of emissions as much as major tropical basins. The countries most harmed by current deforestation patterns tend to be in sub-Saharan Africa. These distributional consequences make coordinated climate action less likely: wealthy non-tropical countries would need to expand agricultural output at the expense of higher-value-added sectors, while tropical countries bear the adjustment costs of reduced agricultural labor demand.

Deforestation policy benefits from complementary interventions that enable factor reallocation. In the tropics, deforestation reduction generates slack agricultural labor and excess land demand elsewhere through sharp price responses. Outside the tropics, expanded agricultural production is necessary to reduce pressure on carbon and biodiversity hotspots. Large-scale payments for ecosystem services—including proposals like the Tropical Forest Forever Facility—address only the compensation margin. My results point to a need for complementary policies that support labor reallocation within tropical countries and agricultural investment in substitute regions. The experience of American trade policy offers a cautionary parallel. Autor, Dorn, and Hanson (2013) document the arguably unforeseen (Bombardini, Li, and Trebbi 2023) distributional costs of trade liberalization. Future deforestation policy requires careful design to avoid these same pitfalls.

References

Abman, Ryan, and Clark Lundberg. 2020. "Does Free Trade Increase Deforestation? The Effects of Regional Trade Agreements." Publisher: The University of Chicago Press, *Journal of the Association of Environmental and Resource Economists* 7 (1): 35–72.

Adamopoulos, Tasso, and Diego Restuccia. 2014. "The Size Distribution of Farms and International Productivity Differences." *American Economic Review* 104 (6): 1667–97.

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

Alix-Garcia, Jennifer, Juliano Assunção, Teevrat Garg, Prakash Mishra, and Fanny Moffette. 2025. "Tradeoffs and synergies for agriculture and environmental outcomes in the tropics." *Forthcoming, Review of Environmental Economics and Policy.*

Alvarez, Fernando, and Robert E. Lucas. 2007. "General equilibrium analysis of the Eaton–Kortum model of international trade." *Journal of Monetary Economics* 54 (6): 1726–1768.

Angelsen, Arild. 1999. "Agricultural expansion and deforestation: modelling the impact of population, market forces and property rights." *Journal of Development Economics* 58 (1): 185–218.

Araujo, Rafael, Francisco J. M. Costa, and Marcelo Sant'Anna. 2025. "Efficient Forestation in the Brazilian Amazon: Evidence from a Dynamic Model." *Review of Economic Studies* 93 (1): 72–105.

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare. 2012. "New Trade Models, Same Old Gains?" *American Economic Review* 102 (1): 94–130.

Assunção, Juliano, Clarissa Gandour, and Romero Rocha. 2023. "DETER-ing Deforestation in the Amazon: Environmental Monitoring and Law Enforcement." *American Economic Journal: Applied Economics* 15 (2): 125–56.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103 (6): 2121–68.

Balboni, Clare, Aaron Berman, Robin Burgess, and Benjamin A. Olken. 2023. "The Economics of Tropical Deforestation." Publisher: Annual Reviews, *Annual Review of Economics* 15 (1): 723–754.

Bernard, Andrew B., Jonathan Eaton, J. Bradford Jensen, and Samuel Kortum. 2003. "Plants and productivity in international trade." ISBN: 0002-8282 Publisher: American Economic Association, *American economic review* 93 (4): 1268–1290.

Bombardini, Matilde, Bingjing Li, and Francesco Trebbi. 2023. "Did US Politicians Expect the China Shock?" *American Economic Review* 113 (1): 174–209.

Borchert, Ingo, Mario Larch, Serge Shikher, and Yoto V. Yotov. 2021. "The International Trade and Production Database for Estimation (ITPD-E)." *International Economics* 166:140–166.

Carleton, Tamma, Levi Crews, and Ishan Nath. 2023. *Agriculture, Trade, and the Spatial Efficiency of Global Water Use.* Working Paper.

Clausing, Kimberly A., and Catherine Wolfram. 2023. "Carbon Border Adjustments, Climate Clubs, and Subsidy Races When Climate Policies Vary." *Journal of Economic Perspectives* 37 (3): 137–62.

Comin, Diego, Danial Lashkari, and Martí Mestieri. 2021. "Structural change with long-run income and price effects." *Econometrica* 89 (1): 311–374.

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

Curtis, Philip G., Christy M. Slay, Nancy L. Harris, Alexandra Tyukavina, and Matthew C. Hansen. 2018. "Classifying drivers of global forest loss." Publisher: American Association for the Advancement of Science Section: Report, *Science* 361 (6407): 1108–1111.

Dominguez-Iino, Tomas. 2023. *Efficiency and Redistribution in Environmental Policy: An Equilibrium Analysis of Agricultural Supply Chains.* Working Paper.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, geography, and trade." Publisher: Wiley Online Library, *Econometrica* 70 (5): 1741–1779.

Fally, Thibault. 2015. "Structural gravity and fixed effects." *Journal of International Economics* 97 (1): 76–85.

Farrokhi, Farid, Eliot Kang, Heitor S. Pellegrina, and Sebastian Sotelo. 2023. *Deforestation: A Global and Dynamic Perspective.* Working Paper.

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

Foster, Andrew D., and Mark R. Rosenzweig. 2021. "Are There Too Many Farms in the World? Labor Market Transaction Costs, Machine Capacities, and Optimal Farm Size." *Journal of Political Economy* 130 (636-680): 000–000.

Fowlie, Meredith, Stephen P. Holland, and Erin T. Mansur. 2012. "What Do Emissions Markets Deliver and to Whom? Evidence from Southern California's NOx Trading Program." *American Economic Review* 102 (2): 965–93.

Fowlie, Meredith L. 2009. "Incomplete Environmental Regulation, Imperfect Competition, and Emissions Leakage." *American Economic Journal: Economic Policy* 1 (2): 72–112.

Gillingham, Kenneth, and James H Stock. 2018. "The cost of reducing greenhouse gas emissions." Publisher: American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203-2418, *Journal of Economic Perspectives* 32 (4): 53–72.

Gouel, Christophe, and David Laborde. 2021. "The crucial role of domestic and international market-mediated adaptation to climate change." *Journal of Environmental Economics and Management* 106:102408.

Grupp, Tristan, Prakash Mishra, Mathias Reynaert, and Arthur Van Benthem. 2023. *An Evaluation of Protected Area Policies in the European Union.* Technical report. Working Paper.

Harari, Mariaflavia. 2020. "Cities in Bad Shape: Urban Geometry in India." *American Economic Review* 110 (8): 2377–2421.

Harstad, Bård. 2020. *Trade and Trees: How Trade Agreements Can Motivate Conservation Instead of Depletion.* CESifo Working Paper. Publisher: CESifo Working Paper.

———. 2023. "The Conservation Multiplier." Publisher: The University of Chicago Press, *Journal of Political Economy* 131 (7): 1731–1771.

Harstad, Bård, and Torben K Mideksa. 2017. "Conservation contracts and political regimes." Publisher: Oxford University Press, *The Review of Economic Studies* 84 (4): 1708–1734.

Hidalgo, F. Daniel, Suresh Naidu, Simeon Nichter, and Neal Richardson. 2010. "Economic Determinants of Land Invasions." *The Review of Economics and Statistics* 92 (3): 505–523.

Hotz, V. Joseph, and Robert A. Miller. 1993. "Conditional Choice Probabilities and the Estimation of Dynamic Models." *The Review of Economic Studies* 60 (3): 497–529.

Hsiao, Allan. 2025. *Coordination and Commitment in International Climate Action: Evidence from Palm Oil.* Working Paper.

"Emissions Trends and Drivers." 2023. In *Climate Change 2022 - Mitigation of Climate Change: Working Group III Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change,* 215–294.

Knight, Katie L., et al. 2022. "Error propagation in pollen-based land-cover reconstruction." *The Holocene* 32:1435–1447.

Lagakos, David, and Michael E. Waugh. 2013. "Selection, Agriculture, and Cross-Country Productivity Differences." *American Economic Review* 103 (2): 948–980.

Lundbäck, Mikael, Henrik Persson, Carola Häggström, and Tomas Nordfjell. 2020. "Global analysis of the slope of forest land." *Forestry: An International Journal of Forest Research* 94 (1): 54–69.

Malhi, Yadvinder, et al. 2012. "Tropical forests in the Anthropocene." *Annual Review of Environment and Resources* 37:205–239.

Nath, Ishan. 2025. "Climate Change, the Food Problem, and the Challenge of Adaptation through Sectoral Reallocation." *Journal of Political Economy* 133 (6): 1705–1756.

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

Perino, Grischa, Robert A Ritz, and Arthur A van Benthem. 2025. "Overlapping Climate Policies." *The Economic Journal* 135 (671): 2122–2160.

Phillips, Dom. 2019. "Bolsonaro Declares 'the Amazon Is Ours' and Calls Deforestation Data 'Lies.'" *The Guardian.*

Restrepo, Verónica, and Gabriel Leite Mariante. 2024. *Does Conservation Work in General Equilibrium?* Working Paper.

Saiz, Albert. 2010. "The Geographic Determinants of Housing Supply." *The Quarterly Journal of Economics* 125 (3): 1253–1296.

Scott, Paul T. 2013. *Dynamic Discrete Choice Estimation of Agricultural Land Use.* Working Paper.

Shapiro, Joseph S. 2016. "Trade Costs, CO2, and the Environment." *American Economic Journal: Economic Policy* 8 (4): 220–54.

Simonovska, Ina, and Michael Waugh. 2014. "The elasticity of trade: Estimates and evidence." *Journal of International Economics* 92 (1): 34–50.

Skiba, Katarzyna. 2023. "Europe's Oldest and Largest Forest Is Now A Major Political Battleground." *WorldCrunch.*

Song, Xiao-Peng, Matthew C. Hansen, Stephen V. Stehman, Peter V. Potapov, Alexandra Tyukavina, Eric F. Vermote, and John R. Townshend. 2018. "Global land change 1982-2016." *Nature* 560 (7720): 639–643.

Sotelo, Sebastian. 2020. "Domestic Trade Frictions and Agriculture." *Journal of Political Economy* 128 (7): 2690–2738.

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

Tombe, Trevor. 2015. "The Missing Food Problem: Trade, Agriculture, and International Productivity Differences." *American Economic Journal: Macroeconomics* 7 (3): 226–58.

Train, Kenneth E. 2009. *Discrete choice methods with simulation.* 2nd ed. Cambridge University Press.

———. 2008. "EM Algorithms for nonparametric estimation of mixing distributions." *Journal of Choice Modelling* 1 (1): 40–69.

UNFCCC. 2013. *Estimation of carbon stocks and change in carbon stocks of trees and shrubs in A/R CDM project activities.* Technical report v04.1. United Nations Framework Convention on Climate Change.

Tables and Figures

Table 1: Summary statistics of the joint distribution of potential yields and forest carbon sequestration.

	Mean (SD)	Q10	Q25	Q50	Q75	Q90
Global Distribution						
Carbon biomass, t/ha	21.37 (30.87)	0.20	0.80	7.64	28.66	60.63
Wheat yields, kg/ha	628.31 (826.34)	0.00	0.00	47.33	1086.67	2067.94
Maximum yields, kg/ha	1821.03 (1562.46)	0.00	0.00	2032.33	3274.50	3812.11
Distance to a city, km	452.40 (552.67)	53.06	111.75	247.97	538.40	1149.89
Global Distribution: Deforested > 10% 1982-2016						
Carbon biomass, t/ha	30.03 (19.48)	9.76	15.95	25.58	39.67	56.38
Wheat yields, kg/ha	689.27 (650.99)	0.00	0.00	715.08	1191.78	1512.28
Maximum yields, kg/ha	3077.74 (1159.75)	1292.25	2684.67	3349.42	3828.08	4109.33
Distance to a city, km	327.91 (309.40)	80.30	144.17	252.36	406.03	646.13
Tropics only: Deforested > 10% 1982-2016						
Carbon biomass, t/ha	40.09 (19.68)	16.36	24.22	36.87	53.87	68.29
Wheat yields, kg/ha	329.09 (596.51)	0.00	0.00	0.00	587.25	1351.94
Maximum yields, kg/ha	3195.30 (418.31)	2893.56	2985.00	3037.64	3245.22	3823.00
Distance to a city, km	530.81 (233.53)	286.00	401.52	520.43	661.08	748.14
Treecover (p.p.)	1982	25.42		2016	24.16	
Biomass (tC/ha)	1982	21.37		2016	61.06	

NOTES: Summary statistics for the distribution of yields and carbon emissions from deforestation. Deforested > 10% limits attention to land whose vegetation cover falls by 10% in the VCF data (area share = 3.34%). Tropics limits to parcels which have a tropical or sub-tropical eco-region of any form in the World Wildlife Fund classification schema (area share = 0.19%). Distance to city refers to cities in the Akbar et al. (2023) data. Maximum yields take point-wise maximum across crops in my analysis: oilpalm, wheat, maize, soybeans, rice, and potatoes (sweet or white). All statistics weighted by area. (p.p.) refers to variables measured in percentage points.

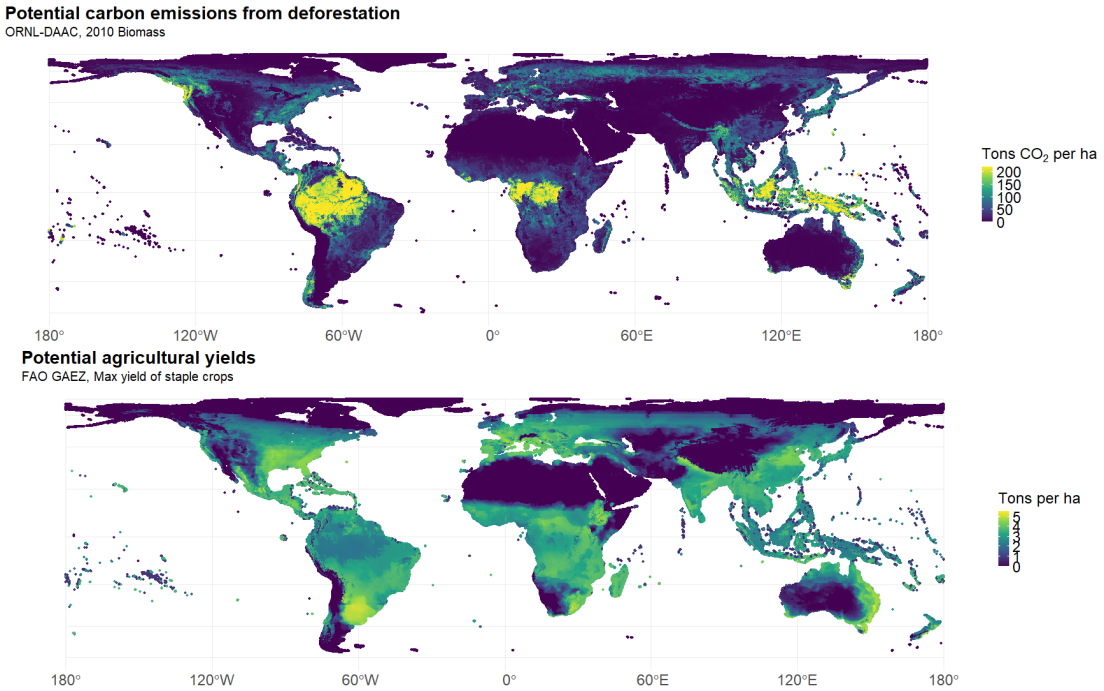

Figure 1: Map of potential CO_2 emissions from deforestation (top) and agricultural productivity (bottom), each in megatons.

NOTES: Top figure plots the biomass extracted from the ORNL-DAAC data in 2010. Bottom figure plots an agricultural potential yield index, corresponding to the maximum of the yield of staple crops discussed in Appendix B. All quantities are scaled in tons per hectare.

Table 2: Supply-side: estimation of the landowners' problem.

Parameter	Estimate (Std. Error)	Standard deviation (Std. Error)
	No heterogeneity	
$\gamma(\omega)$	0.0099 (0.0016)	
Elasticity equivalent	1.652 (0.26)	
Cost per km	-0.20 (0.03)	
	Discrete mixture heterogeneity	
$\gamma(\omega)$	0.0088 (0.0046, 0.0155)	0.0055 (0.0020, 0.0087)
Elasticity equivalent	1.273 (0.705, 2.065)	1.885 (1.006, 2.957)
Cost per km	-0.6273 (-2.0130, -0.2394)	0.2708 (0.0000, 1.3008)

NOTES: No-heterogeneity standard errors clustered at the tile (100 × 100 grid cells) level. Heterogeneous confidence intervals use 75 draws of a Bayesian bootstrap procedure clustered at the tile level. Linear cost per km standard errors via the delta method. Standard deviations in the heterogeneous case calculated by taking standard deviation of grid cell-level, type-weighted average coefficients. Because plot-level cost per km a ratio of travel time and revenue coefficients, I limit to $\gamma(\omega) \geq 10^{-4}$ to prevent overflow.

Table 3: Estimates of Agricultural Trade Elasticity σ

	OLS	Poisson	IV	PPML-IV
$\log(p)$	-0.418	-0.002	-2.270	-3.801
	(0.156)	(0.003)	(1.842)	(0.066)
Observations	78,844	127,116	78,844	197,888
First-stage F	—	—	1907.6	1907.6
First-stage Coefficient	—	—	-0.150	-0.150
σ^A	1.418	1.002	3.270	4.801

NOTES: Dependent variable is trade share (log for OLS/IV, level for Poisson/PPML-IV). All specifications include importer×year fixed effects and cover years 2001–2016. OLS and Poisson estimate via least squares with log price as regressor. IV uses log wheat yields as an instrument. PPML-IV estimates via GMM with a log wheat yields as the instrument. Standard errors clustered by exporter and importer (two-way) for OLS/Poisson/IV; clustered by exporter-importer pair for PPML-IV. First-stage F coefficient for a regression of log prices on log wheat yields. The target parameter, σ^A, is given by $1 - \beta$ where β is the estimated coefficient.

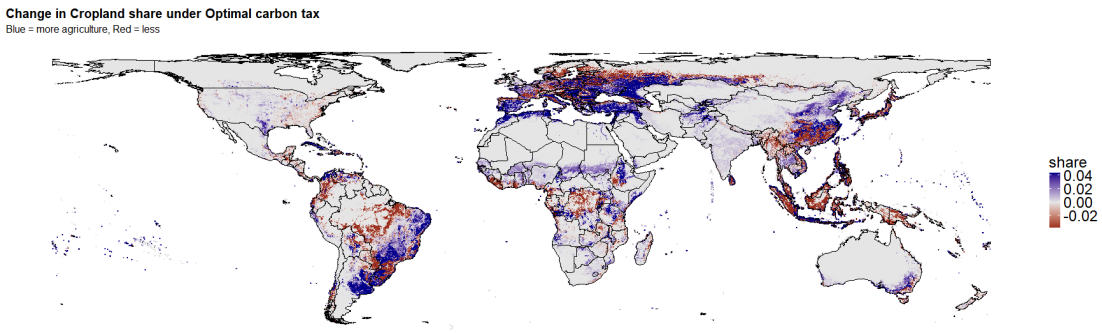

(a) Change in cropland land use shares under optimal tax

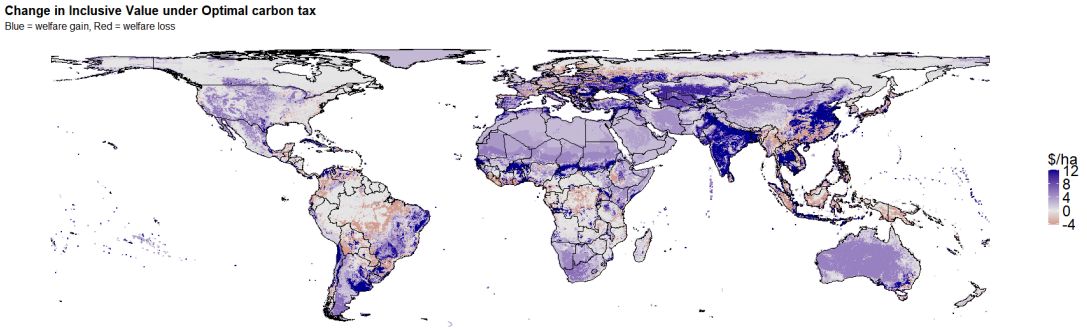

(b) Landowner welfare: inclusive value of land uses

Figure 2: Changes in land use and landowner welfare from optimal deforestation carbon tax.

NOTES: Figure illustrates changes in cropland share (top) and changes in landowner inclusive value (bottom) in the taxed equilibrium relative to business-as-usual (BAU). Inclusive values are deflated by country-level real price index P_i^C. Both color scales are censored at the 5th and 95th percentiles of changes for visual clarity.

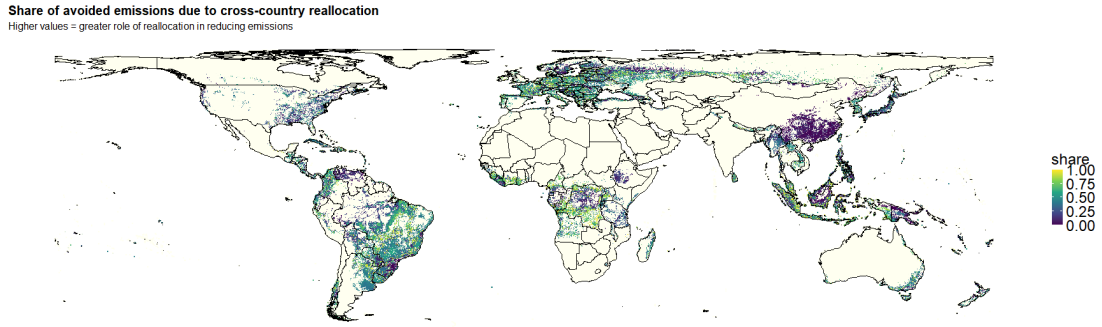

Figure 3: Map of the share of the optimal CO_2 emissions prevented when cross-country reallocation of emissions is not possible.

NOTES: Both optimal carbon tax and cross-country reallocation of emissions computed with a social cost of carbon of $190. No cross-country case maximizes social welfare with a country-level tax rate subject to the constraint that countries emit the same share of global deforestation emissions as emitted 1982-2016. Regions which avoid less than 0.01 t / ha under the global optimal Pigouvian tax are recoded as NA for visual clarity.

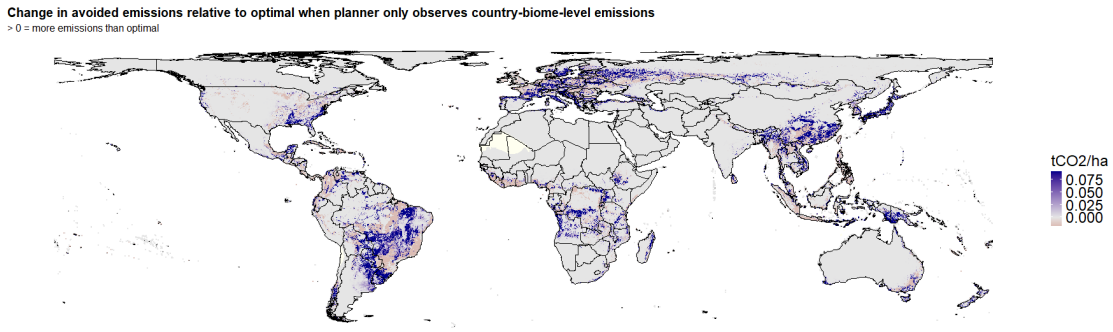

Figure 4: Map of the change CO_2 emissions relative to the global optimum when planner only observes country-biome-level biomass.

NOTES: Optimal carbon tax computed with a social cost of carbon of $190. Carbon emissions changes are winsorized at the 5th and 95h percentiles for visualization.

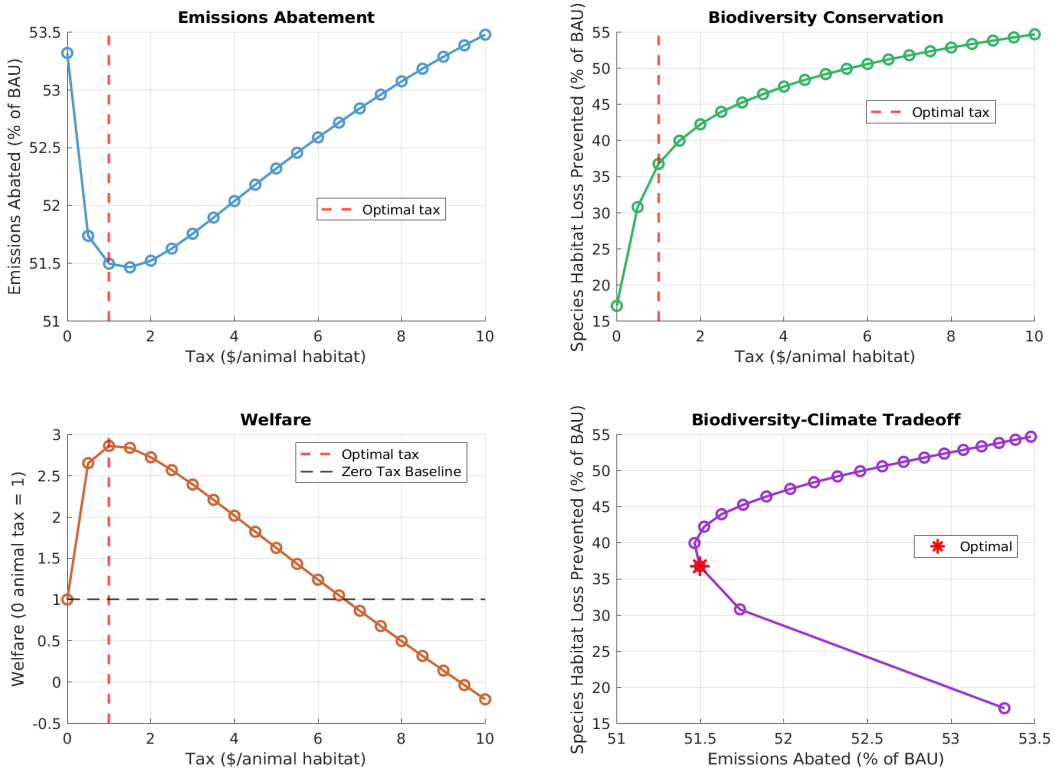

Figure 5: Simulation results for carbon emissions and species habitat loss under increasing marginal tax rates on habitat loss. Optimality at a $1 per animal value of habitat.

NOTES: Habitat loss measured as habitat of threatened species in the IUCN Red List data. Welfare gain indicates gains under a $190 present-value of carbon and $1 present-value of habitat. Emissions and species loss prevented are relative to business-as-usual.

Table 4: Deforestation policy counterfactual simulation results

Metric	Estimate	95% CI
Panel A: Optimal Carbon Tax Policy		
Optimal Carbon Tax ($/ton CO_2)	8.0	[4.0, 21.3]
Emissions Reduction (%)	63.1	[51.9, 76.0]
Food Production Change (%)	0.6	[0.1, 0.7]
Net Welfare Change (% GDP)	0.05	[0.03, 0.09]
GE/PE Response Ratio	1.08	[1.06, 1.10]
Panel B: Targeting Counterfactuals		
Country-Level Aggregation		
Emissions Reduction (%)	48.5	[39.0, 54.9]
Efficiency vs Pigouvian (%)	76.8	[69.6, 79.7]
Country-Biome Aggregation		
Emissions Reduction (%)	57.8	[38.8, 66.4]
Efficiency vs Pigouvian (%)	91.5	[76.7, 100.0]
Panel C: Biodiversity Co-benefits		
Optimal Pigouvian Tax ($/animal habitat)	1.00	[0.00, 1.50]
Species Habitat Loss Prevented (% of BAU)	36.8	[10.3, 45.2]
Emissions Abated (% of BAU)	51.5	[27.9, 64.1]

NOTES: This table reports bootstrap estimates and bias-corrected 95% confidence intervals for carbon tax policies. All % changes are relative to that quantity in business as usual. *Panel A* shows results for the globally optimal uniform carbon tax. GE/PE Response Ratio compares food production in general equilibrium to that in a counterfactual with fixed prices. *Panel B* presents targeting counterfactuals where policymakers enforce the optimal tax rate, but incidence is according to aggregated carbon stock data. *Panel C* shows the optimal tax on habitat loss at a $1/animal habitat valuation of species loss. Bootstrap confidence intervals computed using Efron (1987) bias-correction method with 88 draws.

Table 5: Decomposing level and reallocation effects of taxing deforestation carbon emissions

Panel A: Decompositions (% of Total)		
Metric	Estimate	95% CI
Production Change - Levels	-76.6	[-97.8, 36.5]
Production Change - Cross-Country	50.0	[11.3, 62.3]
Production Change - Within-Country	126.7	[115.6, 150.7]
Emissions Change - Levels	-2.3	[-2.9, 2.0]
Emissions Change - Cross-Country	-15.9	[-17.6, -12.9]
Emissions Change - Within-Country	-81.8	[-86.7, -79.3]
Panel B: Country Quota Allocation		
Mean Tax Rate ($/ton CO_2)	19.0	[10.0, 30.0]
Tax Range (Max-Min, $/ton)	110.6	[18.5, 279.8]
Emissions Reduction (%)	48.0	[34.3, 55.5]
Emissions Reduction vs Optimal (%)	76.0	[54.2, 87.8]
Food Production Change (%)	-0.7	[-0.9, -0.3]
Panel C: Country-Biome Quota Allocation		
Mean Tax Rate ($/ton CO_2)	0.5	[0.3, 1.1]
Tax Range (Max-Min, $/ton)	5.4	[2.6, 14.9]
Emissions Reduction (%)	1.1	[1.0, 1.2]
Emissions Reduction vs Optimal (%)	1.8	[1.5, 2.4]
Food Production Change (%)	-0.0	[-0.0, -0.0]

NOTES: This table reports bootstrap estimates and bias-corrected 95% confidence intervals for decompositions of the role of reallocation. All % changes are relative to that quantity in business as usual, unless listed as "vs Optimal" in which case relative to Panel A of Table 4. *Panel A* decomposes production and emissions changes into three partial equilibrium components according to Equation (14). *Panel B* reports general equilibrium results for the optimal country-specific carbon tax which is constrained maintains baseline country-level deforestation shares. *Panel C* reports general equilibrium results for optimal country-biome specific carbon taxes constrained to maintain baseline country-biome deforestation shares. Bootstrap confidence intervals computed using Efron (1987) bias-correction method with 88 draws.

Appendix: The Global Allocative Efficiency of Deforestation

A Theoretical appendix

A.1 Social welfare

Global welfare consists of two components: consumer surplus and emissions. Define $V_i(Y_i)$ as the value function of the representative consumer in i; the solution to their utility maximization problem in Section 3. The value function is a function of total income in the economy Y_i, which will implicitly depend on tax rates and revenue recycling rules. Carbon tax funds in the European Union Emissions Trading System and carbon markets in the United States are earmarked for domestic causes, so I assume revenues are distributed lump sum within-country. The parameter κ represents the rate of marginal damages: the cost to social welfare of increasing emissions by 1 ton of carbon dioxide. I assume quasilinearity in emissions for simplicity. Global social welfare is characterized by the following optimization problem over an $\Omega \times 1$ dimensional vector of tax rates t on carbon. I will show that optimal taxes are either constant globally or constant within countries.

$$W = \max_{t} \sum_{i=1}^{J} \psi_i V_i(Y_i) - \kappa D \tag{A.1}$$

$$\text{s.t.} \quad Y_i = \mathcal{Y}_i + \mathcal{T}_i \tag{A.2}$$

$$\mathcal{T}_i = \sum_{\omega=1}^{N_i} d(\omega)\mu^A(\omega) \tag{A.3}$$

$$\mathcal{Y}_i = \Pi_i + \sum_{h \in \{A,M\}} w_i^h L_i^h$$

where ψ_i is a set of Pareto weights and where $D = \sum_{\omega=1}^{N} d(\omega)\mu^A(\omega)$ is total global emissions from deforestation. I ex ante restrict the planner to be unable to change tariff rates across countries to focus ideas on the tax as the key allocative mechanism.[1] I present results abstracting from general equilibrium and terms of trade effects.

The remainder of this appendix is organized as follows. Appendix Subsubsection A.1.1 derives key identities regarding the effect of a marginal tax rate on emissions and income, abstracting from terms of trade effects. The key result regarding optimality of a global Pigouvian tax is delivered in Appendix Subsubsection A.1.2. Carbon taxes can deviate from the usual rule of setting tax rates to the marginal damage rate. In my model, either a redistribution motive on the part of the planner or a wage distortion will generate deviations from the usual rule ($w_i^A \neq w_i^M$, as discussed in Tombe 2015; Nath 2025).

1. While I refer to a planner here, I do not solve a centralized planning problem. I instead consider a fictitious coordinated global planner who can implement the stated Pigouvian policy in a free market equilibrium. The planner's solution differs from market clearing in (9) and (10) because the planner solves a public goods coordination problem and maximizes global welfare, whereas individual landowners are maximizing plot-level returns.

A.1.1 Derivation of welfare differentials.

Global emissions elasticity. Global emissions respond to a global tax through the land use margin.

$$\frac{\partial D}{\partial t} = \sum_{\omega=1}^{N} d(\omega) \times \frac{\partial \mu^A(\omega)}{\partial t} =: -\sum_{\omega=1}^{N} d(\omega) e_d(t;\omega)$$

Derivation of change in value function with respect to a change in tax. By the chain rule:

$$\frac{\partial V_i}{\partial t} = \frac{\partial V_i}{\partial Y_i}\frac{\partial Y_i}{\partial t}$$

The first term in this expression is just the marginal utility of income, given by the real price index $\frac{1}{P_i^C}$. The second term, leveraging Equation (A.2), is:

$$\frac{\partial Y_i}{\partial t} = \sum_{\omega=1}^{N_i} \left[\overbrace{(w_i^M - w_i^A) e_d(t;\omega)}^{\text{Labor reallocation}} - \underbrace{e_\pi(t;\omega)}_{\text{Producer surplus loss}} \right] + \frac{\partial \mathcal{T}_i}{\partial t}$$

This result has three components. First, income is rising whenever the manufacturing wage is larger than the agricultural wage. This is the empirically relevant case (Tombe 2015). In my setting, the labor reallocation channel is tied to land reallocation through the Leontief assumption, so the land use elasticity appears in this term as well. Second, income is falling due to eroding landowner profits which have an absolute elasticity $e_\pi(t;\omega)$ to a tax rate. For the focal case of Type-I extreme errors in the landowner profit function, if $y(\omega) = \frac{1}{\gamma(\omega)} \sum_{h\in\{F,A,O\}} \exp[\gamma(\omega)\pi^h(\omega)]$, then

$$e_\pi(t;\omega) = \left| \frac{\partial y(\omega)}{\partial \pi^A(\omega)} \frac{\partial \pi^A(\omega)}{\partial t} \right| = \mu^A(\omega) d(\omega).$$

In the lump-sum redistribution case, the transfer $\mathcal{T}$ is first-order linear in the tax rate:

$$\frac{\partial \mathcal{T}}{\partial t} = D_i + t\frac{\partial D_i}{dt}$$

A.1.2 The optimal global carbon tax.

The planner's problem reduces to an unconstrained maximization problem. The objective function has the following derivative in a single global tax rate.

$$\frac{\partial W}{\partial t} = \sum_{i=1}^{J} \frac{\psi_i}{P_i^C}\left[\frac{\partial \mathcal{Y}_i}{\partial t} + D_i + t\frac{\partial D_i}{\partial t}\right] - \kappa\frac{\partial D}{\partial t}$$
$$\therefore t^* = \left[\sum_{i=1}^{J} \frac{\psi_i}{P_i^C}\frac{\partial D_i}{\partial t}\right]^{-1}\left(\kappa\frac{\partial D}{\partial t} - \sum_{i=1}^{J} \frac{\psi_i}{P_i^C}\left[\frac{\partial \mathcal{Y}_i}{\partial t} + D_i\right]\right)$$

Recovering Pigou. The usual Pigouvian rule is satisfied under a precise set of conditions. The planning problem must place Pareto weights $\psi_i = P_i^C$. The wage distortion across agriculture and manufacturing must be 0 so that $w_i^M = w_i^A$. Under these conditions, the second term cancels out as $\frac{\partial \mathcal{Y}_i}{\partial t} = -D_i$:

$$\begin{aligned}\frac{\partial \mathcal{Y}_i}{\partial t} &= \sum_{\omega=1}^{N_i}\left[(w_i^M - w_i^A)e_d(t;\omega) - e_\pi(t;\omega)\right] \\ &= \sum_{\omega=1}^{N_i} -e_\pi(t;\omega) \\ &= \sum_{\omega=1}^{N_i} -\mu^A(\omega)d(\omega) = -D_i\end{aligned}$$

In this case, $t^* = \kappa$ and we recover the Pigouvian rule.

Relaxing Pareto weighting. With a general vector of Pareto weights ψ_i, type I extreme value errors, but no wage distortion so that $w_i^M = w_i^A$, the second term above will still cancel. However, the coefficient on marginal damages is no longer 1 in the optimal tax formula. In this setting, differences in the elasticity of deforestation are not neutral across countries because the planner has some preference for redistribution. The weight on marginal damages depends on a covariance term.

$$t^* = \kappa \cdot \frac{\sum_{\omega=1}^{N} d(\omega)\epsilon_d(t)}{\sum_{i=1}^{J} \frac{\psi_i}{P_i^C}\sum_{\omega=1}^{N_i} \epsilon_d(t)d(\omega)}$$

Wage distortions. Finally, with cross-sector wage distortions, taxes will be higher when manufacturing wages on average exceed agricultural wages:

$$t^* = \left[\kappa + \sum_{\omega=1}^{N}\left[d(\omega) + \frac{\psi_i}{P_i^C}(w_i^M - w_i^A)\right]e_d(t;\omega)\right]\left[\sum_{i=1}^{J}\frac{\psi_i}{P_i^C}\sum_{\omega\in\Omega_i}\epsilon_d(t)d(\omega)\right]^{-1}$$

Country-level carbon taxes. If the planner can set carbon taxes separately by country, they choose a vector of tax rates which is again linear in the marginal damage of a ton of carbon

emissions.

$$t_i^* = \left[\frac{\psi_i}{P_i^C}\frac{\partial D_i}{\partial t_i}\right]^{-1}\left(\kappa\frac{\partial D}{\partial t_i} - \frac{\psi_i}{P_i^C}\left[\frac{\partial \mathcal{Y}_i}{\partial t_i} + D_i\right]\right)$$

As before, type-I extreme value errors and an absence of wage distortions remove the second term as $\frac{\partial \mathcal{Y}_i}{\partial t_i} = -D_i$. Thus, under these conditions, the planner is indifferent between a country-specific and global carbon tax.

A.1.3 No-reallocation benchmark.

The no-reallocation benchmark modifies the country-level carbon tax problem by fixing country-level emissions shares to their no-tax baseline. Define $s_i^d(t) = \frac{D_i(t)}{D(t)}$ as the share of global deforestation emissions at a tax rate t. The constrained problem is:

$$\begin{aligned} W = \max_t & \sum_{i=1}^{J} \psi_i V_i(Y_i) - \kappa D \\ \text{s.t.} \quad & Y_i = \Pi_i + \sum_{h\in\{A,M\}} w_i^h L_i^h + \mathcal{T}_i \\ & s_i^d(t) = s_i^d(0) \quad \forall i \in \{1, 2, \ldots, J\} \end{aligned} \tag{A.4}$$

The optimal tax rate will depend on the shadow cost of the J emissions constraints in Equation (A.4), given by Lagrange multipliers κ_i^d. The reallocation-constrained optimal tax $\tilde{t}_i$ relates to the unconstrained tax as:

$$\tilde{t}_i = t_i^* + \kappa_i^d \frac{\partial s_i^d(t)}{\partial t}$$

Solving for country-specific tax rates is a fairly high-dimensional problem. With terms-of-trade and general equilibrium effects, other countries' emissions constraints would also appear in the expression for the optimal tax. These second-order relationships across tax rates can make it quite difficult to solve for the optimal tax rate. I dramatically reduce the dimensionality of the problem by solving for the *sum* of the country-specific tax vector which maximizes welfare, $\sum_{i=1}^{J} \tilde{t}_i$. At every sum of country-specific taxes, I find a fixed point in prices, wages, and the tax rate which satisfy both market clearing and the emissions constraint. This approach produces the full set of feasible solutions for any welfare criterion, reducing the number of function evaluations dramatically when comparing across multiple scenarios.

A.2 Identification of counterfactual equilibria

This section derives market clearing conditions for an equilibrium with a Pigouvian tax of t on agricultural deforestation.

Market clearing in the counterfactual equilibrium. The initial market clearing conditions, copying equations (10) and (9) are:

$$p_i^k Q_i^k = \sum_{j=1}^{J} E_{ij}^k$$
$$w_i L_i^M = \sum_{j=1}^{J} E_{ij}^M$$

These same conditions must hold in the new equilibrium. Denote the new equilibrium outcomes by a prime.

$$p_i^{'} Q_i^{A'} = \sum_{j=1}^{J} E_{ij}^{A'}$$
$$w_i^{'} L_i^{M'} = \sum_{j=1}^{J} E_{ij}^{M'}$$

Denote by the hat $\hat{x} = \frac{x'}{x}$ the ratio of the value of x in the new ($'$) and old equilibria. Beginning with the first equilibrium condition, I divide both sides by the starting value of production in agriculture to obtain hat changes:

$$\begin{aligned}
\hat{p}_i \hat{Q}_i^A = \frac{1}{p_i Q_i^A} p_i^{'} Q_i^{A'} &= \sum_{j=1}^{J} \frac{1}{p_i Q_i^A} E_{ij}^{A'} \\
&= \sum_{j=1}^{J} \frac{X_{ij}^A}{p_i Q_i^A} \hat{E}_{ij}^A \\
&= \sum_{j=1}^{J} \gamma_{ij}^A \hat{E}_{ij}^A
\end{aligned}$$

In the last line, γ_{ij}^A denotes the sales share of j in i's agricultural production in the *baseline equilibrium*. I can repeat this same exercise for the manufacturing clearing condition.

$$\hat{w}_i \hat{M}_i = \sum_{j=1}^{J} \gamma_{ij}^M \hat{E}_{ij}^M$$

The hat change in the expenditure $\hat{E}_{ij}^h$ is composed of three terms:

$$\hat{E}_{ij}^h = \hat{\lambda}_j^h \hat{\lambda}_j^{i|h} \hat{Y}_j$$

Hat-changes in outer nest shares are:

$$\hat{\lambda}_j^h = \left(\frac{P_j^{h'}}{P_j^{C'}}\right)^{1-\sigma} \left(\frac{P_j^h}{P_j^C}\right)^{\sigma-1} = \left(\frac{\hat{P}_j^h}{\sum_{h\in\{A,M\}} \lambda_j^h (\hat{P}_j^h)^{1-\sigma}}\right)^{1-\sigma}$$

By a similar token, hat-changes in inner nest shares are:

$$\hat{\lambda}_j^{i|h} = \left(\frac{p_i^{h'}}{P_j^{h'}}\right)^{1-\sigma^h} \left(\frac{p_i^h}{P_j^h}\right)^{\sigma^h-1} = \left(\frac{\hat{p}_i^h}{\sum_{k=1}^{J} \lambda_j^{k|h} (\hat{p}_k^h)^{1-\sigma^h}}\right)^{1-\sigma^h}$$

Finally, the change in income can be written as a weighted sum of changes in landowner income and wage worker income, where I define α_j as the share of income in country j from landowners in business-as-usual $\alpha_j = \frac{\Pi_j}{Y_j}$. Without explicit cross-country labor mobility or population growth, there is no change in the total potential labor supply L_j:

$$\hat{Y}_j = \alpha_j \hat{\Pi}_j + (1-\alpha_j)\hat{w}_j$$

Supply-side identification. Then, a final ingredient in identification focuses on changes in land use shares (an input into $\hat{M}_i$ and $\hat{Q}_i^A$) and land rents, $\hat{\Pi}_j$. To demonstrate identification, I will assume that the error term ϵ^h in the landowners problem, Equation (3), is distributed as a Type-I Extreme value: several other functional form assumptions yield qualitatively equivalent identification results (Chiong, Galichon, and Shum 2016). The scale parameter of shocks is $\frac{1}{\gamma(\omega)}$. Under this assumption, land use shares on plot ω aggregate over parcel-level shocks $\epsilon^h(\omega)$ as:

$$\mu^A(\omega) = \frac{\exp[\gamma(\omega)\pi^A(\omega)]}{\exp[\gamma(\omega)\pi^A(\omega)] + \exp[\gamma(\omega)\pi^F(\omega) + \exp[\gamma(\omega)\pi^O(\omega)]}$$

Then, profits can be linked directly to (interior) land use shares as:

$$\log \frac{\mu^A(\omega,0)}{\mu^F(\omega,0)} = \gamma(\omega)[\pi^A(\omega,0) - \pi^F(\omega,0)]$$

This same expression will hold for counterfactual land use shares and profits:

$$\log \frac{\mu^A(\omega, t_i^{SCC})}{\mu^F(\omega, t_i^{SCC})} = \gamma(\omega)[\pi^A(\omega, t_i^{SCC}) - \pi^F(\omega, t_i^{SCC})]$$

Then, plugging in profit functions, the level of initial land use, prices, and wages is required to solve for a counterfactual allocation.

A similar condition holds for changes in agriculture relative to the other land use. Conditional on identifying $\hat{\mu}$, agricultural production changes $\hat{Q}_i^A$, labor supply changes $\hat{M}_i$, and profit changes $\hat{Pi}_i$ all follow. Changes in agricultural quantities are:

$$\hat{Q}_i = \frac{\sum_{\omega=1}^{N_i} \eta^A(\omega)\mu^A(\omega, t_i^{SCC})}{\sum_{\omega=1}^{N_i} \eta^A(\omega)\mu^A(\omega, 0)}$$

The hat-change in manufacturing labor supply is given by:

$$\hat{M}_i = \frac{L_i - \sum_{\omega=1}^{N_i} \mu^A(\omega, t_i^{SCC})}{L_i - \sum_{\omega=1}^{N_i} \mu^A(\omega, 0)}$$

and finally, changes in aggregate profits are given by changes in the inclusive value:

$$\hat{\Pi}_i = \frac{\sum_{\omega=1}^{N_i} \frac{1}{\gamma(\omega)} \log[\sum_{k\in\{F,A,O\}} \exp[\gamma(\omega)\pi^{k'}(\omega)]]}{\sum_{\omega=1}^{N_i} \frac{1}{\gamma(\omega)} \log[\sum_{k\in\{F,A,O\}} \exp[\gamma(\omega)\pi^{k}(\omega)]]}$$

Summary. The counterfactual system can be solved without a normalization as there are $2 \times J$ conditions in the $2 \times J$ unknowns $(\hat{p}, \hat{w})$. Changes are purely identified off of the parameters $(\gamma(\omega), \sigma, \sigma^A, \sigma^M)$.

B Data appendix

In this section, I describe the process of collating a mix of remotely sensing and economic data into a grid cell level panel.

B.1 Sketch of data construction process

There are multiple types of data which must be synthesized to complete this project. This includes remote-sensing data, the bulk of which is presented in raster (point-cloud) format, administrative data on boundaries, roadways, and waterways which is presented in a vector (geometric) format, and economic data which is often tied to an administrative unit.

To harmonize all of these data sources, I define plots using a rectangular grid overlaid on the world map. Importantly, the raster data comes at varying resolutions, which means that my final

dataset is subject to aggregation bias if my resolution is too low, and can create an uninformative block-diagonal data structure if the resolution is too high. Table A.12 lists key raster data sources and describes their resolution.

I extract the data in 504, 10 × 10 arc-degree tiles in a WGS84 projection. Visualization and areas are computed using a CEA (equal-area) projection. Each tile is overlayed with 10,000 0.1 arc-degree sub-tiles. Finally, tiles are intersected with country terrestrial boundaries from the Global Administrative Boundaries database so that every tile fragment is firmly within a country boundary. Each stage is one of 504 tiles which are 10 degrees-by-10 degrees. Doing this ensures that the entire world is extracted in a timely manner. I overlay neighborhoods over these cells: neighborhoods are 1 degree-by-1 degree tiles. As is displayed in A.12, the motivating driver of this resolution is the GAEZ potential yield data, which is has a 9 km native resolution.

B.2 Land use data

Several satellite-based land use products targeted at measuring deforestation have emerged in recent years. I will rely on two such sources described below: the first, the European Space Agency ESACCI data, is used to estimate the model, and the second, the VCF project, is used to validate the model on long-run deforestation data.

First, the European Space Agency Climate Change Initiative, or ESACCI data, which covers a 1 kilometer × 1 kilometer grid annually from 1995-2019. The ESACCI data is valuable for its high temporal frequency and detailed landcover classification (see A.11 for details on the classification scheme), which I can aggregate into six land uses: forest, grassland, settlement, cropland, vegetation, and bare soil/other. However, ESACCI changes measurement regimes in 2016, making the last 5 years of the data spuriously different from the first 20 years. I drop these data from panel regressions.

Second, I use the vegetation continuous fields product. This "VCF" measure is taken at a 250 meter resolution derived from the MODIS satellite (Song et al. 2018). Prior work has used this data in the context of Indonesian palm oil (Hsiao 2025). It reports land shares of tree canopy, short vegetation, and bare areas. Unfortunately, short vegetation is not mapped one-to-one with cropland; it will also include grasslands and meadows. I have two cross-sections of this data, one indicating the level of tree cover in 2016, and one indicating the change in tree cover since 1982.

B.3 Productivity data

B.3.1 Emissions and biomass

I use a global cross-sectional map of aboveground and belowground woody biomass for the year 2010 from NASA's Oak Ridge National Laboratory Distributed Active Archive Center (ORNL-DAAC). This dataset is derived from LiDAR-calibrated satellite imagery and reports biomass at a native spatial resolution of 300 meters by 300 meters. Biomass is measured in megagrams of carbon per hectare (MgC/ha).

B.3.2 Agricultural productivity.

Agricultural productivity is measured using the FAO GAEZ maps, which have a 9-by-9 km resolution. I restrict attention to the subset of crops used to construct agricultural productivity and output measures in the paper: wheat, rice, maize, oil palm, soybeans, white potatoes, sweet potatoes, and sugar cane. Rice comes in two varietals in the GAEZ: dryland and wetland – I take their maximum in a pixel to compute a single rice productivity. Used extensively in prior economic research, this dataset provides information on the biophysical suitability of different crops and livestock in different regions of the world, as well as information on climate, soil, and other environmental factors (Adamopoulos and Restuccia 2022; Farrokhi and Pellegrina 2023; Costinot, Donaldson, and Smith 2016) The GAEZ reports potential yields for nearly 30 crops at two input (high- and low-input) and water availability (irrigated and rainfed) levels over 30-year epochs. Data for the main results of this paper come from high-input, rainfed yields for the climate epoch 1980-2010.

In the empirical exercise in Section 4, I require a single-dimensional crop index. I follow prior work in using a weighted sum of a subset of crop potential yields, with weights corresponding to land use shares in 1990. In a prior draft of this paper, I estimated results using the first principal component of crop-specific potential yields; results were similar.

B.3.3 Biodiversity measurement

Biodiversity is a notoriously difficult object to quantify. My measure counts up the potential habitat area of the list of threatened species on the International Union for the Conservation of Nature, or IUCN Red List. Potential habitat uses documented sightings of threatened species to form a suitable elevation range, land use range, and temperature and precipitation range which is suitable for life. These potential habitat ranges can adapt over time. I use the most recent vintage of this data, published in the year 2017, summing habitat across all represented species.

B.4 Prices

I aggregate the Food and Agriculture Organization series on producer prices for individual crops. I construct a price for agricultural output as:

$$p_{it} = \sum_k p_{it}^k \mu_{i,1991}^k$$

where $\mu_{i,1991}^k$ are land use shares of crop k in country i in 1991. Note this is exactly the aggregation scheme used for country-level yields. When testing this aggregation, average country-level prices are most strongly influenced by dropping wheat or soybean when crops are dropped one-at-a-time. Importantly, oilpalm, which I argue is a relevant piece of the final average, does not skew the average global price. Using fixed base-year weights ensures that time-series variation in p_{it} reflects

changes in prices rather than endogenous crop reallocation. The weighting year 1991 is chosen to maximize global coverage while preceding the bulk of deforestation and trade liberalization studied in the paper.

FAOSTAT price series exhibit several well-documented data quality problems. In particular:

1. Entry and exit: Countries enter the FAOSTAT price series at different times, especially in the early 1990s following the dissolution of the Soviet Union.

2. Flat price spells: Some country-crop series report identical prices for multiple consecutive years, despite substantial global price movements.

3. Spurious price jumps: Certain series exhibit large, isolated price changes that are not mirrored in global markets.

Absent correction, these features introduce attenuation bias and spurious volatility into the aggregate price index.

I apply a set of transparent, pre-specified cleaning rules to country-crop price series prior to aggregation. These rules are adapted from Lobell, Schlenker, and Costa-Roberts (2011) and applied symmetrically across all countries and crops. A country-crop price series p_{kit} is excluded from the construction of p_{it} if it satisfies any of the following criteria:

1. Flat price criterion: The series contains three or more consecutive years with no nominal price change. (1.08% of country-crop pairs)

2. Excess persistence criterion: The 10-year autocorrelation of the log price exceeds 0.4 or the 20-year autocorrelation exceeds 0.2. (22.7% of country-crop pairs)

Excluded series are treated as missing and re-normalized out of the price index by rescaling remaining weights to sum to one.

B.5 Production and trade flows

For the estimation of the demand model and trade barriers, I use data on international production quantities and trade flows from the International Trade and Production Database for Estimation, or ITPD-E (Borchert et al. 2021). The ITPD-E is a product of the United States Trade Commission assembled to allow for explicit estimation of gravity-type demand models. It includes flows in values of commodities at the Harmonized System (HS)-2 level for 1991-2021 at an annual frequency, between 265 countries. This amounts to 170 industries grouped into four broad sectors: Agriculture is used as the A sector for this paper. I omit the mining and forestry broad sectors (HS2 codes 27-35), so that there is a mapping from my manufacturing sector to the same in the ITPD-E. A key feature of this data, as compared to comparable trade datasets like CEPII BACI used in prior agricultural macroeconomics and trade research, is it limits the use of interpolation or

other statistical techniques. It also includes production data, which facilitates estimating domestic absorption.[2]

B.6 Tariffs

I rely on the Market Access Map, or MacMAP data, to construct a time-varying measure of tariffs in terms of their ad valorem equivalents. The data is provided at the HS6 level. I aggregate it up to the HS4 code for each of my crops. Details are provided in Appendix Table A.1.

Table A.1: Mapping of major crops to HS4 codes.

Crop name	**HS4 code**
Wheat	1001
Rice	1006
Maize	1005
Oilpalm	1207
Soybeans	1201
White potatoes	0701
Sweet potatoes	0714
Sugar cane	1212

B.7 Other data sources

There are several other data sources which are used throughout the paper. Appendix Table A.12 summarizes key raster and geospatial datasets, provides the spatial resolution of these data, and gives a brief description of each. Unless otherwise noted, I use travel time to cities from Nelson et al. (2019) as the key transportation cost control.

Country-level travel costs in dollars per ton kilometer are taken from a World Bank report which provides cross-sectional freight cost estimates aggregating across a number of studies in trade (Herrera Dappe, Lebrand, and Stokenberga 2024). Travel costs were extracted by providing Figures 1.6 and 1.7 to Chat GPT version 4.0 and asking it to extract the values of each bar. Extracted values were visually verified against the bar graph. Country-level estimates of the median cost per ton-km of truck-based freight are then converted to a dollar cost of freight travel. I assume travel times are set at 60 km/hr and multiply freight rates by travel time. Then, the dollar cost of truck freight is:

$$\text{freight dollar cost}, \tau^f(\omega) = \text{travel time, hours} \times 60\text{km/hr} \times \$ \text{ cost per ton-km} \times \eta^A(\omega), \text{ in t/ha}$$

I use a variety of data from the Food and Agricultural Organization (FAO). These data all are country-by-year panels for (at minimum) 1990-2021. Data include information on value added

2. In the agricultural sector, 57% of the 265 ITPD-E countries have at least one HS-2 code level industry's domestic production attached. In manufacturing this number drops to 32%.

per worker, value add shares of labor in agriculture and forestry activities (which I use directly as a measure of agricultural value add absent further forestry-specific data), gross production in thousands of dollars USD, pesticide use in kilograms per hectare, estimated emissions of carbon dioxide in kilograms per hectare, fertilizer usage for each of three chemicals (Urea/nitrogen, phosphates/phosphorus, and potassium/potash) in kilograms per hectare, land use and landcover shares including information on irrigation usage, and finally a suite of food security indicators. All data are publicly accessible via the FAOSTAT portal.

B.8 Global market access calculations

I calculate market access variables of the following form. Define population density $x(\omega)$. Then, population access for a given plot ω is given by:

$$\text{Population access}(\omega) = \sum_{\omega' \neq \omega} x(\omega)\delta(\omega, \omega')^{-1.96}$$

where $\delta(\omega, \omega')$ measures the geodesic distance in kilometers between the two grid cells and 1.96 is a distance elasticity (Duranton, Morrow, and Turner 2014). To compute this at global scale is a computationally punishing task as it requires calculating a 1.6 million-dimensional square matrix of distances. Instead, I limit attention to country-level computations. For all countries, this distance matrix is computable in this form except for Canada, the United States, and Russia. For these three countries, I use chunked processing to compute distance matrices across tiles. Each chunk comprises a subset of origins and the universe of destinations, limiting the number of computations held in memory at a given time.

B.9 Regional aggregation

I organize some countries into groups to manage the size of the equilibrium computation problem. Additionally, some countries have data irregularities which otherwise prevent computation. I construct an aggregated Europe which includes the formal European Union, the United Kingdom, Switzerland, and Serbia. I aggregate the Congo River Basin due to inconsistent data availability within the Basin: this is Cameroon, the Central African Republic, the Democratic Republic of the Congo, the Congo, Equatorial Guinea, and Gabon.

I perform two strict exclusions, meaning that the associated land area is dropped from the final analysis dataset. First, I exclude all GADM ISO3 codes corresponding to airports, military bases, research sites, and other non-sovereign territories. These are: XAD, XCL, XKO, XPI, ZNC, and Z01-Z09. Excluded areas have total physical area 306,152 sq. km. Next, I exclude eight countries in the GADM boundaries data who report 0 crop area in the ESACCI data as of 2000. Most are island nations with very small area. These are the Bouvet Islands, Heard & McDonald Islands, Iceland, Monaco, South Georgia & South Sandwich Islands, Svalbard & Jan Mayen, St. Pierre &

Miquelon, and Tokelau.

There are 161 countries with an area less than 300,000 square kilometers. Small countries are both relatively small as a share of global production and pose a computational issue, as very few physical grid cells will govern the totality of agricultural supply. Absent some smoothing, these countries can create "bang-bang" equilibrium behavior as taxes can end up entirely removing agricultural supply. I group these into a constructed "rest-of-world," or RoW henceforth. The largest excluded countries are the Philippines and New Zealand with approximately 296 and 269 thousand sq km of area respectively.

I additionally include in this rest-of-world: Korea, Japan, Algeria, Iraq, Jordan, Niger, Pakistan, Tunisia, Yemen, Sri Lanka, Angola, and Chad. These countries are a small share of global food production of all included crops and pose computational issues due to limited land area which is clearly able to be cultivated. Finally, Cote d'Ivoire, Turkmenistan, and Botswana lack crop-level production or land use data in FAOSTAT; Uganda and Liberia lack price data; Sudan, South Sudan, Belarus, and Costa Rica lack trade data; and Argentina, Guyana, Syria, and Uzbekistan lack consumption data.

To conduct this aggregation, for moments which are averages, I assign to the country aggregates (Europe, Congo River Basin, and Rest-of-World) a weighted average of available country-level data. For crop-level data, I use crop production quantities as weights. For aggregate data, I use real GDP as weights. For data which is at the field level (but not crop level), like aggregate agricultural shares, I weight by land area.

C Land use estimation appendix

C.1 Distance instrument

In this section, I develop an instrumental variables strategy which leverages long-run, persistent changes in agricultural profitability in the cross-section. Cross-sectional variation has the advantage relative to my tariff instrument in Appendix Section C.4 of being longer-run and thus better simulating responses to persistent shocks like a committed carbon tax.

Starting with the regression in (13), I measure transportation costs $\gamma_\tau \tau(\omega)$ using the freight cost measure $\tau^f(\omega)$ described in Appendix Section B.7. This yields a single right-hand side variable, net returns of deforestation:

$$Y^{AF}(\omega) = \gamma(\omega)(p_i \eta^A(\omega) - \tau^f(\omega)) - \gamma(\omega)\xi^F(\omega)$$

Regions which are more profitable for deforestation and less profitable as forest can also have better infrastructure. Then, the unobserved amenity $\xi^F(\omega)$ will confound an ordinary least squares of $\gamma(\omega)$. Better infrastructure lowers freight costs but is associated positively with the amenity; I thus expect an attenuation bias.

As discussed in Section 4, I leverage an instrument to resolve this identification challenge. My instrumental variables approach combines on prior work on Braziliian deforestation (Souza-Rodrigues 2019) and work on urban growth (Saiz 2010; Harari 2020). My instruments are straight-line distance to city and the share of a grid cell which has a steep slope (above 15 degrees). The first stage regresses a left-hand side measured in dollars on a combination of instruments:

$$p_i \eta^A(\omega) - \tau^f(\omega) = \beta_1 \text{distance to city}(\omega) + \beta_2 [\text{slope}(\omega) \geq 15^\circ] + u(\omega)$$

Relevance relies on the following argument: if land is shifted 1 km further from a city or is slightly more unforgiving in terrain, then it is relatively less profitable for agriculture than forest use exactly as mediated by freight costs. On steep slopes, while natural vegetation continues to grow up to slopes of 30 degrees, 15 degrees is sharply associated with a drop off in agricultural activity globally. Note that these terrain variables are relevant both through $\eta^A(\omega)$ – the base yield of agricultural land – and the land availability channel. As evidence, I present a visual reduced form in Appendix Figure A.1 which demonstrates a sharp response of land use to slope even after controlling for accessibility through city distance.

Exclusion here is largely threatened by a story about selection. In general, the identification argument relies on comparisons between a plot which is 1 km further away from a city, or 1p.p. more steep than its counterfactual comparison plot. If these comparisons select on plots which have fundamentally different mixes of forest cover or other forest attributes, then these unobservable forest amenities might still attenuate $\gamma(\omega)$ relative to its true value. I include controls for nearby population access, the population density of a grid cell, the share of non-forest-nor-agricultural land, and increasingly fine administrative fixed effects as shifters of forest amenity demand which

should not show up in my estimates of γ. Further, in Appendix Figure A.2, biomass – a plausible proxy of a forest amenity – is uncorrelated with steep slopes above a 5% share after controlling for selection. Below 5%, there is a slight positive correlation.

The first stage for locations with data on freight costs is shown in Appendix Table A.2. Across all specification tested, 1 km increase in distance to a city results in \$0.13 to 0.17 fewer dollars per hectare on average. Net returns are responsive to steep slopes, but the coefficient is imprecise until I include all controls, as net returns also affect returns to non-agricultural land use. Interpreting the coefficient on steep slopes in column (5), a grid with 1 p.p. higher land share of 15 degrees or more has a \$0.13 decrease in expected returns per hectare, or the equivalent of being 1 km further from market. The coefficient on distance to ports is not significantly different from the coefficient on distance to city, with similar overall magnitudes.

Taking these insights to the data, I estimate a two-stage least squares model which recovers a global average deforestation elasticity $\gamma(\omega) = \gamma$. Appendix Table A.3 shows the response of the logit log odds ratio to the instrumented value of agricultural production variable. The first-stage F-statistic is consistently strong. The implied average elasticity, depending on the precise choice of fixed effects, ranges between 1.7-5.6, which implies that a 1% increase in the value of agricultural production –\$3.33/ha – increases the share of croplands by 1.6-5.6%.

Because I lack data on a number of countries to dollarize transportation costs, I re-estimate the model with an implicit dollarization of transportation costs on the full data. This model is:

$$Y^{AF}(\omega) = \gamma(\omega)[p_i \eta^A(\omega) - \gamma_\tau \text{distance to city}(\omega)] - \gamma(\omega)\xi^F(\omega)$$

Thus, I instrument potential yields with distance to city and slope steepness, but control for a linear interaction of straight-line distance and yields. In this specification, dollarization comes purely from price and yields rather than transport costs. The first stage purely captures the differential potential yield response to steep slopes, as I explicitly partial out variation from straight-line distances. Despite the different sources of variation, estimates are similar but more stable, with the global deforestation elasticity in column (2) around 1.4. Marginal costs per ton-kilometer in this specification are around \$0.20 per km, which is lower than the average reported transportation cost in my dollarized data. Differences will reflect both measurement error in freight costs and functional form differences: my dollarized costs are purely linear in t-hours, whereas realized travel costs may in reality exhibit concavities further from market.

C.1.1 Sample selection

In this sectionn, I model agricultural entry following the discrete choice literature on firm entry and exit to account for land with 0 cropland shares (Bresnahan and Reiss 1991; Hopenhayn 1992). Entry costs are recovered from the participation decision using a two-stage approach (Heckman 1979), where the first stage estimates the probability of observing positive cropland. This model introduces a nonlinearity in entry costs for the first landowner on a grid ω relative to the continuum of landowners, creating local increasing returns to scale. The second stage re-estimates Equation (13) accounting for a selection adjustment.

The goals of this section are twofold. First, I validate that sample selection towards interior land use decisions does not bias my main estimation approach. Second, I provide a structural interpretation for controls as entry cost shifters (see Appendix C.1 for details on controls).

The first stage is a probit model which regresses the probability of planting any crop, $\mathbb{I}[\mu^A(\omega) > 0]$ on controls. Following my arguments in Appendix C.1, I control for population access and population density. Appendix Figure A.3 illustrates the strong relationship between positive cropland shares and population access. I allow the first stage to depend on the value of agricultural output to lend a dollar value interpretation to coefficients. Entry is identified off of variation in population access and the probit functional form assumption, generating the estimating equation:

$$\mathbb{I}[\mu^A(\omega) > 0] = \chi_1[p_i \eta^A(\omega) - \tau^f(\omega)] + \chi_2 \log(\text{population access}(\omega)) + \chi_3 \bar{\mu}^A_{c(\neg\omega)} + \chi_i + \chi_b + \chi_e(\omega)\chi(\omega)$$

where $\chi(\omega)$ is a residual entry cost, which is expressed in dollars as $\chi(\omega)/\chi_1$. $\bar{\mu}^A_{c(\neg\omega)}$ represents the average agricultural land share in the second-order administrative unit c in which the observation ω is located, excluding grid cell ω itself. I include country χ_i, biome χ_b, and ecoregion χ_e fixed effects.

The probability of entry from this probit model $\hat{p}^e(\omega)$ is used to calculate an inverse Mills ratio $\frac{\phi[\hat{p}^e(\omega)]}{\Phi[\hat{p}^e(\omega)]}$ where ϕ and Φ here refer to the normal density and distribution functions. The inverse mills ratio acts as a control in place of population access in Equation (13), where this regression now limits attention to entered grids $\mu^A(\omega) > 0$. I use my instruments from Appendix C.1.

Analyzing results of the entry probit in Appendix Table A.5, the average global cost of entering agriculture is \$92 (SE = \$11.1) per hectare. Entry costs on average are thus 44% of global average revenues in Table A.13. This average entry cost reflects the average difference in implied profitability between land which produces non-zero cropland and cropland: its large relative magnitude suggests some selection exists. However, introducing the inverse mills ratio selection correction does not meaningfully alter estimates of the global average $\gamma(\omega)$. Unpacking this result, I find that variation which drives the $0-1$ entry decision, which is a moment which is complete omitted from the main analysis, does not meaningfully alter the estimated relationship between agricultural returns and forest shares. The coefficient in Panel B is statistically indistinguishable from the coefficient in Appendix Table A.3, column (2). I thus conclude that selection is not on average a source of

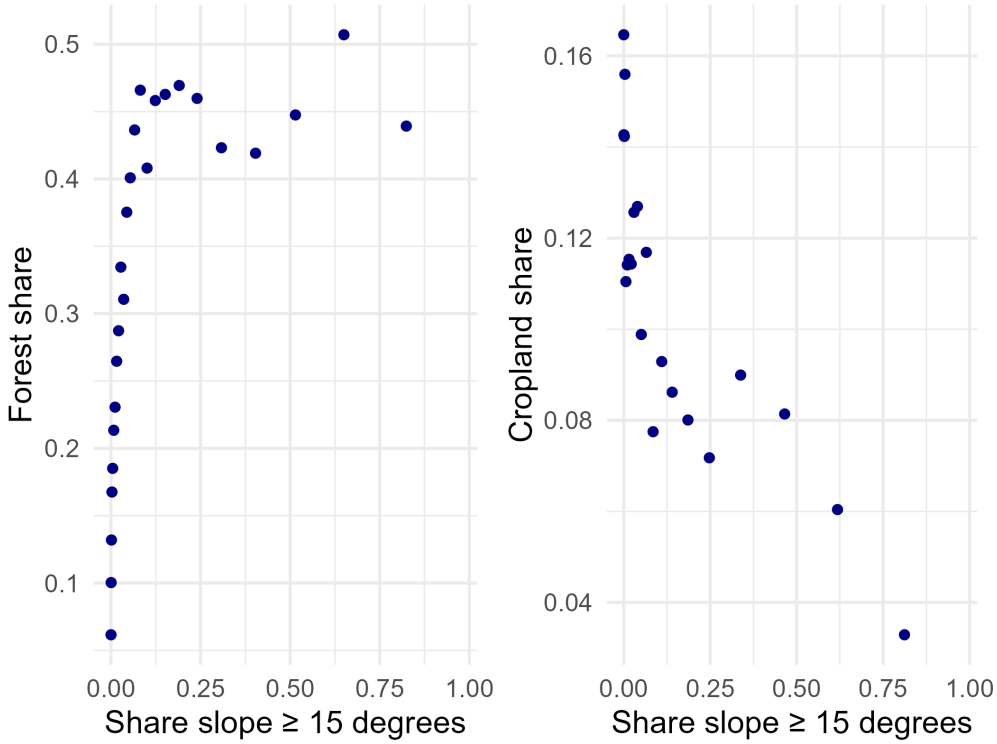

Figure A.1: Visual reduced form: binscatter of forest and cropland shares against the share of grid cell with steep slope

NOTES: Binned regression on a stratified random sample of 100 grid cells by country. Both figures control for a straight-line distance to city.

statistical bias in the target parameter. An alternate interpretation is that, after including controls for population access in the main analysis, I reduce bias from selection on interior shares relative to a model without controls.

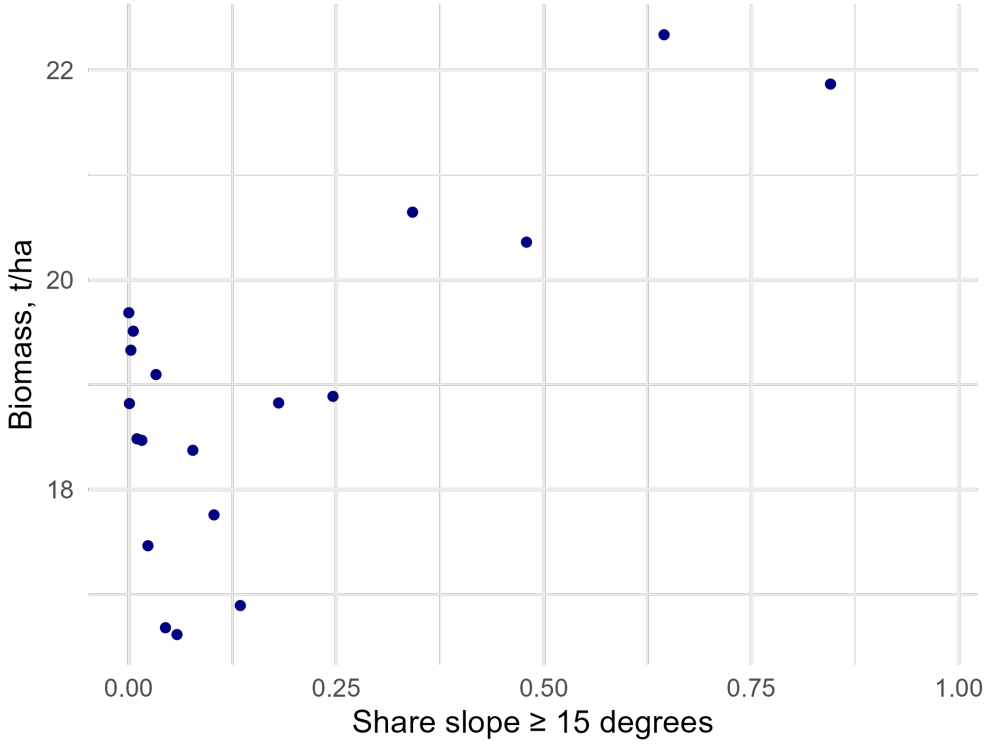

Figure A.2: Exclusion evidence: binscatter of biomass in 2010 on share of grid cell with steep slopes after controling for land selection.

NOTES: Binned regression on a stratified random sample of 100 grid cells by country. Controls for vegetation continuous fields measure of treecover and non-treecover vegetation, population density in 1975, and population access measure.

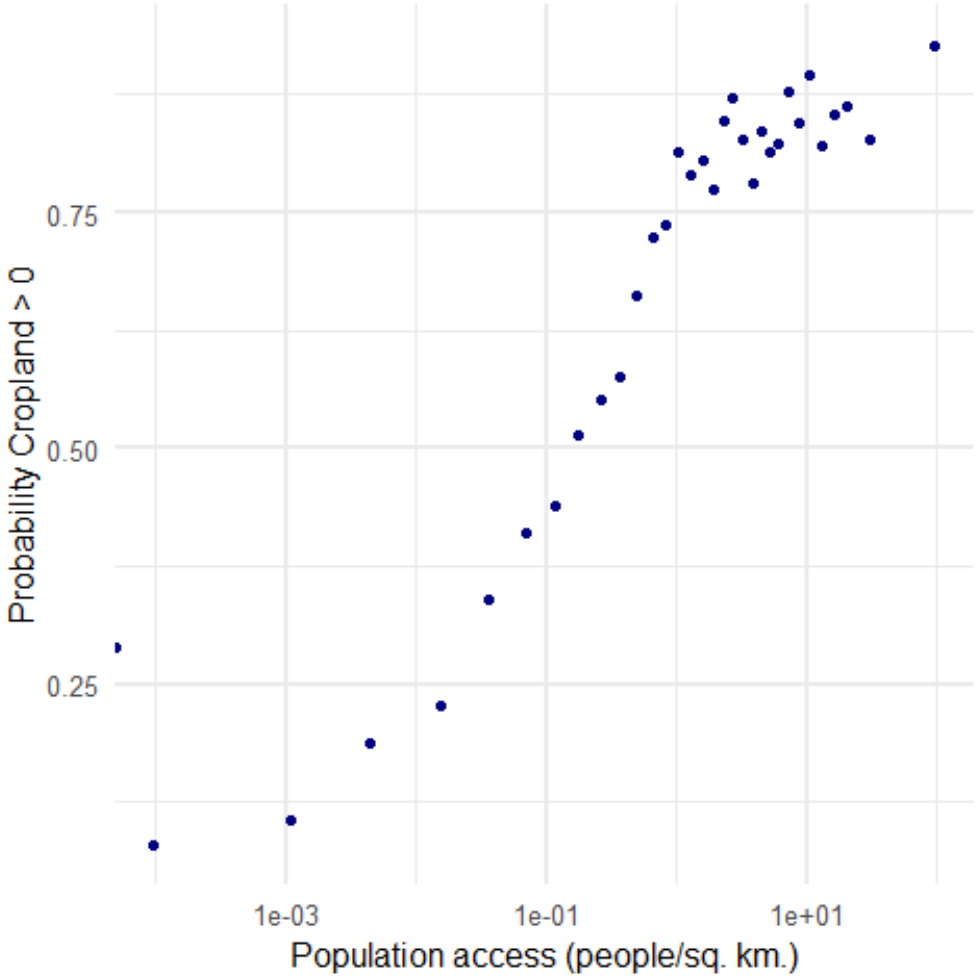

Figure A.3: Sample selection: probability of having nonzero Cropland share by grid-cell level population access measure.

NOTES: Binned regression on a stratified random sample of 100 grid cells by country. No controls; weighted by area. Population access is a within-country, inverse squared-distance-weighted average of population relative to a grid cell.

Table A.2: First stage of agricultural returns on Souza-Rodrigues (2019) instruments.

	Value of agricultural production ($/ha)				
	(1)	(2)	(3)	(4)	(5)
Distance to nearest city (km)	-0.1778***	-0.1773***	-0.1414***	-0.1351***	-0.1312***
	(0.0523)	(0.0519)	(0.0436)	(0.0416)	(0.0407)
Fraction grid with slope = 15 degrees		-8.013	-10.98	-8.112	-13.57*
		(8.082)	(8.130)	(8.409)	(8.113)
Port distance (km)			-0.1117***	-0.1091***	-0.1071***
			(0.0388)	(0.0380)	(0.0365)
Population access index				0.4680**	0.4644**
				(0.2344)	(0.2319)
1975 Population density				0.0088	0.0263
				(0.0130)	(0.0171)
other					-24.88***
					(8.920)
R^2	0.94902	0.94905	0.94957	0.94965	0.95005
Observations	254,149	254,149	254,149	254,101	254,101
country fixed effects	✓	✓	✓	✓	✓
state fixed effects	✓	✓	✓	✓	✓
county fixed effects	✓	✓	✓	✓	✓

NOTES: Observations are grid cells with cropland and forest shares $\in (0, 1)$ and a valid dollar value of transportation costs. Standard errors clustered at the tile level (100 $\times$100 grid cell tiles). Country, state, and county fixed effects reflect ADM-0, ADM-1, and ADM-2 fixed effects. Value of agricultural production refers to the sum of potential yield $\times$ prices and dollarized transport costs. Elasticity measures report the logit elasticity implied by the coefficient on value of agricultural production.

Table A.3: Two-stage least squares of land use share on agriculture returns using Souza-Rodrigues (2019) instruments.

	log(Cropland/(Forest))			
	(1)	(2)	(3)	(4)
(Intercept)	-4.089***			
	(0.8975)			
Value of agricultural production ($/ha)	0.0049*	0.0085***	0.0186***	0.0288***
	(0.0026)	(0.0025)	(0.0049)	(0.0096)
Population access index	0.0756***	0.0131	0.0079	0.0339***
	(0.0184)	(0.0187)	(0.0150)	(0.0097)
Population density, 1975	0.0003	0.0039**	0.0043***	0.0047***
	(0.0015)	(0.0016)	(0.0013)	(0.0014)
other	2.662***	4.098***	3.662***	3.013***
	(0.4678)	(0.4358)	(0.4388)	(0.3435)
R^2	0.02218	0.12930	0.15105	0.32900
Observations	254,101	254,101	254,101	254,101
F-test (1st stage), Value of agricultural production ($/ha)	9,039.3	14,137.3	5,023.9	2,064.7
Elasticity	0.9648	1.660	3.635	5.647
Mean Value ($/ha)	333.0	333.0	333.0	333.0
country fixed effects		✓	✓	✓
state fixed effects			✓	✓
county fixed effects				✓

NOTES: Observations are grid cells with cropland and forest shares $\in (0, 1)$ and a valid dollar value of transportation costs. Standard errors clustered at the tile level (100 $\times$100 grid cell tiles). Country, state, and county fixed effects reflect ADM-0, ADM-1, and ADM-2 fixed effects. Value of agricultural production refers to the sum of potential yield $\times$ prices and dollarized transport costs. Elasticity measures report the logit elasticity implied by the coefficient on value of agricultural production.

Table A.4: Two-stage least squares of land use share on agricultural returns without dollarizing transportation costs.

	log(Cropland/(Forest))			
	(1)	(2)	(3)	(4)
(Intercept)	-3.565***			
	(0.5223)			
Value of agricultural production ($/ha)	0.0060***	0.0099***	0.0066***	0.0096***
	(0.0013)	(0.0016)	(0.0011)	(0.0020)
Population access index	0.0466***	0.0008	0.0257**	0.0259**
	(0.0148)	(0.0119)	(0.0104)	(0.0115)
Population density, 1975	-0.0008	0.0021	0.0009	0.0040***
	(0.0012)	(0.0014)	(0.0011)	(0.0012)
other	2.434***	3.771***	2.634***	2.308***
	(0.4062)	(0.3754)	(0.2821)	(0.2911)
Distance (km) x yield (t/ha)	-0.0028***	-0.0028***	-0.0017***	-0.0047***
	(0.0006)	(0.0006)	(0.0004)	(0.0009)
R^2	0.02093	0.13536	0.32773	0.46384
Observations	384,571	384,571	384,571	384,571
F-test (1st stage), Value of agricultural production ($/ha)	35,850.0	52,236.0	42,155.2	26,857.4
Elasticity	0.8653	1.419	0.9430	1.380
Marginal cost of 1 km	-0.3347	-0.2081	-0.1921	-0.3515
country fixed effects		✓	✓	✓
state fixed effects			✓	✓
county fixed effects				✓

NOTES: Observations are grid cells with cropland and forest shares $\in (0, 1)$ and a valid dollar value of transportation costs. Standard errors clustered at the tile level (100 $\times$100 grid cell tiles). Country, state, and county fixed effects reflect ADM-0, ADM-1, and ADM-2 fixed effects. Value of agricultural production refers to the potential yield $\times$ prices. Elasticity measures report the logit elasticity implied by the coefficient on value of agricultural production.

Table A.5: Results of the two-step Heckman selection model.

	Estimate	Std. Error
Panel A: Entry Costs (USD per hectare)		
Mean entry cost	164.33	20.20
Observations	940,424	
Panel B: Outcome Equation (log(Cropland/Forest))		
Agricultural returns	0.0082	0.0031
Observations	244,012	

NOTES: Panel A reports entry costs estimated from selection equation (probit model of agricultural entry) normalized by agricultural value coefficient. Standard errors are calculated using the delta method and reflect the standard error of the agricultural value coefficient. Panel B shows the outcome equation coefficient on instrumented net agricultural returns, with instruments distance to city and fraction slope $> 15\%$. Inverse Mills ratio included to correct for selection bias. Country fixed effects in both stages. Standard errors clustered by tile.

C.2 Nonparametric heterogeneity

This section builds on the dollarized transportation cost design described Appendix Section C.1. In this section, I allow for global heterogeneity in the deforestation elasticity. I approximate the global distribution of land use elasticities with a discrete mixture model. Discrete mixtures can approximate arbitrary distributions of $\gamma(\omega)$ without requiring parametric forms of heterogeneity (Train 2008). Because deforestation elasticities represent both variation in unobservables related to institutions and technology, I prefer this nonparametric approach to parametric equivalents. Put differently, the data chooses salient dimensions of heterogeneity.

Define three discrete types of landowners indexed by a type $t \in \{H, M, L\}$. Without loss, the high type $t = H$ is a subset of landowners who are more elastic to price incentives, type M is less price sensitive, and type $t = L$ are the least price sensitive, so that the scale parameters obey $\gamma^H > \gamma^M > \gamma^L$.[3] Type probabilities are denoted ρ^H, ρ^M, ρ^L.

I use an Expectation-Maximization algorithm to estimate a logit with nonparametric discrete-type heterogeneity following Train (2008). I allow for type probabilities to vary at the country-by-ecoregion level. This level of aggregation maintains spatially fine distinctions across types of forest and different forest-crop frontiers, but refrains from overfitting the data (Dingel and Tintelnot 2020). Because the Hotz-Miller inversion implies a linear estimating equation, the maximization step amounts to a type-probability weighted logit regression.

EM solutions can be sensitive to initial conditions. I use a "distance-squared" initialization to ensure that I accurately cover the range of possible starting coefficients γ^H, γ^L. This initialization resamples the bootstrap distribution of $\gamma(\omega)$ (I use 30 draws) and randomly selects a single draw as the first type. Other types are drawn from the remaining bootstrap distribution with a probability weight proportional to their squared distance from the coefficient γ of the initial draw, ensuring sufficient coverage of the search space. In some bootstrap draws, the EM converges to a small mass on a type with a negative coefficient on revenues: in such cases, I redraw this type and continue the EM algorithm.

Turning to results in Appendix Table A.6, 25% of the sample is very elastic to prices. They are nearly ten times as elastic as the average and median producer. The most elastic producers tend to be in agriculture-dominant regions with high population densities and low costs of getting to market.

In Figure A.4, I illustrate the geographic distribution of types globally. For plots which are not in the analysis sample (where Cropland or Forest is exactly 0), I impute elasticities both for this map and in the main analysis. pixels first by using country-level average elasticities. Remaining missing values are assigned biome-level elasticities: this means that missing tropical forest can be imputed with the value of tropical forest in other countries (though this is quite rare). Finally, any parcels which do not have an elasticity in the same country or the same biome are assigned the global average elasticity. Regions which are more forested – including the depths of the Amazon – are less

3. Estimates of many more types – ranging from 5 to 50 – resulted in very limited quantitative or qualitative changes to key results. I choose three types for parsimony.

elastic to prices than regions in population centers. This heterogeneity is completely driven by the data; the algorithm does not require a coherent or spatially correlated type probability. However, the data suggests that the Brazilian coast is far more likely to respond to a tax-based incentive than the more carbon-rich Brazilian interior.

Table A.6: Results of estimating Equation (13) with a discrete mixture model.

Variable	Type 1	Type 2	Type 3
Value of agricultural production ($/ha)	0.0038 (0.0027)	0.0112 (0.0072)	0.0203 (0.0057)
Travel cost (t-km/ha)	-0.0023 (0.0016)	-0.0036 (0.0034)	-0.0059 (0.0039)
IV F-stat (1st stage)	109864.67	110083.24	109571.00
Type probability	0.4208 (0.1403)	0.2965 (0.1289)	0.2526 (0.0977)
Derived Statistics			
Elasticity equivalent	0.7250 (0.5278)	2.0755 (1.2669)	3.8010 (1.0022)
Marginal cost of 1 km ($/ha)	-1.02 (1.55)	-0.42 (0.46)	-0.27 (0.15)
Type-weighted averages of observables			
Agricultural value ($/ha)	319.0218	323.4890	351.5137
Travel cost (t-km/ha)	195.2837	205.3507	181.7519
Population access	3.1711	3.6627	3.9575
Population density, 1975	3.2013	3.6642	3.9832
Other land share	0.4795	0.4002	0.3487
Cropland share	0.1773	0.2073	0.2549

NOTES: Bootstrap results from 119 draws of the EM algorithm with 3 types. Estimates show mean (standard error) across bootstrap draws. Types are sorted by their coefficient on the value of agricultural production, ensuring Type 1 always has the lowest elasticity and Type 3 the highest. IV F-statistic is the first-stage Cragg-Donald F-statistic for the instrument validity. Elasticity equivalent averages logit elasticity formula over pixels for cropland shares. Marginal cost of 1 km rescales travel cost coefficient by the coefficient on returns. Type probabilities are area-weighted means. Convergence rate: 96.2%.

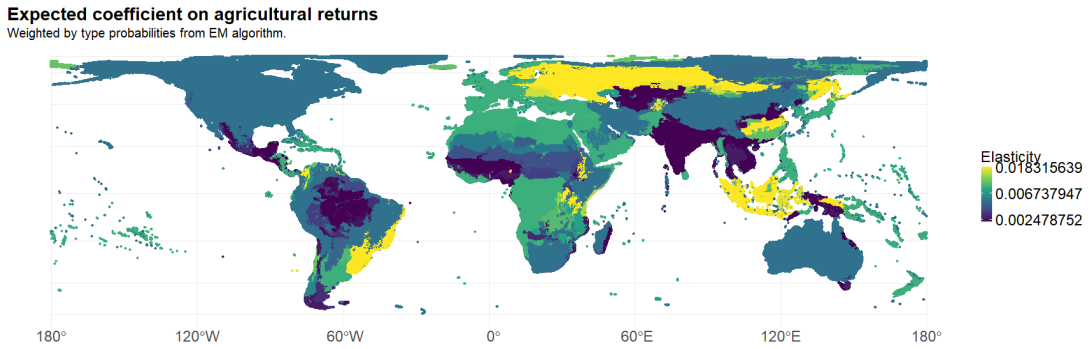

Figure A.4: Map of expected coefficient on agricultural returns.

NOTES: Reports type-probability-weighted sum of the coefficients reported in Appendix Table A.6 for each pixel. Observations which are not in-sample for EM algorithm are imputed using a biome-level average.

C.3 Spatial first differences

Spatial correlation in unobservables $\xi^F(\omega)$ threaten the exclusion restriction for the distance and slope instruments in Appendix Sections C.1. While my controls generate robust results for the deforestation elasticity, unobservable amenities may continue to be correlated with distance to city or terrain variables. A pernicious example might be local variation in the wage, where local labor markets are unobserved and spatially correlated. High-yield regions will have systematically different wages. If the productivity of local labor markets is correlated with distance to city or local slopes (e.g., through sorting), then exclusion may fail. Here, I introduce an alternate "spatial quasi-differences" identification strategy which weakens the previous identification assumption (extending Druckenmiller and Hsiang 2019).

Define $\{x_\omega, y_\omega\}$ as the coordinates for ω in Cartesian space. Define $d\omega$ as the neighboring plot at $\{x_\omega, y_\omega\} + \vec{e}$ for a step $\vec{e}$ in some direction (e.g., in the x-direction, $\vec{e} = \{1, 0\}$). Neighbors share a component of their unobservable return, $\xi^F(\omega)$. For example, neighbors share the same wage w_i. I assume this spatially persistent component of the structural error obeys an AR(1) process with coefficient ρ:

$$\xi^F(\omega) = \rho\xi^F(d\omega) + v(\omega)$$

where the remaining portion of the unobservable, $v_t(\omega)$, is as-good-as-random. This assumption motivates a spatial quasi-difference to remove the effect of the persistent unobservable. Defining $\breve{x}(\omega) = \rho x(\omega) - x(d\omega)$,

$$\breve{Y}^{AF}(\omega) = \gamma(\omega)[\breve{\pi}^A(\omega)] + (\rho - 1)v_i + v(\omega) \tag{A.5}$$

Spatial quasi-differencing thus removes the persistent unobservable by assumption, leaving only the innovation process $v(\omega)$.[4]

This spatial quasi-differencing strategy relaxes the requirement that the distance-to-city and terrain-slope instruments satisfy the exclusion restriction conditional on observables alone. Suppose the placement of infrastructure is correlated with unobserved local fundamentals, such as labor market productivity or other amenities, that are imperfectly captured by population controls. I assume these unobservables are smooth at a sufficiently local scale. By quasi-differencing outcomes and instruments across neighboring grid cells, the contribution of these spatially correlated unobservables is differenced out to a first order.

Under this assumption, identification does not require global orthogonality between the instruments and unobserved fundamentals. Instead, it requires that there exists residual, high-frequency spatial variation in straight-line distance to cities and in the fraction of land

4. I prefer quasi-differencing to an overt spatial first difference because satellite outcomes are measured with error. Further, as deforestation is a rare event, temporal AR(1) approaches return estimates of ρ close to 1 even over longer panels. The OLS estimator is no longer asymptotically normal at these unit roots due to nonstationarity (e.g., shock effects are increasing in time) (Arellano and Bond 1991).

with slope exceeding 15 degrees that is orthogonal to the remaining, locally smooth unobserved heterogeneity. Intuitively, while economic fundamentals evolve smoothly across space, fine-scale variation in geography induces discontinuous local variation in market access and effective production costs, generating plausibly exogenous shifts in the interaction of prices and potential yields.

Estimation proceeds by GMM, where the instruments discussed in Appendices C.1 and C.4 would be used to weight residuals of Appendix Equation (A.5) to get estimates of ρ and $\gamma(\omega)$. I use spatial differences of included and excluded instruments in a two-step GMM estimator. In spatial differences, the GMM places higher weight on slope-related variation. Slope variation is less spatially correlated, and thus provides greater power than distance to city in differences.

Appendix Table A.7 shows results. First, the spatial correlation parameter ρ is close to 1. For reference, logit shares $Y(\omega)$ demonstrate a within-country, spatial AR(1) correlation coefficient of 0.42 globally. High autocorrelation in returns, for which $\rho = 0.95$, drives ρ. Second, the coefficient on returns is approximately constant across runs. The resulting elasticity is very close to the preferred estimates in Table 2. Transport costs are slightly higher in differences, sitting at \$0.5-1 per ton-kilometer. In Appendix Table A.4, this figure was closer to \$0.20 per km.

Thus, while spatial correlation does not appear to meaningfully attenuate the counterfactual-relevant deforestation-revenue elasticity, spatial correlation in unobservables does understate transportation costs. This result is consistent with agricultural returns being larger in regions with higher transportation costs.

Table A.7: Spatial First Difference Estimates

	Western Neighbor (1)	Eastern Neighbor (2)	Southern Neighbor (3)	Northern Neighbor (4)
Panel A: Parameters				
ρ	0.9901	0.9552	0.9639	0.9801
	[0.9855, 0.9932]	[0.9496, 0.9602]	[0.9582, 0.9689]	[0.9756, 0.9837]
Gross returns (IV)	0.0137	0.0081	0.0156	0.0087
	[0.0102, 0.0171]	[0.0057, 0.0104]	[0.0120, 0.0192]	[0.0060, 0.0114]
Transport costs	-0.0094	-0.0048	-0.0101	-0.0050
	[-0.0116, -0.0073]	[-0.0067, -0.0029]	[-0.0125, -0.0076]	[-0.0070, -0.0031]
Panel B: Model Fit				
First-stage F-stat	14791.47	16045.38	16325.45	16159.73
Observations	303,422	303,412	297,986	298,000
Country FE	Yes	Yes	Yes	Yes
Clustered SE	Tile	Tile	Tile	Tile

NOTES: Estimates from 2-step GMM spatial first difference models. ρ is the spatial correlation parameter. Gross returns are instrumented using spatially differenced distance to city. Transport costs are included as a control. Controls for ADM-2 fixed effects, population density, population access, and other land use shares, all spatially differenced. 95% confidence intervals in brackets. Standard errors clustered at the tile level.

C.4 Tariff shifter

In this section, I explore changes in tariffs abroad as a shifter of domestic agricultural prices. Tariff changes abroad generate variation in demand for goods produced domestically. Both in the framework of this paper and broader work on international trade, to a first order an increase in a tariff as applied by an importer effectively shrinks the market size of an exporter. The first stage of my instrument takes the following form.

$$p_{it} = \beta_1 ADV_{ijt} + \beta_2 ADV_{ijt} \times \mathbb{I}[\lambda_{ij0} > \text{Median}_k(\lambda_{ik0})] + \mu_{ij} + \mu_{jt} + \epsilon_{ijt}$$

where ADV_{ijt} represents the ad valorem tariff equivalent between exporter i and importer j in year t, p_{it} represents prices in the exporter, λ_{ij0} refers to the trade share of i in importer j in a base period 0, and μ_* are fixed effects. Fixed effects partial out the starting levels of bilateral trade μ_{ij} and control for importer-specific supply shocks μ_{jt}.

Intuitively, the first stage represents a continuous difference-in-differences design on tariffs comparing more exposed (as measured through baseline trade shares) exporters to less exposed exporters. Exposure here captures the idea of "large" importers, as measured by their trade share (a theory-consistent measure, as larger equilibrium price effects result from lower autarky shares and thus larger bilateral shares). Identification relies on the idea that aggregate supply shocks in a given country i do not strongly predict tariff changes in partner countries j. Identification would be compromised if, for example, positive production trends on average predicted increasing tariffs.

My narrative predicts that a higher tariff reduces demand for goods from i: $\beta_1 < 0$. I also argue that $\beta_2 < 0$: this requires a stronger assumption that larger importers are not significantly more inelastic to changes in price. If higher import shares are explained by higher idiosyncratic preference for a particular variety of a product, this relationship would be weaker or reversed. Because the agricultural commodities in my study are likely more physically substitutable than other products, $\beta_2 < 0$ is a more appropriate prediction than in other settings.

Columns (1) and (2) of Appendix Table A.8 corroborate my hypothesss. On average, a 1 basis point increase in ad valorem equivalent tariffs decreases exporter producer prices by \$0.27, or a 0.3% decrease based on average 2001 prices. Above-median importers experience a 50% increase (\$0.13) larger impact on prices.

This story could be consistent with tariffs coinciding with positive supply shocks, however. If importers apply tariffs differentially to regions which experience supply gluts to protect domestic goods, then tariffs would appear to predict a contemporaneous price decrease. As an imperfect test of this identification threat, I re-estimate the first stage with log-quantities on the left hand side. In columns (3) and (4) of Appendix Table A.8, I find that, if anything, exporter quantities decrease in years with an increase in ad valorem tariff increase, and there is no marginal interaction effect of a larger importer. Decreases in quantities are inconsistent with a supply-driven tariff on net. In column (4), the average effect of a 1 basis point increase in tariffs on equilibrium quantities is a

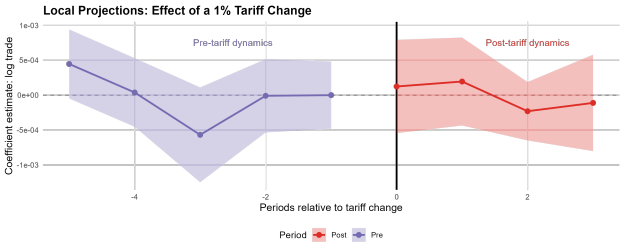

Figure A.5: Estimates from a local projections model of trade flows on tariffs.

NOTES: Local projections conducted using a 3-year time step based on availability of MacMAP data, 2001-2016. Sample composition thus differs across observations: observations closer to 0 have more dyad-years available. Standard errors are two-way clustered at the importer-and-exporter level.

0.0026% decrease.

I also analyze the dynamics of production around a tariff increase. I run a local projections model which looks at exports between countries prior to tariff shocks:

$$X_{ij,t+h} = \beta^h ADV_{ij,t+h} + \mu_{ij}^h + \mu_{jt}^h + \epsilon_{ijt}^h$$

Conceptually, I compare the time path of trade among countries which experience a tariff increase of 1 percentage point to those who experience no such tariff increase. Fixed effects control for the initial level of trade and protection within a dyad (i, j) and global trends in demand (j, t) across exporters to refine attention to the exporter-specific tariff shock. The results are plotted in Appendix Figure A.5. Results again demonstrate no clear evidence of rising productivity ex ante. Post-tariff dynamics are also indistinguishable from 0. Because tariffs are relatively aggregated, finer tariff data or finer information on a crop-level may result in stronger correlations. In aggregate, however, tariffs are inconsistent with a productivity shock for agriculture.

The first stage can be recast as a shift-share instrument to aggregate variation to the level of an individual exporter-year. The shift share IV I consider has the following form:

$$p_{it} = \beta_1 \sum_{j=1}^{J} \omega_{ij} ADV_{ijt} + \beta_0 t + \mu_i + \epsilon_{ijt}$$

Weights ω_{ij} are either: baseline trade shares among all countries, or baseline trade shares if the importer is in the top half of country-level imports and 0 otherwise. Fixed effects ensure identification comes off of changes in prices on changes in the instrument relative to baseline levels. I control for a global trend to partial out the average trend in commodity prices over the period – country-level trends in prices are too demanding. Appendix Table A.9 shows results. The first stage is reasonably strong, with a 1p.p. increase in trade-weighted tariffs resulting in a 1.2% fall in prices. Given that a 1p.p. increase in trade-weighted tariffs is approximately 6% of baseline tariffs

in 2001, this amounts to a price elasticity of around 0.2. Responses are similar in columns (1) and (2), where (2) drops below-median countries in the instruments' weighting. Identification is thus primarily driven by local effects of high-trade share countries rather than a generic tariff increase. Column (3) upweights importers with higher initial tariffs. While precision is lower, the smaller coefficient magnitude is consistent with higher initial tariff countries having smaller decreases in demand in relative terms. This behavior is consistent with concavity in the elasticity of demand.

As a potential richer test, I note that one could leverage the timing of tariffs relative to the growing season of crops. If tariffs are set after crops are planted in a given country, then tariffs would physically be unable to affect supplied quantities, completely sidestepping supply uncertainty in this paper. Because I lack the precise data on harmonized tariff changes, I leave this to future work.

Table A.8: First stage effects: changes in agricultural tariffs and agricultural prices.

	Producer price		log(Q)	
	(1)	(2)	(3)	(4)
Ad valorem equivalent tariff	-26.73*	-5.125	-0.0966*	-0.0261
	(14.48)	(9.144)	(0.0500)	(0.0278)
Above median imports	26.70	-14.83	0.1627	0.0890
	(33.99)	(27.49)	(0.1130)	(0.0546)
Ad valorem equivalent tariff × Above median imports	-13.24*	-14.54*	0.0040	-0.0054
	(7.649)	(7.787)	(0.0297)	(0.0223)
R^2	0.76596	0.78507	0.98390	0.98682
Observations	52,566	52,566	52,566	52,566
F-test	118.64	132.41	2,215.0	2,714.0
Average Price Elasticity	-0.41	-0.15		
Avg. price, 2001	76.46	76.46	76.46	76.46
Avg. output, 2001	14,186,824.4	14,186,824.4	14,186,824.4	14,186,824.4
Change in Y: going from 0 average AVE	-6.28	-2.24	-0.02	-0.01
exporter-importer fixed effects	✓	✓	✓	✓
exporter fixed effects	✓	✓	✓	✓
importer fixed effects	✓		✓	
importer-year fixed effects		✓		✓

NOTES: Observations are dyad-years. Standard errors are two-way clustered at the importer-by-exporter level. Prices, quantities, and ad valorem tariffs are all land-share weighted averages with land share weights taken in 2001. Average price elasticity divides coefficient by mean prices, conditional on observe prices above $10 to clean out erroneously small prices in the FAO data. Change in Y going from 0 to average AVE measures the average price and $\log(Q)$ change implied by multiplying the coefficients on tariffs by the average observed ad valorem equivalent.

Table A.9: Shift-share first stage: weighted average ad valorem tariff on prices.

	Producer price		
	All	Above median	High initial tariff
	(1)	(2)	(3)
Change in price (USD/t) for 1p.p increase in tariff	-0.9662*	-0.9654*	-0.7392**
	(0.5579)	(0.5569)	(0.3498)
year	13.43***	13.43***	13.10***
	(2.536)	(2.536)	(2.399)
R^2	0.83198	0.83198	0.83282
Observations	320	320	320
F-test (projected)	30.300	30.300	31.080
Avg. price, 2001	77.37	77.37	77.37
Avg. Z, 2001	16.58	16.59	16.59
exporter fixed effects	✓	✓	✓

NOTES: Observations are exporter-years. Standard errors are clustered at the exporter level. Column (1) RHS is a trade-share weighted sum of tariffs using trade from 2001. Column (2) is the same but only weights trade among the top 50% of importers for a given exporter, renormalizing so weights sum to 1 within this group. (3) reuses the full set of trade weights from (1) and multiplies trade shares by ad valorem rate in 2001, again renormalizing so weights add to 1. Year shows a coefficient on a linear trend.

Table A.10: Two-stage least squares estimates: cropland-forest response to tariff-induced revenue shocks.

	Log Cropland to Forest Ratio			
	OLS	IV		
	(1)	(2)	(3)	(4)
Revenue (USD / ha)	0.0006*	0.0012**	0.0013**	0.0013**
	(0.0003)	(0.0005)	(0.0006)	(0.0006)
Transport cost (t-km / ha)	-0.0005	-0.0006	-0.0007	-0.0007
	(0.0004)	(0.0004)	(0.0004)	(0.0004)
Avg. Pop. Density (1975)	0.0129***	0.0115***	0.0113***	0.0112***
	(0.0035)	(0.0034)	(0.0033)	(0.0032)
All Other Land Uses	2.185***	2.299***	2.323***	2.325***
	(0.6648)	(0.6423)	(0.6404)	(0.6413)
R^2	0.14570	0.14691	0.14732	0.14741
Observations	2,128,649	2,120,387	2,120,387	2,120,387
F-test (1st stage), Revenue (USD / ha)		555,034.1	665,714.3	1,723,002.1
Elasticity Equivalent	0.1110	0.2310	0.2530	0.2550
Transportation cost ($/t-km)	-0.8775	-0.5378	-0.5102	-0.5076
country fixed effects	✓	✓	✓	
year fixed effects			✓	
country-year fixed effects				✓

NOTES: Observations are grid-years. Standard errors are clustered at the country level. IV in (2)-(4) uses shift share of tariffs weighted by baseline trade shares × potential yields $\eta^A(\omega)$ to instrument country-level prices × potential yields. F-test reflects strength of this combined instrument relative to first-stage on price alone in A.9. Elasticity equivalent calculated using logit formula, but restricting shares to cropland share within forest-cropland. Transportation costs are dollarized by dividing transportation costs by revenue coefficient.

C.5 Bayesian bootstrap

The nonparametric heterogeneity in Appendix Section C.2 produces a global, grid cell level distribution of land use elasticities. To generate confidence intervals for counterfactuals, I develop a resampling approach which allows me to recover global, plot-level distributions of deforestation elasticities using a Bayesian bootstrap. The Bayesian bootstrapping approach allows for complete global coverage by reweighting observations systematically rather than dropping them entirely. The procedure can be succinctly described as follows. For a given bootstrap run indexed b,

1. Draw plot-level weights $w^b(\omega)$ from a Dirichlet distribution. Dirichlet weights are correlated (clustered) at the tile (10×10 grid cell block) level and the country level.
2. Initialize the EM algorithm using coefficients from weighted logit regressions in the k-means++ algorithm described in Appendix Section C.2. For n types, $10n$ regressions are run on randomly weighted subsets of the data to recover n maximally "far apart" types.
3. Solve for the distribution of EM coefficients using the iterative logit algorithm in Train (2008)
4. Hierarchically interpolate EM coefficients for land which had no interior land use shares. EM coefficients are assigned as: (1) the mean of coefficients in the same country; (2) if still missing, the mean of coefficients in the same ecoregion in other countries; (3) if still missing, the global mean elasticity.
5. This global plot-level distribution is then stored as data which can be used to calibrate the full equilibrium model

With a bootstrap approach, there is some risk that an individual bootstrap run places small weight on an implausible (e.g., negative or extremely small) coefficient on agricultural revenues. Such coefficients can be numerically very difficult to resolve: the equilibrium solution method will be unstable when even some small subset of land has arbitrarily large inclusive values. Thus, in any bootstrap run which has a type of land use parameter $\gamma(\omega)$ which is less than 10^{-4}, ten times smaller than the smallest parameter draw in the preferred specification, I redraw that type from the initial k-means++ distribution.

When simulating the model over bootstrap draws, simulation follows a similar algorithm.

1. Draw a dataset of EM land use elasticities, the output of the previous algorithm.
2. Initialize the calibrated baseline equilibrium based on these elasticities using the algorithm discussed in Section 5.3
3. Determine a welfare-maximizing tax rate using a golden search over the range of taxes $[0, 20]$

This returns the outputs necessary for the bootstrapped confidence intervals in Tables 4 and 5.

D Additional tables and figures

D.1 Metadata

Table A.11: Reclassified labeling of the European Space Agency Climate Change Initiative data used to categorize land uses. 6 total categories were constructed, including excluded areas (NA).

ESACCI		This Paper	
Code	Description	Description	Code
0	No data	No data	NA
10	Cropland, rainfed	Cropland	2
11	Herbaceous cover	Other	4
12	Tree or shrub cover	Other	4
20	Cropland, irrigated or post-flooding	Cropland	2
30	Mosaic cropland (> 50%) / natural vegetation	Cropland	2
40	Mosaic natural vegetation	Forest	1
50	Tree cover	Forest	1
60	Tree cover	Forest	1
61	Tree cover	Forest	1
62	Tree cover	Forest	1
70	Tree cover	Forest	1
71	Tree cover	Forest	1
72	Tree cover	Forest	1
80	Tree cover	Forest	1
81	Tree cover	Forest	1
82	Tree cover	Forest	1
90	Tree cover	Forest	1
100	Mosaic tree and shrub (> 50%) / herbaceous cover (< 50%)	Other	4
110	Mosaic herbaceous cover (> 50%) / tree and shrub (< 50%)	Other	4
120	Shrubland	Other	4
121	Shrubland evergreen	Other	4
122	Shrubland deciduous	Other	4
130	Grassland	Cropland	2
140	Lichens and mosses	Other	4
150	Sparse vegetation	Other	4
151	Sparse tree (< 15%)	Other	4
152	Sparse shrub (< 15%)	Other	4
153	Sparse herbaceous cover (< 15%)	Other	4
160	Tree cover	Forest	1
170	Tree cover	Forest	1
180	Shrub or herbaceous cover	Other	4
190	Urban areas	Other	3
200	Bare areas	Other	5
201	Consolidated bare areas	Other	5
202	Unconsolidated bare areas	Other	5
210	Water bodies	No data	NA
220	Permanent snow and ice	No data	NA

Table A.12: Data sources for all raster (point-cloud) data used in this project. Table lists the resolution at which each dataset is available.

Data source	Resolution	Description
European Space Agency Climate Change Initiative	300 m	Globally estimated land cover maps provided annually, 1995-2019
Food and Agriculture Organization Global Agro-ecological Zones	9 km	Potential yield data and land resources (soil, climatology). High-yield, rainfed conditions. Units: 10 kg dry weight per hectare.
Travel time to cities (Nelson et al. 2019)	1 km	Global travel-time accessibility indicators to cities of various populations for the year 2015. Use the most expansive city definition.
World Port Index	(NA)	Shapefile of the location of world ports
City locations data (Akbar et al. 2023)	(NA)	Shapefile of major cities
World Wildlife Fund (Olson et al. 2001)	1 km	Describes bio- and eco-regions of the world (used to determine "tropical" forest regions)
Distributed Active Archive Center for Biogeochemical Dynamics, NASA	300 m	Measures aboveground biomass in Megagrams Carbon per hectare in 2010.
International Union for Conservation of Nature and Natural Resources: Red List of Threatened Species	5 km	Measures the ranged species rarity of all red list species in each cell.
Global Human Settlement Layer	30m	Maps for 1975 and 1990 on presence of built structures. Derived from Landsat image collections.

D.2 Descriptives

Table A.13: Summary statistics of agricultural returns and carbon.

	Mean	SD	Q25	Q50	Q75
Global Land					
CO_2, t/ha	36.83	53.21	1.37	13.17	49.39
Potential revenue (V_A), USD/ha	212.81	227.01	0.00	169.74	314.60
Transportation cost (τ), USD/ha	1.08	4.33	0.02	0.09	0.53
Net return per t CO_2	5.072	8.243	0.458	1.740	5.528
Deforested Land (VCF change $< -1\%$)					
CO_2, t/ha	49.77	44.14	18.87	37.82	66.33
Potential revenue (V_A), USD/ha	247.03	207.37	117.77	207.70	322.22
Transportation cost (τ), USD/ha	0.98	4.06	0.06	0.17	0.56
Net return per t CO_2	2.922	5.778	0.569	1.310	2.602

NOTES: Area-weighted summary statistics. CO_2 computed as biomass per hectare times 44/12 times 0.47. Potential revenue (V_A) is producer price index times yield. Transportation cost (τ) calculated as travel time to market times yield times per-km transport cost. Net return per t CO_2 is ($V_A - \tau$) divided by biomass. Includes only land with at least 1 t CO_2 emissions to prevent dividing by 0. Deforested land defined as vegetation cover change less than -1% between 1982-2016. All monetary values in USD per hectare.

Table A.14: Yield statistics on deforested land for top 10 countries by emissions, 1982-2016.

Country	Avg Yield (t/ha)	Q10 Y/E (t/tCO_2)	Q90 Y/E (t/tCO_2)
Brazil	3.11	0.002	0.02
Canada	1.50	0.001	0.01
Congo - Kinshasa	3.46	0.002	0.01
Russia	1.53	0.001	0.01
Paraguay	3.99	0.006	0.02
United States	1.98	0.001	0.02
Bolivia	3.72	0.003	0.02
Indonesia	2.59	0.001	0.01
Argentina	4.27	0.006	0.10
Zambia	3.51	0.004	0.01
Everyone else	3.52	0.003	0.03

NOTES: Yield statistics for the top 10 countries from Table A.16, in the same order. Average yield shows area-weighted mean maximum potential yield on deforested land, where maximum is across set of in-sample crops. Q10 and Q90 show 10th and 90th percentiles of yield-per-emissions ratio (tons yield per ton CO_2). Statistics weighted by deforested area and limited to areas with positive yields and biomass.

Table A.15: Top 10 countries by deforestation emissions, 1982-2016: vegetation cover.

Country	1980 Share (p.p)	Vegetation change, 1982-2016 (p.p)	Non-decreasing share (p.p)
Brazil	58	-3.8	59.4
Canada	32	0.3	79.0
Congo - Kinshasa	64	-1.3	62.3
Russia	31	4.2	86.4
Paraguay	66	-18.4	10.1
United States	26	2.5	86.2
Bolivia	56	-3.2	60.4
Indonesia	70	2.2	82.5
Argentina	19	-2.7	63.5
Zambia	32	-4.7	40.5
Everyone else	15	1.6	90.2

NOTES: Top 10 countries ranked by share of global deforestation emissions. Forest cover levels and change from VCF data. All statistics weighted by area. Non-decreasing share refers to area with vegetation change ≥ 0.

Table A.16: Top 50 countries by deforestation emissions, 1982-2016.

Country	Emissions (tCO_2/km^2)	Area (1000 km^2)	Share (%)	Country	Emissions (tCO_2/km^2)	Area (1000 km^2)	Share (%)
Brazil	61	39540.7	34.4	Peru	118	240.7	0.4
Canada	44	12049.7	7.6	Nicaragua	78	349.9	0.4
Congo - Kinshasa	84	5073.5	6.0	Mexico	32	841.0	0.4
Russia	41	8763.3	5.1	Papua New Guinea	165	159.9	0.4
Paraguay	42	7377.1	4.4	Chile	41	514.9	0.3
United States	50	5956.7	4.2	Malaysia	112	179.8	0.3
Bolivia	67	4223.8	4.0	South Africa	27	711.9	0.3
Indonesia	108	2367.1	3.6	Mongolia	26	746.3	0.3
Argentina	29	8401.3	3.4	Uganda	42	421.1	0.3
Zambia	57	3685.3	3.0	Malawi	40	385.3	0.2
Angola	60	3375.7	2.9	India	52	292.1	0.2
Mozambique	42	4694.5	2.8	France	46	311.5	0.2
Australia	48	3725.4	2.5	Laos	71	199.0	0.2
Tanzania	40	3266.4	1.9	Honduras	82	166.1	0.2
Madagascar	58	2002.9	1.7	Cameroon	157	77.5	0.2
Myanmar (Burma)	76	940.2	1.0	Kenya	44	229.4	0.1
Congo - Brazzaville	166	370.5	0.9	Ghana	43	222.2	0.1
Cambodia	57	830.0	0.7	Belize	76	102.1	0.1
Gabon	211	217.5	0.7	Thailand	48	159.0	0.1
Colombia	95	443.6	0.6	Côte d'Ivoire	45	167.2	0.1
Venezuela	69	506.3	0.5	Nigeria	43	146.4	0.1
Vietnam	57	581.8	0.5	Zimbabwe	31	185.6	0.1
Ethiopia	53	595.9	0.4	North Korea	35	154.3	0.1
Guatemala	50	617.4	0.4	Solomon Islands	130	41.1	0.1
China	20	1534.0	0.4	Guyana	114	38.7	0.1

NOTES: Top 50 countries ranked by share of global deforestation emissions, 1982-2016. Emissions calculated as biomass carbon content converted to CO_2 as in Equation (1). Area shows total deforested area in thousands of square kilometers. Share represents percentage of global deforestation emissions. Only includes areas with forest loss.

Table A.17: Balance table across observations with interior land shares in 2000 and those without.

	Non-interior		Interior		Differences	
	Mean	Std. Dev.	Mean	Std. Dev.	Diff. in Means	Std. Error
Biomass, t C/ha	20.46	34.789	23.89	26.227	3.43	6.095
Maximum potential yields, kg/ha	1157.60	1401.480	2799.13	1219.705	1641.52	714.063
Distance to nearest city, km	560.92	542.709	193.79	224.069	-367.13	49.436
Distance to nearest port, km	611.49	489.317	514.82	493.784	-96.66	131.332
Belowground biomass, t C/ha	86.00	94.967	99.89	77.508	13.89	25.481
Population density per sq. km., 1975	0.64	8.724	3.61	22.368	2.97	0.921
Vegetation change in %, 1982-2016	0.71	4.553	2.01	9.386	1.31	0.514

NOTES: Interior land use grid cells must have agricultural and forested shares within $(0, 1)$ in the 2001 ESACCI data. Among 1,591,999 terrestrial grid cells, 596,609 have interior shares of target land uses, with area share of 43%. All statistics are area-weighted.

D.3 Additional figures: Pigouvian tax counterfactual

Table A.18: Weighted average income changes by factor group and weighting scheme.

Weighting scheme	Total	Manufacturing	Agricultural	Landowners
Emissions	0.06	0.13	-4.93	1.93
Income	0.06	0.01	-0.34	0.41
Agricultural production	0.01	-0.06	-0.30	0.70

NOTES: This table shows weighted average percent changes in real income by factor group under different weighting schemes. Emissions weighting uses baseline country-level emissions shares. Income weighting uses baseline country-level total income shares. Agricultural production weighting uses baseline country-level agricultural production shares. All income values are deflated by country-level consumption price indices.

Table A.19: Summary of production and GDP changes by country under optimal carbon tax.

Country	% Change, Emissions	% Change, Agriculture	% Change, real income
Global	-63.1	0.6	0.06
Landowners			0.22
Agricultural workers			0.26
Manufacturing workers			0.03
Australia	-56.8	0.4	0.05
Bolivia	-47.7	-2.0	-0.11
Brazil	-59.5	-0.2	0.31
Canada	-54.0	1.9	0.01
Chile	-64.5	-0.0	0.07
China	-34.9	0.4	-0.09
Colombia	-69.3	-1.0	0.42
Congo River Basin	-77.7	-9.4	-1.40
Ethiopia	-36.8	-0.1	-0.06
Europe	-63.8	1.0	0.44
Indonesia	-83.2	-3.1	0.64
India	-23.0	0.2	-0.24
Iran	-45.4	0.2	-0.03
Kazakhstan	-13.3	0.1	0.01
Kenya	-45.2	-0.0	0.00
Madagascar	-57.9	0.1	0.09
Mexico	-15.5	-0.1	0.00
Mali	-0.1	-0.0	0.12
Myanmar (Burma)	-46.1	-7.0	0.24
Mongolia	-37.8	0.3	0.09
Mozambique	-21.0	-0.4	-0.08
Mauritania	0.1	0.3	0.15
Malaysia	-84.7	-0.2	0.93
Namibia	0.0	0.3	0.07
Nigeria	-10.0	-0.2	-0.09
Peru	-68.7	0.0	0.05
Russia	-64.0	-0.2	0.04
Rest-of-World	-55.2	0.8	-0.14
Thailand	-23.2	0.2	-0.09
Tanzania	-43.5	-0.7	-0.03
Ukraine	-53.3	-0.1	0.22
United States	-48.5	1.9	-0.03
Venezuela	-37.9	-4.9	0.26
Vietnam	-21.1	-0.5	0.28
South Africa	-20.9	0.4	0.09
Zambia	-28.8	-0.3	-0.04
Zimbabwe	-13.3	-0.4	0.02

NOTES: This table shows the countries with the largest changes in emissions share under the global optimal carbon tax. % change shows the percent change in production. Share shows the baseline emissions share as percentage of global production. Real income is price index-deflated total income in that country.

D.4 Additional figures: biodiversity and co-benefits

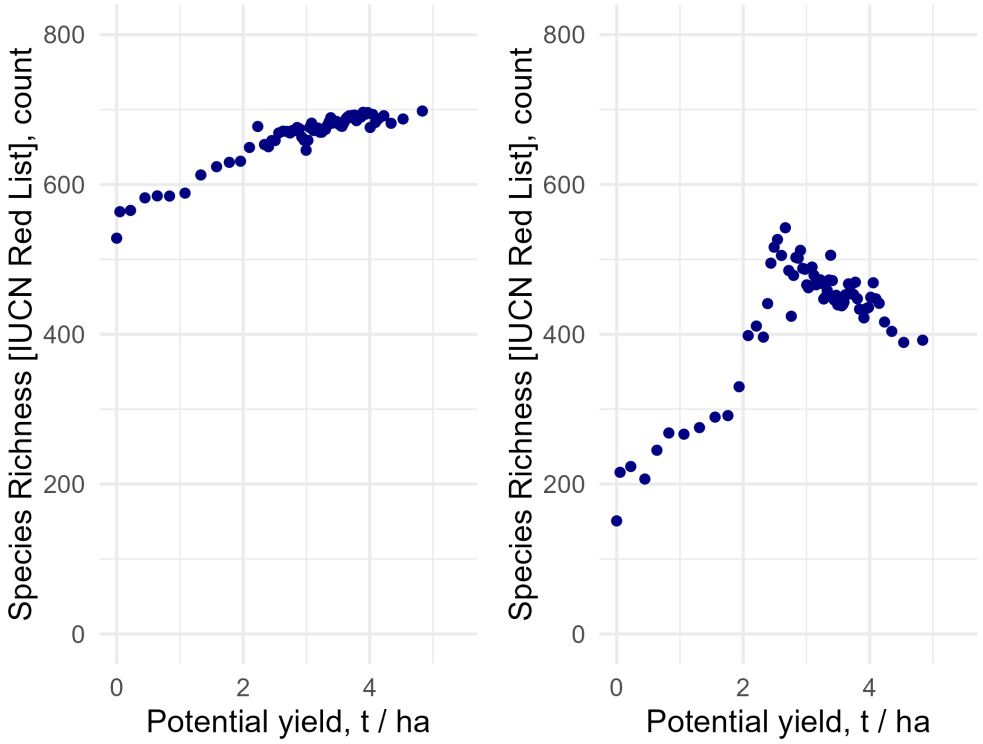

Figure A.6: Binned regression of agricultural yields and rare species counts among land deforested since 1982. Left controls for country fixed effects; right additionally controls for cropland share in 2000.

NOTES: Binned regression conducted on stratified random sample of 100 grid cells per country. Maximum agricultural potential yields indicate pointwise maximum at grid cell level across crops. Cuts out two bins above 1 t/ha.

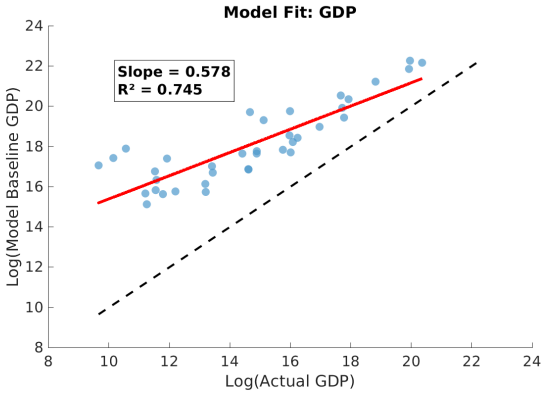

(a) Fit: model vs. data on GDP

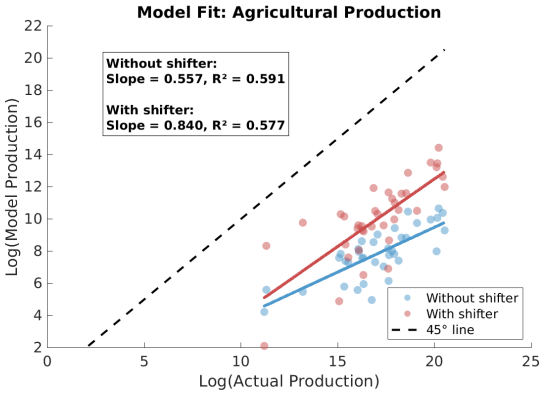

(b) Fit: model vs. data on crop output

Figure A.7: Fit of model on data for held-out moments.

NOTES: Dashed line indicates 45 degree line. Data on crop output sums output across included crops from FAO data. Data on GDP from Penn World Table, summing GDP in the rest-of-world.