

Does Conservation Work in General Equilibrium?*

Veronica Salazar Restrepo[†]

IIES

Gabriel Leite Mariante[‡]

CUNEF Universidad

August 3, 2026

Deforestation for agriculture accounts for roughly a fifth of global CO₂-equivalent emissions. Because most forest conservation policies are spatially targeted, their local gains can be undone in general equilibrium if the deforestation they prevent simply relocates—leakage. We quantify leakage for two flagship Brazilian policies, the Priority List and Protected Areas, using an open-economy quantitative spatial model in which costly deforestation endogenously supplies agricultural land, inverted on regional data and disciplined by micro estimates and reduced-form policy effects. Leakage is small: over 2003–2018, general-equilibrium adjustment undoes 3.7–8.8% of the deforestation avoided by the Priority List and 6.4–14.7% of that avoided by Protected Areas, and less in carbon terms. What makes leakage small is where conservation is targeted. Restricting regions where agricultural land is valuable leaves cheap suppliers available to absorb displaced demand; restricting regions where land is cheap to clear removes the cheapest suppliers, leaving residual supply steep and inducing intensification rather than relocation. Brazilian deforestation is predominantly supply-driven, explaining both low leakage and why the Priority List leaks less. This mechanism yields a model-free diagnostic of leakage risk: targeted regions’ share of the market value of newly cleared land relative to their share of cleared area.

*For their generous mentorship, we thank Oriana Bandiera, Gharad Bryan, Robin Burgess, and Daniel Sturm. For all their comments and questions, we thank seminar audiences at LSE’s Development Economics Work in Progress, LSE’s Energy & Environmental Economics seminar, LSE’s International Trade Workshop, the BQSE Seminar, and conference audiences at STEG-CEPR, LSE’s Environment Week, EMUEA, EAERE, and OxDev. For their time and insightful discussions, we thank Gabriel Ahlfeldt, Clare Balboni, Tim Doberman, Matthias Doepke, James Fenske, Douglas Gollin, Allan Hsiao, Clement Imbert, Juliana Jaramillo Echeverri, Xavier Jaravel, Per Krusell, Isabella Manelici, Prakash Mishra, Pol Simpson, Xavier Tierney, Fabian Bald, Julien Wolfersberger.

[†]veronica.salazar-restrepo@iies.su.se

[‡]g.leitemariante@cunef.edu

1 Introduction

Tropical deforestation is among the human activities with the highest environmental impact: forest clearing and the subsequent agricultural use of cleared land are responsible for about one-fifth of global CO₂-equivalent emissions over the past two decades (IPCC 2019). The most widespread instruments against it are spatially targeted: they intensify protection in specific areas, making agricultural land scarcer there. Precisely because they are local, these policies may leak—the deforestation they prevent in targeted areas may relocate to untargeted ones, undoing their aggregate effect. Leakage concerns are consequential: strategies to mitigate deforestation are excluded from major green finance mechanisms, such as the United Nations Clean Development Mechanism, on exactly these grounds.¹ Yet leakage is intrinsically hard to measure: it is a general-equilibrium phenomenon, and the reduced-form comparisons used to evaluate conservation policies assume it away.

Brazil offers an ideal setting to confront this problem. It holds a third of the world’s remaining rainforests, accounts for a third of yearly tropical deforestation, and over the past decades has implemented two large-scale, spatially targeted conservation policies with rich accompanying data. The first is the Priority List (2008–): municipalities with the highest deforestation rates were singled out for intensified command-and-control enforcement. The second is the continuous establishment of Protected Areas—Conservation Units and Indigenous Territories—where deforestation is banned outright and which, in contrast to the Priority List, typically cover regions with high forest cover and little ongoing deforestation.

We first document the local effects of both policies. That both bind locally is well established—for the Priority List by Assunção and Rocha (2019) and Cisneros et al. (2015), and for Protected Areas and Indigenous Territories by an extensive prior literature (Soares-Filho et al. 2010, Nolte et al. 2013, Baragwanath and Bayi 2020). Our estimates reproduce these findings and extend them to the margins the model requires. The Priority List sharply reduced deforestation in listed municipalities—total loss of natural vegetation falls by roughly 36%—while local agriculture did not contract: production shifted from pasture toward crops, value added per hectare rose, and cattle ranching densified. Through a regression discontinuity design at protection

¹The UN defines deforestation leakage as “*The unexpected loss of anticipated carbon benefits due to the displacement of activities in the project area to areas outside the project, resulting in carbon emissions.*”

boundaries, we find that Protected Areas have effectively conserved forest. We use these estimates as inputs to the analysis that follows rather than as a contribution in themselves: however well identified, they cannot measure leakage. Identification relies on untreated regions being unaffected by treatment elsewhere; if deforestation relocates, the estimates conflate reduced deforestation in targeted municipalities with displaced deforestation, and no reduced-form comparison can reveal how much relocates or where.

Answering that question requires a model of how agriculture is redistributed across space in response to local policies. We build an open-economy quantitative spatial model of Brazil in which deforestation is an economic sector: clearing forest is a costly activity with decreasing returns that endogenously supplies land as a factor of production for agriculture, so each region has an upward-sloping supply curve of new land whose position reflects its deforestation productivity—how cheaply its forest converts into farmland. On the demand side, regions produce differentiated agricultural varieties—temporary crops, permanent crops, and pasture-based cattle, each with a different land intensity—alongside a non-agricultural good; consumers have non-homothetic preferences over food; workers migrate across regions and switch sectors subject to frictions; and goods are traded domestically and with the rest of the world subject to trade costs, with world demand for Brazilian varieties less than perfectly elastic. Agricultural land accumulates through deforestation and depletes the remaining forest, so the model delivers a full dynamic path of land use region by region, and leakage emerges from prices: conservation policy tightens the market for new land, and where the displaced demand lands depends on who the remaining cheap suppliers are.

We take the model to the data in three steps. First, we estimate the key parameters outside the model from Brazilian micro data: the deforestation supply elasticity from a shift-share design using world commodity-price shocks, the land-demand system from land shares and expenditure data, and migration and trade frictions from bilateral flows. Second, we invert the model on the observed 2002–2018 panel of deforestation, land use, employment, and incomes, recovering each region’s productivity in agriculture, non-agriculture, and deforestation year by year—so the calibrated economy reproduces the observed deforestation exactly, and the recovered fundamentals can be validated against data the inversion never targets (land prices, incomes, and external correlates of clearing costs). Third, we discipline the policies themselves

with the reduced-form estimates: the counterfactuals impose the policy-induced supply shift estimated in Section 3, so the model answers a question the reduced form cannot: what national deforestation would have been had the policy never existed, with all price, migration, and trade responses switched on, and what it would have been under the policy had prices not responded at all. Leakage is the gap between the two.

Our headline results are that leakage is small, and smaller for the policy that targets the places with the highest observed deforestation, as is the case for the Priority List. Over 2003–2018 the Priority List avoided roughly 80,000 km² of deforestation gross of general-equilibrium effects, an order of magnitude more than the Protected Areas established over the same period. Leakage depends on the elasticity with which new land is supplied, and no single design pins that elasticity down, so we report headline leakage as a range rather than a point. The baseline is the all-Brazil estimate of the deforestation technology ($\delta = 0.41$), which matches the model’s coverage and implies a supply elasticity of new land of $\varepsilon_S = \delta/(1 - \delta) \approx 0.70$, close to the value calibrated for global deforestation by Farrokhi et al. (2025); the upper end is the Legal-Amazon frontier estimate ($\delta = 0.62$), implying $\varepsilon_S \approx 1.63$, so that clearing responds more than proportionally to the price of land. The range is thus generous rather than cautious: its upper end sits above the elasticities this literature has estimated, and because leakage is increasing in ε_S , that end bounds leakage from above. Across the range, cumulative leakage runs from 3.7% to 8.8% of avoided deforestation for the Priority List and from 6.4% to 14.7% for Protected Areas; we re-solve the full model at both ends wherever a quantitative claim depends on it, and Appendix Table 20 extends the sensitivity analysis to every other parameter of the model, one at a time (Section 7). Two conclusions are invariant across that range: leakage remains small—at least 80% of the deforestation reduction estimated without accounting for leakage remains—and the Priority List leaks less than Protected Areas. Carbon leakage is smaller still, because the targeted regions are among the most carbon-dense. A further clarification concerns geography: about 60% of the Priority List’s general-equilibrium effect occurs *within* the listed municipalities themselves—price responses that dampen local gains without moving clearing anywhere else—so clearing actually displaced to untargeted regions is smaller still.

The paper’s central contribution is to explain *why* leakage is small, and why it differs across the two policies, with one key economic property of the setting: Brazil-

ian deforestation is predominantly supply-driven, rather than demand-driven. In the model, spatial variation in deforestation decomposes exactly into a supply term—how cheap land is to clear—and a demand term—how valuable agricultural land is. We show theoretically that leakage is governed by this decomposition: holding observed deforestation fixed, leakage is large when targeted deforestation is demand-driven and small when it is supply-driven. The logic is akin to that of tax incidence. If the targeted regions deforest heavily because clearing is cheap there, a ban removes the economy’s cheapest suppliers of new land; the residual supply curve is steep, and the induced land-price increase contracts demand—through intensification and substitution toward land-light goods—rather than relocating clearing. If the targeted regions instead deforest heavily because their land is valuable, the cheap suppliers remain untargeted and production relocates to the next-best region. In the data, the big deforesters are poor, remote, pasture-intensive places with low revenue per hectare: regions with low clearing costs, not regions with high land demand. Deforestation productivity alone explains almost all of the spatial variation in deforestation—far more than the demand-side fundamentals—and regions with low clearing costs systematically have low agricultural land value.

An increasingly widespread availability of satellite-image data has contributed to a growing wave of empirical research on the economics of deforestation ([Balboni et al. 2023](#)). We relate first to the reduced-form literature evaluating conservation policies, and then to the structural literature on trade, space, and the environment.

The reduced-form literature finds that both instruments bind: blacklisting cut deforestation in listed municipalities by around 50% ([Assunção and Rocha 2019](#), [Cisneros et al. 2015](#))², and Protected Areas and Indigenous Territories conserve forest where they are established, with effectiveness varying by governance regime and location ([Soares-Filho et al. 2010](#), [Nolte et al. 2013](#), [Qin et al. 2023](#)) and, for Indigenous Territories, accruing once demarcation confers full collective property rights ([BenYishay et al. 2017](#), [Baragwanath and Bayi 2020](#)). This literature has increasingly attended to spillovers. [Pfaff and Robalino \(2017\)](#) summarize theory and evidence on conservation spillovers measured in the vicinity of protected areas; [Fuller et al. \(2019\)](#) systematically review the evidence on leakage from Protected Areas and find

²The Legal Amazon (*Amazônia Legal*) is an administrative region defined in Brazilian law, comprising the nine states of the Amazon basin and covering roughly five million km²—about 60% of Brazil’s territory. It is the standard unit for Brazilian deforestation policy and satellite monitoring, and the sample frame for our municipality-level reduced-form analysis in Section 3.

its incidence highly context-dependent; [Alix-Garcia et al. \(2012\)](#) estimate that leakage undoes about 4% of the reductions from a Mexican PES scheme; [Robalino et al. \(2017\)](#) document heterogeneity in the leakage of Costa Rican parks by market access. Identification of spillovers in reduced form, however, requires an ex-ante stance on where leakage can go—typically nearby—while the most suitable substitute for protected land may be in an entirely different part of the country. Our general-equilibrium approach internalizes exactly the spillovers that bias these designs, places no prior restriction on where displaced clearing lands, and—because we find leakage is small—largely vindicates the reduced-form toolkit in this setting. A related strand shows how deforestation responds to the demand for agricultural land, via commodity prices, credit, and market access ([Assunção et al. 2015](#), [Berman et al. 2023](#)); our supply-versus-demand decomposition builds directly on that insight.

Structurally, we contribute to the literature that models land use and deforestation in general equilibrium ([Costinot et al. 2016](#), [Souza-Rodrigues 2019](#), [Sotelo 2020](#), [Pellegrina 2022](#), [Pellegrina and Sotelo 2025](#), [Farrokhi and Pellegrina 2023](#), [Araujo et al. 2026](#), [Hsiao 2026](#)). That work has approached deforestation from scales above and below the one at which conservation policy is made. Above it, [Farrokhi et al. \(2025\)](#) ask how agricultural trade liberalization redistributes forest across countries and [Mishra \(2026\)](#) how far the global location of clearing departs from what a carbon price would deliver; neither evaluates a policy that was implemented. [Domínguez-Iino \(2026\)](#) follows the incidence of environmental regulation along agricultural supply chains, where the margin of interest is bargaining between farmers and processors rather than the location of clearing. Below it, [Tsuda et al. \(2025\)](#) model agglomeration and congestion externalities among neighboring communities on a single frontier. [Akerman \(2025\)](#) shares our country and asks a demographic question—how much forest Brazil’s fertility decline saved—over a horizon long enough for that transition to run its course.

Conservation policy, however, is assigned to sub-national jurisdictions while its consequences are absorbed by a national market for land and labor. That is the scale at which we work, and it is what lets us ask: how much of the deforestation that an implemented, spatially targeted policy prevents is undone by reallocation elsewhere in the same economy, and what property of the setting determines the answer. To that end we add labor markets, migration, and structural change to spatial models of land use ([Eckert and Peters 2022](#), [Lagakos and Waugh 2013](#), [Alvarez](#)

2020, Bustos et al. 2016), model deforestation as an endogenous land-supply sector rather than a discrete conversion choice (Roberts and Schlenker 2013, Scott 2013), and discipline the policy shock with the reduced-form design of Section 3. The result that leakage is smallest precisely for policies targeting the highest-deforestation regions—and the supply-driven mechanism behind it—is, to our knowledge, new; it is the general-equilibrium counterpart to the optimal targeting of enforcement characterized in partial equilibrium by Assunção et al. (2023). The mechanism also connects to work on deforestation and structural transformation in other frontier economies (Gollin and Wolfersberger 2025), where the margin we emphasize—whether clearing happens because land is cheap to clear or because it is valuable once cleared—governs how agricultural expansion responds to policy and to market access.

The remainder of the paper proceeds as follows. Section 3 documents the local effects of the two policies. Section 4 develops the theory of when leakage is small and shows, in a set of stylized facts, that Brazilian deforestation is supply-driven. Section 5 presents the quantitative model and its estimation. Section 6 reports the main counterfactual results and draws out the policy-design implications of targeting, strength, and horizon. Section 7 presents mechanism decompositions, robustness, and the policies’ effects on structural transformation and incomes. Section 8 concludes.

2 Data

Spatial units and land use We use data at different levels of spatial granularity. Land use information comes from satellite data collected by the MapBiomass project³ between the years of 1985 and 2022 at a native 30m $\times$ 30m resolution. We aggregate their original 28 detailed categories of land use into six: forests, other natural ecosystems, pasture, permanent crops, temporary crops, and other non-vegetated areas (e.g. cities, rivers, etc). For the reduced-form regression-discontinuity analysis of Protected Areas, we lay a fine 1 km hexagonal grid over the territory (each hexagon has an area of roughly 0.87 km²) and take each hexagon to be the unit of analysis. We code a cell as forest if the 30m MapBiomass class at its *centroid* is forest, giving a

³MapBiomass is a multi-institution collaborative initiative that processes satellite images into publicly available datasets on land use for the Brazilian territory

binary outcome whose conditional mean is the probability that a point is forested.⁴ For the quantitative model, we use data at the level of Brazilian micro-regions, considered to be broadly equivalent to local labor markets or commuting zones⁵. Data on Brazil’s administrative boundaries is publicly available via the Brazilian Institute of Geography and Statistics (IBGE).

Conservation policies We gather geo-coded information on the boundaries of localized conservation policies from Brazilian government sources. Information on Indigenous Territories comes from the National Foundation of Indigenous Peoples (FUNAI)⁶. Information on the Conservation Units, including their year of establishment, comes from the website of the Ministry of the Environment.

Environmental indicators We use data on environmental indicators from a variety of sources. Temperature and soil moisture data for the years 1985-2019 come from TerraClimate⁷. Precipitation data for years 2000-2019 comes from ERA5⁸. Above and Below Ground Biomass Carbon Density comes from the UN Environment Programme World Conservation Monitoring Centre (UNEP-WCMC) carbon biomass dataset, which contains a snapshot of biomass carbon density for 2010. We use these data, alongside land use data for 2010, to assign average carbon density to natural ecosystems in different micro-regions in order to approximate the carbon emissions under different simulated scenarios.

⁴Cells this small rarely straddle a Protected-Area boundary, and centroid sampling avoids the arbitrary threshold a within-cell majority rule would require.

⁵There are 557 micro-regions in continental Brazil (the model excludes the single oceanic-island micro-region, which has no land margin), and they are often considered to be equivalent to a US commuting zone, as the majority of people live and work within the same micro-region ([Dix-Carneiro and Kovak 2017](#))

⁶FUNAI provides information on the year in which they have undergone each of the five stages of approval (study, delimitation, declaration, homologation, and regularization). We consider the homologation as their year of creation, as this is the stage of judicial approval officially recognizing the territory as an indigenous land.

⁷TerraClimate is a publicly available global gridded dataset of monthly climate and climatic water-balance variables, constructed by combining high-resolution climatological normals with time-varying data from reanalysis products. The temperature data are the minimum and maximum monthly temperatures, and soil moisture is measured in mm and comes from a one-dimensional soil water-balance model; both are collected at a 4,638.3 m resolution.

⁸ERA5 is the fifth-generation global atmospheric reanalysis produced by the European Centre for Medium-Range Weather Forecasts (ECMWF). The data gathered are the monthly aggregate precipitation in mm at 11,132 m resolution.

Agricultural data We consider three types of agricultural activities: pasture, temporary crops and permanent crops. For each type, we use data from the 2006 Agricultural Census on area, number of workers employed, and revenues. To create a yearly panel of crop revenues, we use data from the Municipal Survey on Agricultural Production (PAM)⁹. For pastures, no direct data on revenues exists at a disaggregated level for the entire country, so we leverage cross-sectional data on other sources that cover the dimensions governing yearly revenue changes. We combine: (i) yearly estimates of the number of cows in each municipality from the Municipal Survey on Livestock Production (PPM), (ii) yearly national variation in the price of beef from a daily time series curated by the School of Agricultural Studies of the University of São Paulo¹⁰, and (iii) yearly state-level data on the average weight of slaughtered cows from the national statistical office’s Quarterly Survey of Slaughtered Animals¹¹.

Local GDP and labor markets For model validation, and as ancillary outcomes of the municipality-level reduced-form analysis of the Priority List municipalities, we use data on local GDP, employment and wages. Local GDP data for the years 2002-2019 are publicly available via the Brazilian Institute of Geography and Statistics (IBGE). Data on local wages comes from the population censuses of 2000 and 2010, where respondents self-report monthly income for different activities - we take aver-

⁹The PAM has yearly revenue estimates for 71 different crops, which we aggregate into either permanent or temporary. We adjust our yearly revenue estimates such that, for the year 2006, they are equal to revenue measures from the Agricultural Census, likely to be more accurate. We then estimate relative yearly variation in each municipality as:

$$\widehat{Revenue}_{rt}^k = Revenue_{r2006}^{k, AgCensus} \times \frac{Revenue_{rt}^{k, PAM}}{Revenue_{r2006}^{k, PAM}}$$

for each $k \in \{\text{permanent crops, temporary crops}\}$.

¹⁰Available for public consultation at: <https://www.cepea.esalq.usp.br/br/indicador/boi-gordo.aspx>

¹¹Given the average age of a cow at the time of slaughter is 3 years, we estimate yearly cattle revenues by multiplying the 3-year lagged change in weight, number of cows and price of beef by the 2006 revenue data from the Agricultural Census. This extrapolation over time requires the assumption that the yearly changes in the price of beef can be considered as constant across all municipalities, and that yearly changes on average slaughtered weight can be considered constant across all municipalities within a given state. For region r in state $s(r)$, we have:

$$\widehat{Revenue}_{rt}^{\text{pasture}} = Revenue_{r2006}^{k, AgCensus} \times \frac{p_t^{\text{beef}}}{p_{2006}^{\text{beef}}} \times \frac{kg/cow_{s(r)t-3}}{kg/cow_{s(r)2003}} \times \frac{\# cows_{rt-3}^{PPM}}{\# cows_{r2003}^{PPM}}$$

ages at the micro-region level split by agricultural and non-agricultural sectors. Data on sectoral employment is calculated by considering the number of workers in agriculture and non agriculture for both of these years and, for each micro-region, linearly extrapolating these values to approximate the time series from 2003 to 2019. We also use the 2010 census to obtain information on internal migration, as respondents are also asked in which municipality they lived 5 years prior to data collection.

Internal trade We use data on interstate trade flows at two points in time: 1999 and 2017. Data for 1999 comes from [de Vasconcelos \(2001\)](#) (as cited by [Morten and Oliveira \(2024\)](#)) as a country-level matrix of the inter-state trade in goods and services¹². Data for 2017 comes from the National Council of Fiscal Policy (CONFAZ), which relies on data from the Electronic Invoice (NF-E) to establish the interstate commercial balance¹³. Given potentially large gaps in trade flow records (in 1999, five states, mostly Amazonian, did not have any records of inter-state trade) we use these data to estimate a distance-elasticity of trade rather than to match observed trade flows. Moreover, our analysis is at the micro-region level, not at the state level, so we use the state-level-estimated distance elasticity of trade costs to approximate the iceberg trade costs that decay with distance. We provide more details in subsection [D.6](#).

Household expenditures Finally, we use the 2017-2018 Household Expenditure Survey (POF) to calibrate the preference parameter that governs non-homotheticities in individual preferences. This survey contains individual-level data on income and itemized expenditures, which we use to measure the correlation between individual income levels and the share of their expenditure that goes to food. We use household expenditure shares on food items as a proxy for consumption of agricultural goods.

¹²These data come from administrative records collected by the author from each state’s office. At the time, records were not centralized nor digitalized and it is likely that there were more gaps and inconsistencies.

¹³The NF-E contains details of goods departure, including its destination and value, as well as information about the entry of merchandise. CONFAZ aggregates these data to calculate interstate exports and imports.


3 Local Effects of Conservation Policies

Brazil’s success in slowing deforestation over the first two decades of the century is widely attributed to the policies adopted after 2004 (Burgess et al. 2019, Assunção et al. 2015). Some of these were “blanket-level” policies applying to the entire territory, such as the adoption of a unified *Action Plan for the Prevention and Control of Deforestation in the Legal Amazon* (PPCDAm) and changes in the legal limits for deforestation on private land (DOU 2012). Among the locally targeted instruments—the ones subject to leakage concerns—two stand out. The first is the “Priority Municipalities” policy, or Priority List, enacted in 2008 (DOU 2007): municipalities with the highest deforestation rates—initially 36 municipalities accounting for around 45% of the previous year’s deforestation—were singled out for a package of sanctions and scrutiny triggered by listing itself: priority in satellite-based monitoring and in IBAMA field operations, with more frequent inspections and fines; embargoes on farms caught clearing illegally, whose output the soy- and beef-supply-chain agreements then refuse to buy; stricter conditions on subsidized rural credit; and georeferenced land-registration requirements, with exit from the list conditional on deforestation falling below a threshold and on most private land being environmentally registered (Assunção and Rocha 2019). The list has been updated yearly, covering roughly 40 municipalities and one million square kilometers by the end of our sample period. The second is the continuous establishment of Protected Areas—Conservation Units and demarcated Indigenous Territories—where deforestation is banned outright. These have been established gradually since the country’s return to democracy in 1985 and, in contrast with the Priority List, typically cover regions with high forest cover and little ongoing deforestation. Figure 1 maps both policies; Appendix A.1 presents their evolution over time.¹⁴ The data used throughout the paper are described in Section 2.

This section documents the local effects of both policies, which motivate the theory and the model that follow. Three facts emerge. First, the Priority List sharply reduced deforestation in listed municipalities. Second, it did so while agricultural

¹⁴Given that these policies are effective locally, why not extend them to the whole country? Enforcement is fiscally costly—placing 80% of the Brazilian Amazon under existing policies would cost at least 1.7 billion USD per year (da Silva et al. 2022)—conversion of forest to agriculture generates private profits, implying a non-zero optimal level of deforestation from Brazil’s perspective, and restricting deforestation carries political costs, since its proceeds are often captured by local elites.

Figure 1: Spatial distribution of Priority List municipalities (green), Conservation Units (orange), and Indigenous Territories (yellow) as of 2021.



Note: This figure illustrates the spatial targeting of the three localized conservation policies studied in this paper over the past decades. The left panel shows, in green, the municipalities that have been added to the Priority List since its start in 2008. The right panel shows the location of protected areas (which do not necessarily coincide with municipal borders) according to its type: orange areas are Conservation Units, yellow areas are demarcated Indigenous Territories.

value added *rose*: production shifted from pasture toward crops and intensified on the remaining land. Third, Protected Areas bind—forest cover is discretely higher just inside their borders. What these local estimates cannot tell us is the aggregate effect of the policies. Identification relies on the stable unit treatment value assumption (SUTVA): untreated regions must be unaffected by treatment elsewhere. If restricting the supply of new agricultural land in targeted regions raises the demand for deforestation in the control regions—leakage—that assumption fails, and local estimates cannot separate reduced deforestation from relocated deforestation. Quantifying that requires the general equilibrium analysis of Sections 4–6; Appendix A.4 then uses the model to bound the resulting bias in this section’s estimates, and finds it small.

3.1 The Priority List

We estimate the effect of entering the Priority List with a synthetic difference-in-differences (SDiD) design (Arkhangelsky et al. 2021) on the panel of Legal Amazon

municipalities over 1995–2019:

$$(\log) \text{ Forest Loss}_{mt} = \delta_t + \gamma_m + \beta \text{Priority}_{mt} + \epsilon_{mt}, \quad (1)$$

where Priority_{mt} indicates that municipality m is on the list in year t , and δ_t and γ_m are year and municipality fixed effects.¹⁵ Because the policy explicitly targeted deforestation hotspots, treated municipalities were on different trends by design; the SDiD estimator reweights control units so that pre-treatment trends match, and the identifying assumption is that matched pre-trends would have continued absent the policy. Outcomes are in logs, so municipality-years with zero forest loss are dropped.¹⁶

Table 1 reports the results. Total loss of natural vegetation falls by 0.44 log points, or 36%, entirely accounted for by primary vegetation loss; regrowth of secondary vegetation rises, and the point estimate on secondary vegetation loss is positive though imprecise, consistent with more land cycling through fallow. Figure 2 plots the corresponding event study for the 2008 cohort: estimated effects are indistinguishable from zero in every pre-listing year, drop sharply to between -0.5 and -0.8 log points in the years following listing, and attenuate toward the end of the sample, when the federal administration reversed environmental enforcement and cut the budgets of environmental agencies (Burgess et al. 2019).

What the listing shock is. The Priority List did not arrive in isolation: its 2008 enactment coincided with Resolution 3,545, which conditioned subsidized rural credit on environmental compliance across the Amazon biome, followed the 2006 soy moratorium, preceded the 2009 cattle agreements, and sat within the broader escalation of PPCDAm enforcement. Our design absorbs the biome-wide component of these measures through the year effects and the synthetic controls, which are Legal-Amazon municipalities subject to the same blanket policies; what loads onto β is the com-

¹⁵Implemented with the `sdid` routine of Clarke et al. (2024), which handles staggered adoption without using already-treated units as controls. Standard errors are bootstrapped (50 repetitions) and clustered at the municipality level. The dynamic (event-study) version replaces $\beta \text{Priority}_{mt}$ with $\sum_{\tau} \beta_{\tau} \text{Priority}_{m,t-\tau}$.

¹⁶This is immaterial for the estimate: zero-loss years are only 63 of the 20,200 Legal Amazon municipality-years in the sample (0.3%), spread across 25 municipalities, all never-listed—no treated municipality has a zero-loss year. Because the synthetic difference-in-differences estimator is computed on a balanced panel, these 25 municipalities are simply excluded from the donor pool of controls (leaving 721 controls for the 2008 event study), which leaves every treated unit’s estimate unaffected.

ponent that binds *differentially* in listed municipalities. Much of that differential component is the Priority List by construction—embargoes, monitoring priority, and the credit and registration conditions are consequences of listing—but concurrent instruments that concentrate where deforestation is highest, such as the supply-chain agreements, may contribute to it. We therefore read β , and the counterfactuals built on it in Section 6, as the effect of the listing-anchored enforcement bundle concentrated on listed municipalities, rather than of the administrative act of listing alone. For the leakage question this distinction is second-order: general-equilibrium relocation depends on the size and location of the supply shock to targeted regions, not on the label of the instrument that produced it—though it does matter for extrapolating the counterfactuals to a hypothetical stand-alone listing policy.


Table 1: Effects of the Priority List on deforestation and regrowth (SDiD)

Outcome	ATT	Std. err.	Obs.
(log) total veg. loss	-0.445***	(0.119)	20,137
(log) primary loss	-0.443***	(0.131)	20,126
(log) secondary loss	0.161	(0.134)	19,390
(log) secondary gain	0.210**	(0.084)	20,093

Note: Staggered synthetic difference-in-differences estimates, as in equation (1), of the effect of the Priority List on deforestation and reforestation outcomes (total loss of natural vegetation, loss of primary vegetation, loss of secondary vegetation, and gain of secondary vegetation), all in logarithms. The sample includes municipalities in the Brazilian Legal Amazon, 1995–2019; municipality-years with a zero outcome are dropped by the log transformation. Standard errors are calculated via bootstrap with 50 repetitions, clustered at the municipality level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Agricultural intensification. The reduction in deforestation was not accompanied by a contraction of local agriculture—listed municipalities moved to less land-intensive production. Figure 3 shows event studies for agricultural land use and intensification. The area in pasture falls persistently after listing, while the area in crops rises. Agricultural value added per hectare of agricultural land rises by roughly 15%, and the number of cattle per hectare of pasture rises by roughly 20%: less land is used, more value is produced on it, and ranching, the most land-intensive activity, densifies. Appendix Table 10 reports crop-by-crop revenue, area, price, and yield responses; the expansion is concentrated in maize and soy.

Figure 2: Dynamic effects of the Priority List on deforestation




Note: This figure shows the synthetic difference-in-differences event study of (log) total loss of natural vegetation for the municipalities that joined the Priority List in 2008 (36 treated), with never-listed Legal Amazon municipalities as the donor pool (721 controls, balanced 1995–2019). Each point is the yearly gap between treated municipalities and their synthetic control, relative to the time-weighted pre-treatment baseline (Arkhangelsky et al. 2021). The vertical line marks entry into the Priority List; the shaded area is the 95% confidence interval from a bootstrap clustered at the municipality level (50 repetitions).

These responses discipline how we should think about the general equilibrium effects of conservation policy. Restricting the supply of new agricultural land raises its price, and producers respond on the margins available: economizing on land by substituting toward labor and other inputs, shifting toward land-saving activities, and—potentially—relocating land-intensive production to other regions. The first two margins are visible here; the third is exactly the leakage that local estimates cannot detect. That producers economize on land so readily when it becomes scarcer is a first hint that displaced clearing need not be large.

3.2 Protected Areas

Protected Areas do not follow municipal borders, so we estimate a spatial regression discontinuity in forest cover at protection boundaries. Working on a fine 1 km grid, we code each cell with a binary forest indicator (Section 2) and compare forest cover just inside and just outside protection boundaries, conditioning on a smooth (quadratic

Figure 3: Effects of the Priority List on agricultural land use and intensification




Note: Synthetic difference-in-differences event studies for the 2008 Priority List cohort, constructed as in Figure 2: never-listed Legal Amazon municipalities form the donor pool, points are yearly treated-minus-synthetic-control gaps relative to the time-weighted pre-treatment baseline, the vertical line marks listing, and shaded areas are 95% confidence intervals from a bootstrap clustered at the municipality level (50 repetitions). Panels (a) and (b) cover balanced municipalities over 1995–2019. Panel (c) shows (log) agricultural value added (from the system of regional accounts) per hectare of agricultural land, available 2002–2018. Panel (d) shows (log) cattle heads per hectare of pasture, starting in 2000, when herd coverage becomes sufficiently complete.

spline) function of distance to the border and on municipality-year fixed effects. The full specification, estimates, and validity tests are in Appendix A.3. The probability that a cell is forested is roughly 13 to 15 percentage points higher just inside a Protected Area than immediately outside it (Appendix A.3); Figure 4 shows the discontinuity, both in the raw data and after residualizing on municipality-year fixed effects.

We read this border contrast as evidence that protection binds, rather than as a causal estimate of its effect. The reverse experiment corroborates this reading: where protected status has been formally erased, deforestation accelerates (Moretz-Sohn 2026). A placebo exercise around the borders of *future* Protected Areas finds no discontinuity more than ten years before establishment, but a significant one within five years of it (Appendix Figure 14), and the border gap, though present before designation, visibly widens after it—consistent with lengthy demarcation disputes suppressing deforestation before formal establishment, and with prior occupation by indigenous communities whose livelihoods preserve the forest before official indigenous-territory status is granted. None of our quantitative results depend on a causal reading: the counterfactual analysis of Section 6 models Protected Areas as enforced removals of convertible land, using their observed near-zero internal deforestation, not the border estimate.

Figure 4: Discontinuity in forest cover at Protected Area borders



Note: Each dot is the mean share of land in forest among hexagons in 1 km bins of distance to the border of the nearest Protected Area, pooling all years 1985–2022; negative (positive) distances are inside (outside) the protected area, and the solid lines are quadratic fits on each side of the border. The left panel shows raw means; the right panel shows means of residuals from a regression of forest share on municipality-year fixed effects. Because Protected Areas are established gradually and are not, as a rule, revoked, a hexagon’s distance to the nearest border is essentially non-increasing over time.

From local estimates to general equilibrium. Both designs identify *local* effects: a comparison of treated regions with controls assumed to be unaffected by the policy. If conservation policy displaces deforestation toward untreated regions, the

controls are treated too, local estimates overstate the policy’s aggregate effect, and no reduced-form comparison can reveal where—or how much—deforestation relocates. The remainder of the paper studies the drivers of deforestation leakage and quantifies it for these two policies.

4 Theory

The pessimistic prior about leakage comes from a “flat-economy” benchmark: if all land were equal, workers and goods moved freely, and the demand for agricultural land were perfectly inelastic, banning deforestation in one region would have no aggregate effect—clearing would relocate one-for-one, and leakage would be 100%. This section shows which departures from that benchmark determine leakage, and derives the paper’s organizing result: *leakage depends on whether the targeted deforestation is supply-driven or demand-driven*. Throughout, $\delta \in (0, 1)$ is the elasticity of new land with respect to the resources invested in clearing it—the degree of decreasing returns in the deforestation technology introduced below—so that $\varepsilon_S \equiv \delta/(1 - \delta)$ is the implied price elasticity of deforestation supply; ε_D denotes the price elasticity of the aggregate demand for newly cleared agricultural land. Proofs and derivations are in Appendix B.

4.1 A deforestation supply equation

Suppose that clearing forest in region r is a competitive activity with decreasing returns: investing I_r^D units of the local final good, whose price is p_r , yields $T_r^D = Z_r^D (I_r^D)^\delta$ units of new agricultural land, which sells at the local land price q_r . Here Z_r^D is the region’s *deforestation productivity*—how cheaply a hectare of forest is turned into farmland, capitalizing soil and terrain suitability, distance to roads and markets, revenues from clearing itself, and the intensity of enforcement. (Section 5 microfoundations this technology and embeds it in the full quantitative model; free entry delivers equation (6) there.) Optimal deforestation then satisfies, in logs,

$$\ln T_r^D = \underbrace{\frac{1}{1-\delta} \ln Z_r^D}_{\text{supply: local clearing cost}} + \underbrace{\varepsilon_S \ln \left(\delta \frac{q_r}{p_r} \right)}_{\text{demand: land rent relative to prices}}, \quad \varepsilon_S = \frac{\delta}{1-\delta}. \quad (2)$$

Cross-sectional variation in deforestation decomposes *exactly* into a supply term—the deforestation productivity Z_r^D —and a demand term—the real value of agricultural land q_r/p_r , which embeds agricultural productivity, market access, and trade. The question “how big is leakage?” reduces to “which term generates the spatial variation in T_r^D among the *targeted* regions?”

4.2 A frictionless benchmark

Strip the economy to its core: goods are perfect substitutes across origins and labor is freely mobile, so there is a single national market for newly cleared land, clearing at one real price v . Deforestation is then a *pure* Z^D phenomenon—a region deforests more if and only if clearing is cheaper there. Let national demand for newly cleared land have price elasticity $\varepsilon_D > 0$ (unpacked in Section 4.4), and consider a ban covering a set of regions that account for the *targeted deforestation share* $b \in (0, 1)$ of baseline national clearing. Leakage L is the increase in non-targeted deforestation per unit of deforestation directly removed.

Proposition 1 (Leakage in the frictionless economy). *To first order in the policy,*

$$L = \frac{(1 - b) \varepsilon_S}{(1 - b) \varepsilon_S + \varepsilon_D} \quad (3)$$

Leakage is (i) increasing in the deforestation supply elasticity $\varepsilon_S = \delta/(1 - \delta)$; (ii) decreasing in the aggregate land-demand elasticity ε_D ; and (iii) decreasing in the targeted share b .

Part (iii) is the targeting result in analytic form: *for a policy of given scope, targeting the regions with the most deforestation maximizes b and thereby minimizes leakage per unit of avoided deforestation.* The logic is standard tax incidence in the national market for new land. The ban removes the targeted regions’ share b of supply, leaving residual price-responsiveness $(1 - b)\varepsilon_S$. Because the biggest deforesters are, in this benchmark, necessarily the highest- Z^D , cheapest-clearing regions, what remains is expensive, inelastic supply: the induced land-price rise chokes off demand—intensification, substitution toward land-light goods—instead of relocating clearing. Target expensive suppliers instead, and the cheap, price-elastic ones remain to absorb the displaced demand, so leakage is large. Demand elasticity alone caps leakage below

one: even an arbitrarily small policy ($b \rightarrow 0$), which faces the most elastic residual supply and is therefore the worst case for leakage, leaks only $\varepsilon_S/(\varepsilon_S + \varepsilon_D)$.

Remark 1 (Why leakage grows with the horizon). The demand relevant for equation (3) is the demand for *newly cleared* land—a flow—whereas farmers demand the *stock* of agricultural land. If stock demand has elasticity ε^{stock} , the implied flow-demand elasticity over a horizon of H years is of order $\varepsilon_D(H) \approx \varepsilon^{stock} \cdot T^A/(H \bar{T}^D)$. In Brazil the stock-to-flow ratio $T^A/\bar{T}^D \approx 53$, so at short horizons ε_D is enormous and leakage negligible; as the horizon lengthens the effective demand elasticity falls and leakage rises toward its long-run level. This is exactly the time path of cumulative leakage the full model produces endogenously through land accumulation (Section 6).

What the frictionless benchmark cannot address is the source-of-variation question: with a single national land price, deforestation is supply-driven *by construction*. Distinguishing supply- from demand-driven deforestation is intrinsically a story about frictions.

4.3 Imperfect substitution

Now let goods be imperfect substitutes across origins (elasticity $\sigma > 1$), as in the quantitative model. Land prices v_r are then regional, so a region can deforest heavily for *either* of two reasons: cheap clearing (high Z_r^D) or valuable land (high v_r , driven by high agricultural productivity or market access). This is what makes the supply- and demand-driven cases distinguishable. Take the policy target (region 1) with baseline deforestation ratio $x \equiv T_1^D/T_2^D > 1$ relative to an untargeted region 2. The same observed x is consistent with a continuum of underlying economies:

Lemma 1 (Observationally equivalent economies). *Fix $x > 1$. For each relative land price $\omega \equiv v_1/v_2 \in [x^{-1/\sigma}, x^{1/\varepsilon_S}]$ there is exactly one pair of fundamentals $(Z_1^D/Z_2^D, Z_1^A/Z_2^A)$ consistent with equilibrium ratio x . The lower bound $\omega = x^{-1/\sigma}$ is the pure supply-driven economy ($Z_1^A = Z_2^A$): region 1 deforests more and its land is cheaper. The upper bound $\omega = x^{1/\varepsilon_S}$ is the pure demand-driven economy ($Z_1^D = Z_2^D$): region 1 deforests more because its land is more valuable. In between, ω is strictly increasing in the demand-driven share of $\ln x$.*

The economies in Lemma 1 are indistinguishable to a reduced-form observer who sees only deforestation. They differ in one equilibrium observable—the relative price of land—and, as the next result shows, imply very different leakage.

Proposition 2 (Leakage and the source of variation). *Consider a marginal conservation policy in region 1 (a small reduction in Z_1^D). Leakage, defined as $L = -dT_2^D/dT_1^D$, equals*

$$L = \frac{1}{x} \cdot \frac{\varepsilon_S (\sigma - \varepsilon_D) \lambda_1}{\sigma(\varepsilon_S + \varepsilon_D) - \varepsilon_S (\sigma - \varepsilon_D) \lambda_1}, \quad (4)$$

where λ_1 is region 1's share of agricultural expenditure—equivalently, as shown below, its share of the market value of newly cleared land. Provided $\sigma > \varepsilon_D$, leakage is positive and, holding the observed deforestation ratio x fixed, L is strictly increasing in λ_1 —hence in the demand-driven share of region 1's deforestation. As $\sigma \rightarrow \infty$, (4) converges to the marginal ($b \rightarrow 0$) frictionless formula of Proposition 1: the limit forces $\omega \rightarrow 1$ in Lemma 1, so $\lambda_1/\lambda_2 \rightarrow x$ and the expenditure share ceases to carry independent information.

Corollary 1 (A sufficient statistic). *Given the targeted regions' share of deforested area, leakage is governed by their share of the market value of newly cleared land, λ . Supply-driven targets are big in quantity but cheap (λ low relative to area share $\Rightarrow$ low leakage); demand-driven targets are expensive (λ high $\Rightarrow$ high leakage). The relative land price of targeted regions is thus an observable diagnostic for the leakage risk of a conservation policy—computable, as Section 4.5 shows, without estimating the model at all.*

It is worth being explicit about what λ is, since the whole result turns on it. Land is the only factor here, so market clearing reads $v_r T_r = \lambda_r X$: the CES share of agricultural expenditure falling on region r 's variety is also region r 's share of the market value of the land it clears,

$$\lambda_r = \frac{v_r T_r}{v_1 T_1 + v_2 T_2}.$$

The λ_1 of (4) therefore measures how much of the economy's total willingness to pay for newly cleared land is directed at the targeted region. Its natural comparator is that region's share of cleared area, $b = T_1/(T_1 + T_2)$ —the object that governs leakage in the frictionless benchmark of Proposition 1, where a single national land price makes the two shares identical. With regional land prices they come apart, and their

ratio is exactly the relative price of the targeted land,

$$\frac{\lambda_1}{b} = \frac{v_1}{\bar{v}}, \quad \bar{v} \equiv \frac{v_1 T_1 + v_2 T_2}{T_1 + T_2},$$

where $\bar{v}$ is the average value of a hectare cleared anywhere in the economy. This is the content of Corollary 1 and the reason the diagnostic is observable rather than structural: a region that supplies many hectares but little value ($\lambda_1/b < 1$) is one where land is cheap to clear, the supply-driven case; a region whose hectares are worth more than the average ($\lambda_1/b > 1$) is being cleared because its land is valuable, the demand-driven case. Both objects are measured, not calibrated—Section 4.5 computes λ/b for the Priority List from observed clearing and agricultural-census land values alone.

The sign condition. The factor $\sigma - \varepsilon_D$ in (4) is worth dwelling on, because it delimits the regime the paper is in. Two forces act on an untargeted region when conservation makes land scarcer in region 1. Substitution across origins, governed by σ , pulls production toward region 2 and raises its clearing. Contraction of aggregate demand, governed by ε_D , works the other way: the composite price of agricultural land rises, and the economy buys less newly cleared land *in total*. Which dominates is exactly $\sigma \gtrless \varepsilon_D$. If demand for land were more price-elastic than goods are substitutable across origins, leakage would be *negative*—conservation in one region would lower clearing everywhere, a conservation multiplier rather than a leak. That is a real regime of the model, not a curiosity, and it is the reason the positive leakage we measure is a finding rather than an assumption. In the calibration the inequality holds with room to spare, for a structural reason. Demand for land is a *derived* demand, so to first order its elasticity is a cost-share-weighted average of two objects: the elasticity of substitution between land and labor in production—unity, since technology is Cobb–Douglas—and the price elasticity of demand for the goods that land produces. The latter is itself a share-weighted average of substitution across agricultural goods ($\theta = 2$) and world demand for Brazilian varieties ($\zeta = 6$), damped by a weak income margin (Engel elasticity $\eta \approx 0.5$). An average of terms none of which exceeds $\zeta = 6$ cannot exceed ζ , so $\varepsilon_D < \sigma = 9.6$. Two caveats keep this heuristic rather than a theorem: the weighting is exact only to first order, and ε_D here is the elasticity of demand for the land *stock*, as in the static economy of this section. The check that

matters is therefore the model itself, which delivers strictly positive leakage in every cell of the parameter sweep of Section 7—including at $\sigma = 6$, the lowest value we consider.

The economics mirrors the frictionless incidence logic, now operating through expenditure. Banning a *cheap* (supply-driven) target releases little expenditure relative to the area removed, and what it releases lands on regions where land is more expensive and supply steeper. Banning an *expensive* (demand-driven) target releases a lot, and the cheap suppliers are still standing to absorb it—so leakage can be large enough that demand-driven conservation *increases* total deforested area.

4.4 The demand elasticity ε_D

Both propositions assign a central role to ε_D , the price elasticity of aggregate demand for new agricultural land. In the full model ε_D is not a parameter but an equilibrium object, built from consumer substitution away from food (weak, given PIGL preferences with Engel elasticity $\eta \approx 0.5$); substitution across agricultural goods of differing land intensity, so that when land gets expensive consumption and production shift from land-intensive pasture-based beef toward land-light crops; factor substitution within each activity (intensification—more labor and inputs per hectare); and world demand for Brazilian varieties (ζ , Appendix E). The middle two channels are not hypothetical: they are exactly the adjustment margins the reduced form finds *inside* Priority List municipalities—pasture area falls, crop area and yields rise, cattle herds intensify (Section 3.1, Figure 3). That intensification is direct evidence that ε_D is substantial—one of the two pillars of small leakage. The other pillar—supply-driven targeting—is an empirical claim about Brazil, to which we now turn.

4.5 Stylized facts: Brazilian deforestation is supply-driven

The theory just developed reduces the leakage question to an empirical one: does spatial variation in deforestation come from the supply side (heterogeneity in how cheap land is to clear, Z_r^D) or the demand side (heterogeneity in the value of agricultural land, q_r/p_r , driven by agricultural productivity and market access)? This subsection presents the raw spatial patterns, which point toward the supply side: Brazil’s big deforesters are poor, remote, pasture-intensive places where land is cheap. Section 5 then makes the claim precise using the model’s inverted fundamentals—a correlates

regression on the recovered supply and demand terms—once we have described how those fundamentals are recovered.


Deforestation in Brazil is strikingly concentrated in space: out of the 557 micro-regions, the top 20 account for almost 40% of total yearly deforestation. Figure 5 shows what these places look like. Areas with more deforestation have lower GDP per capita (panel a) and lower market access (panel b).¹⁷ They have lower agricultural revenue per hectare (panel c), and a higher share of agricultural area dedicated to pasture (panel d), the most land-intensive activity. Every one of these correlations is the signature of the supply-driven case: low q_r/p_r , high Z_r^D . If Brazilian deforestation were demand-driven, the deforestation frontier would look like the opposite—productive, well-connected regions with expensive land.

The sufficient statistic of Corollary 1 can, moreover, be computed directly from observed data, with no model object anywhere in the calculation. Weighting each micro-region’s pre-policy clearing (2003–2007) by its 2006 agricultural-census land value, the regions that would enter the Priority List account for 36% of Brazil’s cleared area but under 20% of the market *value* of that clearing: an observed value-to-area ratio of 0.55, with the targeted hectares worth less than half ($0.44\times$) the average cleared hectare outside the List. A same-size set of the *highest*-land-value regions shows what demand-driven targeting would look like in the same data: 10% of clearing value on only 1.6% of cleared area, a ratio of 6.3. In the raw data, then, Brazil’s flagship policy targeted precisely the cheap-clearing configuration that Proposition 2 says leaks least—a fact established before any model is estimated.

These are only correlations and shares in raw data. Section 5 makes the claim precise once the model’s fundamentals are in hand: deforestation productivity alone explains most of the spatial variation in deforestation, while the demand-side fundamentals add little (Table 4); and the inverted fundamentals reproduce the observed value-to-area ratio of the Priority List (0.58 against the raw-data 0.55) without ever targeting it—the sufficient statistic is measured, not calibrated. By Propositions 1 and 2, a policy that targets these regions—as the Priority List does—should generate little leakage; Sections 5 and 6 quantify how little.

¹⁷Market access is defined by the first order approximation in Donaldson and Hornbeck (2016), $MA_d = \sum_o (\tau_d^o)^{1-\sigma} L_o$, where τ_d^o is the iceberg trade cost between o and d as estimated in Appendix D.6.

Figure 5: High-deforestation regions are poor, remote, low-revenue, and pasture-intensive



Note: This figure illustrates the correlation between the average yearly deforestation rate and economic characteristics: log-GDP per capita (panel a), market access (panel b), agricultural productivity (panel c, defined as the ratio between agricultural revenue and agricultural area), and the share of agricultural land dedicated to pasture (panel d). The dots represent the mean of the y-axis variable for 50 bins of log-yearly deforestation.

5 Quantitative Model

The theory of Section 4 identifies what governs leakage; measuring it requires an equilibrium model rich enough to be taken to Brazilian data. This section presents that model. Relative to the stripped-down economy of Section 4, the quantitative model adds every friction and margin that could plausibly *amplify* leakage or reroute it: costly trade across 557 micro-regions and with the rest of the world, migration subject to frictions, occupational choice with sectoral wage gaps, three agricultural sectors with heterogeneous land intensities, non-homothetic preferences, and the accumulation of agricultural land over time. Because Brazil's deforestation-linked com-

modities are largely exported, the model is an *open* economy: the rest of the world is an explicit trading partner that bids for Brazilian output at the world market, so conservation policy can relocate export supply internationally as well as within Brazil. The rest-of-the-world block, its calibration, and the open-vs-closed comparison are detailed in Appendix E; the closed economy is the special case in which the foreign block is switched off. Model details and derivations are in Appendix C.

Overview Brazil is modeled as a sequence of static spatial equilibria linked by the laws of motion of land and labor. The economy consists of 557 micro-regions¹⁸ indexed by r , which differ in their sectoral productivities, land endowments, and amenities, plus a rest-of-the-world region. There are $K + 1$ final good sectors: $K = 3$ agricultural commodities (pasture-based cattle products, temporary crops, permanent crops) that use land and labor with different land shares, and one non-agricultural sector that uses only labor. Deforestation is a sector that uses a composite investment good to produce agricultural land. Trade is subject to iceberg costs; consumer preferences are non-homothetic (PIGL, Boppart 2014); final goods are CES composites of regional varieties (elasticity σ) and the agricultural good a CES composite of the K commodities (elasticity θ).

5.1 Land use change: deforestation and regrowth

Each region r is endowed with L_{r0} workers, T_{r0}^A units of agricultural land, and T_{r0}^N units of terrestrial natural ecosystems which can be converted into agricultural land. Assuming a fixed forest regrowth rate ρ , and letting T_{rt}^D be the level of deforestation, agricultural land evolves according to

$$T_{rt+1}^A = T_{rt}^A(1 - \rho) + T_{rt}^D.$$

For most of this section we omit the time subscript and treat the equilibrium as static.

¹⁸This is typically considered the aggregation level most closely associated with a local market; see for example Dix-Carneiro and Kovak (2017).

5.2 The market for deforestation

Deforestation is a costly investment in the production of agricultural land. Deforesters are atomistic agents in a perfectly competitive market who access forested land, pay the cost of clearing it, and sell it at the price of agricultural land. If an amount I_r^D of the final good, bought at price p_r , is invested in deforestation, it delivers

$$T_r^D = Z_r^D (I_r^D)^\delta, \quad (5)$$

units of agricultural land, where Z_r^D is the region-specific deforestation productivity and $\delta \in (0, 1)$ governs the (within-year decreasing) returns to scale: forest closer to the edge is cheaper to access and clear, so each additional dollar spent deforesting is less productive.

The market has two defining features. First, land in natural ecosystems is open access, so the future value of the forest is fully discounted.¹⁹ Second, there is free entry of deforesters: entry occurs until marginal cost equals the value of agricultural land q_r , so that

$$T_r^{D*} = \underbrace{(Z_r^D)^{\frac{1}{1-\delta}}}_{\text{Local Factors}} \underbrace{\left(\frac{\delta q_r}{p_r}\right)^{\frac{\delta}{1-\delta}}}_{\text{Equilibrium Effects}}. \quad (6)$$

This is the supply equation (2) that anchors Sections 4 and 4.5. The deforestation productivity Z_r^D reflects local environmental conditions, enforcement of anti-deforestation policies, revenues from the act of deforestation itself (e.g. timber), and the option value of standing forest. The current area of unprotected forest is a strong predictor of Z_r^D , and a log-log model describes their relationship remarkably well, so we let

$$Z_r^D = \overline{Z_r^D} (T_r^F)^\psi.$$

Deforestation as claim-staking This market structure is deliberately tailored to how frontier clearing actually happens. Much of Brazilian deforestation is not a titled landowner’s choice of how to use his own land but the act by which possession

¹⁹Around 50% of deforestation in the Amazon over the past few years has happened in untitled public lands. Of the remainder, about half has happened in rural settlements and half in private properties—and even within private property, rights over forested land are insecure, as forest was historically regarded as “unproductive” and thus subject to the ownership claims of squatters. See [Alston et al. \(1999\)](#) for an in-depth description of property rights in the Brazilian Amazon frontier.

of untitled public land is established: clearing is what converts open-access forest into a marketable claim (Alston et al. 1999). Modelling deforestation as a specialized, competitive clearing sector that produces agricultural land and sells it at q_r —to landowners who then rent it to agricultural producers—captures this frontier division of labor, in which those who clear are often not those who ultimately farm. The natural alternative, a forward-looking landowner solving for the optimal timing of conversion of his own forest, presumes exactly the secure property rights over standing forest that the frontier lacks; under open access the future value of the forest is competed away, which is why free entry and full discounting are the defining institutional features of the margin we model rather than approximations to it. One implication deserves transparency. The return to claim-staking includes the expected value of eventual formalization: where cleared public land is more likely to be converted into private ownership, clearing is more profitable at any given land rent, and this capitalizes into the recovered Z_r^D . We make no attempt to hide that Z_r^D is an umbrella term for heterogeneous factors—terrain and climate, enforcement, timber revenues, and titling expectations alike (Section 5.6 shows that observable climate, biomass, and enforcement account for about 40% of its variation). The crucial point for the paper’s decomposition is the boundary, not the composition: Z_r^D collects what cannot be rationalized by the demand-side fundamentals—the productivity of agricultural land and access to markets—and it is that boundary, supply versus demand, on which the leakage theory of Section 4 turns.

5.3 Technology, workers, and preferences

Production The local non-agricultural good is Cobb-Douglas in labor and a complementary fixed factor H_r (capital and structures) with labor share γ_{NA} , $Y_r^{NA} = Z_r^{NA}(H_r)^{1-\gamma_{NA}}(L_r^{NA})^{\gamma_{NA}}$, so non-agricultural value added is non-agricultural labor income scaled by $1/\gamma_{NA}$ (the analogue of $va_r^{Ak} = v_r T_r^{Ak}/\alpha_k$ in agriculture). The fixed-factor rents $(1 - \gamma_{NA})$ of value added accrue to local residents and so are part of regional income and welfare. Agricultural goods k are Cobb-Douglas in land and labor with regional productivities Z_r^{Ak} and land shares α_k :

$$Y_r^{Ak} = Z_r^{Ak}(L_r^{Ak})^{1-\alpha_k}(T_r^{Ak})^{\alpha_k}.$$

All markets are perfectly competitive. The rental rate of agricultural land v_r equals its marginal product; under adaptive expectations (agents assume the future rental rate equals today's) and discount rate β , land is priced at its expected present value, $q_r = \frac{1}{1-\beta(1-\rho)}v_r$.

Occupational choice and wage gaps To allow for the agricultural/non-agricultural wage gaps we consistently find in the data, workers draw idiosyncratic sectoral productivities (z_i^A, z_i^{NA}) from a joint (Frank copula) distribution as in [Lagakos and Waugh \(2013\)](#) and [Alvarez \(2020\)](#), and choose the sector with the higher *amenity-weighted* effective wage $G_{s,r} w_{s,r} z_i^s$, where $G_{s,r}$ is a region–sector amenity (a taste term entering the choice, not output) normalized to $G_{A,r} \equiv 1$. The amenity rationalizes the portion of the agricultural/non-agricultural wage gap that selection alone cannot: with $G_{NA,r} < 1$ a compensating agricultural amenity keeps workers in agriculture despite higher non-agricultural pay. Since the draws are correlated, average incomes are not equalized across sectors. Appendix [C](#) gives the resulting labor supplies.

Migration As in [Eckert and Peters \(2022\)](#), individuals born in origin o choose a destination d given migration costs μ_{od} , destination utility $V(e_d, p_d)$, amenities B_d , and idiosyncratic Frechet taste shocks with dispersion ϵ , yielding gravity migration shares and the law of motion of population.

Preferences Consumers have PIGL preferences over agricultural and non-agricultural goods ([Boppart 2014](#)), with indirect utility

$$V(e, \vec{p}) = \frac{1}{\eta} \left(\frac{e}{(p^A)^\phi (p^{NA})^{1-\phi}} \right)^\eta - \nu \log \left(\frac{p^A}{p^{NA}} \right), \quad (7)$$

so that the agricultural expenditure share declines with real income with Engel elasticity η . The agricultural good is a CES aggregate of the K commodities with elasticity θ , and each final good a CES aggregate of regional varieties with elasticity σ , traded subject to symmetric iceberg costs $\tau_d^o = (\text{distance} + 1)^\kappa$, so that the share of good g consumed in region d that is produced in region o is

$$\pi_d^{go} = \frac{(\tau_d^o)^{1-\sigma} (p^{go})^{1-\sigma}}{\sum_{r=1}^R (\tau_d^r)^{1-\sigma} (p^{gr})^{1-\sigma}}. \quad (8)$$

We take $\sigma = 9.6$ from Domínguez-Iino (2026) (his across-county elasticity for the same setting)—at the higher end of the literature, giving trade and spatial spillovers ample scope—and report robustness to alternative values.

Market clearing Revenues of each origin-good pair equal domestic expenditures across all destinations plus rest-of-world expenditure on that variety:

$$p^{go}Y_o^g = \sum_{d=1}^R \pi_d^{go} X_d^g + \pi_{RoW}^{go} X_{RoW}^g, \quad (9)$$

where the trade shares π_d^{go} follow from CES demand and foreign demand responds to prices with elasticity ζ (Appendix E). A competitive equilibrium is a set of prices, wages, land rents, occupational choices, deforestation levels, and expenditure shares such that consumers and firms optimize, marginal products of labor are equalized across sectors within the occupational-choice structure, and all markets clear with trade balance; Appendix C states the definition formally.

5.4 Taking the model to the data

We apply two data-driven strategies. First, we calibrate the model’s structural parameters through empirical analyses that are consistent with the model structure but do not rely on its equilibrium solution (Table 2). Second, we invert the model to structurally estimate the regional fundamentals—the TFPs of the various activities, including deforestation, and amenities—so that each year over 2003–2018 the model exactly rationalizes observed deforestation, land use, labor shares, and agricultural revenues (Appendix C.3). Appendix D.9 shows that the calibrated model fits non-targeted moments—land prices, labor income, sectoral value added—well.

The deforestation supply elasticity δ The load-bearing parameter is δ : by Proposition 1, leakage is increasing in $\varepsilon_S = \delta/(1 - \delta)$, the supply elasticity of new agricultural land. From equation (6), ε_S is the elasticity of deforestation with respect to real land rents, which we estimate from

$$\log(T_{rt}^D) = \alpha_r^1 + \alpha_t^2 + \frac{\delta}{1 - \delta} \log v_{rt} + \epsilon_{rt}, \quad (10)$$

Table 2: Summary of structural parameters calibration

Parameter		Value	Source/Method
<i>Deforestation Function</i>			
δ	Deforestation returns to scale	0.41	Shift-share IV of deforestation on land rents (Appendix D.1)
ψ	Elasticity of deforestation TFP to natural area	0.26	Regression derived from steady-state deforestation
ρ	Forest regrowth rate	0.01	Observed reforestation rates
<i>Final Goods Production Functions</i>			
$\{\alpha_k\}_k$	Land share in ag. activities	(0.74;0.47;0.85)	From 2006 ag. census (temp. crops; perm. crops; pasture)
γ_{NA}	Non-ag. labour share	0.66	Calibrated to the standard $\approx 2/3$ labour share (Gollin 2002, Guerriero 2019) (Appendix D)
χ_A, χ_{NA}, ι	Workers' sectoral productivity shocks distribution	(1.55;1.52;12.8)	Alvarez (2020)
<i>PIGL preference parameters</i>			
ϕ	Ag. share in price index	0.01	Eckert and Peters (2022)
ν	PIGL Preference parameter	0.12	Eckert and Peters (2022)
η	Engel elasticity	0.506	Expenditure Survey Data (2017/18)
<i>Elasticities of substitution</i>			
σ	Between origins (regions)	9.6	Domínguez-Iino (2026)
ζ	Rest-of-world (foreign) demand	6	Domínguez-Iino (2026)
θ	Between ag. goods	2	Costinot et al. (2016), Dominguez-Iino (2023)
<i>Trade and Migration parameters</i>			
$\{\mu_d^o\}_{o,d}$	Bilateral migration utilities	residuals	Migration flows 2005-2010
ϵ	Dispersion of tastes	1.22	Migration flows and incomes from 2010 census
κ	Trade costs distance elasticity	0.105	Gravity equation of trade flows ($\sigma = 9.6$)
<i>Other Parameters</i>			
β	Discount rate	0.95	NA

where the municipality and year fixed effects control for constant differences between municipalities and for macroeconomic shocks common to all regions, and the standard specification additionally conditions on the lagged remaining stock of natural vegetation—a predetermined depletion control that absorbs the post-frontier margin (municipalities with little forest left cannot respond to rents) and licenses estimation on all Brazilian municipalities rather than the Legal Amazon alone. Because rents are an equilibrium object, our preferred estimates instrument them with a shift-share (Bartik) index that isolates demand-driven variation—world commodity prices interacted with each municipality’s predetermined base-year crop mix. For all Brazilian municipalities over 2008–2018 the IV land-supply elasticity is 0.68 (s.e. 0.28), implying $\delta \approx 0.41$; on the Legal-Amazon frontier subsample, where the land margin is active, it is 1.66 (s.e. 0.43), implying $\delta \approx 0.62$ (Table 3). Appendix D.1 details the construction, a two-level robustness battery (start years, exposure weighting—including geography-predicted (Fiszbein) exposure shares—leave-one-out shocks, the depletion control, and moving-average prices, which raise δ to 0.70–0.75), and an honest discussion of inference: the design rests on few commodity shocks, so we lean on the share-exogeneity reading for the point estimate and do not rest any headline claim on its reduced-form precision (Appendix D.1.4).

For the main specification, δ is set to 0.41, the pooled all-Brazil estimate with the depletion control (Table 3), corresponding to a land-supply elasticity of about 0.7—comparable to the calibration for global deforestation in Farrokhi et al. (2025) (0.9) and the crop-price elasticity of pan-tropical deforestation in Berman et al. (2023). No single reduced-form design pins δ tightly: the frontier (Legal Amazon) shift-share gives $\delta = 0.62$ with the depletion control, dynamic-panel estimates on the frontier are as low as 0.07, and persistent-price specifications push δ up to 0.75 (Appendix D.1). Because leakage is increasing in δ , we re-solve the model across this range and report headline leakage as a range rather than a point: 3.7–8.8% for the Priority List and 6.4–14.7% for Protected Areas between the all-Brazil baseline and the frontier estimate (Section 7, Appendix Table 20). We retain $\delta = 0.41$ as the baseline because the model spans all of Brazil and generates regional differences in the land margin endogenously through forest availability rather than through a region-specific δ ; the frontier estimate marks the upper end of the identified range, and the small-leakage conclusion—and the ranking between the two policies—survives everywhere within it.

Table 3: Deforestation elasticity δ : shift-share IV (baseline)

Dependent Variables:	log deforestation T^D	log land rent v	log deforestation T^D	
	OLS	First stage	Reduced form	IV
Model:	(1)	(2)	(3)	(4)
<i>Variables</i>				
log land rent v	0.0990*** (0.0199)			0.6817** (0.2805)
log remaining forest $_{t-1}$	0.7388*** (0.1053)	0.1418*** (0.0405)	0.7610*** (0.1075)	0.6644*** (0.1070)
Bartik price index Z		0.3452*** (0.0284)	0.2353** (0.0977)	
<i>Fixed-effects</i>				
Municipality	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	56,898	56,898	56,898	56,898
R ²	0.93065	0.94821	0.93060	0.92738
Within R ²	0.00635	0.01121	0.00565	-0.04042

Clustered (Municipality) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: All Brazilian municipalities, 2008–2018. Implied $\delta = \hat{\beta}_{IV}/(1 + \hat{\beta}_{IV}) = 0.41$; Legal-Amazon subsample: $\hat{\beta}_{IV} = 1.66$ (s.e. 0.43), $\delta = 0.62$. Outcome is gross natural-vegetation loss (Map-Biomass transitions). $\log v$ is the model land-rent index, instrumented by the base-year-exposure $\times$ commodity-price Bartik index; $\log \text{forest}_{t-1}$ is the lagged remaining stock of natural vegetation, a predetermined depletion control (Appendix D.1). First-stage F is unit-clustered; it is the appropriate statistic under the share-exogeneity reading of the design, but the identifying variation rests on few commodity shocks—exposure-robust inference is much weaker (Appendix D.1.4). Standard errors clustered by municipality. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$.

Remaining parameters The parameter ψ , governing how deforestation productivity falls as forest is depleted, is estimated from the steady-state relationship between forest-to-agriculture ratios and forest area (Appendix D.2, Table 13); our point estimate is $\psi = 0.26$ (the low-deforestation steady-state column, identified jointly with $\delta = 0.41$). Because ψ and δ enter this steady-state relationship jointly, we identify ψ from micro-regions whose deforestation has fallen below the regrowth rate—those plausibly at their forest steady state—which is what separates the two (Appendix D.2). The regrowth rate $\rho = 1\%$ is the national average of MapBiomass secondary-forest growth over anthropic land, 2002–2019, applied uniformly (the micro-region mean is 1.5%, the median 0.8%). Land shares α_k come from the 2006 agricultural census; the occupational-choice parameters from Alvarez (2020); the Engel elasticity η from the 2017–18 household expenditure survey; migration costs and ϵ from migration gravity on the 2010 census; and trade costs from a gravity equation on inter-state trade flows (Appendix D).


Model inversion Figure 6 maps the estimated deforestation productivities Z_{rt}^D ; the corresponding maps of agricultural TFPs, non-agricultural TFPs, and amenities are in Appendix D.8. Given the calibrated parameters, the fundamentals exactly rationalize the observed data year by year. The inverted fundamentals include the region–sector amenity $G_{NA,r}$ and the (now data-disciplined) non-agricultural productivity Z_r^{NA} : the observed agricultural employment share pins $G_{NA,r}$, while the observed non-agricultural value added (its labor portion $\gamma_{NA} va_r^{NA}$ per worker) pins Z_r^{NA} , so the model reproduces the spatial distribution of non-agricultural income. Appendix C gives the inversion system.

5.5 The sources of variation in deforestation

With the fundamentals in hand, we can now make precise the stylized facts of Section 4.5 and answer the question the theory poses: is spatial variation in deforestation supply- or demand-driven?

Correlates of deforestation Table 4 regresses observed deforestation on the (log) inverted fundamentals, under year and micro-region fixed effects throughout. The supply-side fundamentals alone attain a pseudo- R^2 of 0.955 (column 3); the full de-

Figure 6: Estimated Deforestation TFPs



Note: These maps illustrate the spatial distribution of the total deforestation TFP $Z_{rt}^D = \bar{Z}_{rt}^D (T_{rt}^F)^\psi$ across Brazil for the years 2003, 2010, and 2019, where ψ governs the dependence of the TFP on unprotected forest area. The color scale is consistent across the three years.

mand side—sectoral TFPs and market access—reaches only 0.853 (column 2), barely above what the demand TFPs manage without market access (column 1). With everything included (column 4), the coefficient on $\ln \bar{Z}_r^D$ is 1.67—almost exactly the model-implied elasticity $1/(1 - \delta) = 1.69$ at the calibrated $\delta = 0.41$ in equation (2)—while the demand-side fundamentals are an order of magnitude smaller, with pasture productivity, the land-intensive margin, the only sizable one and positive as the model predicts. The individual demand-side coefficients should not be over-read: the inverted fundamentals are strongly correlated with one another and with $\bar{Z}_r^D$, so several change sign as conditioning sets change across columns (e.g. the crop TFPs and market access between columns (2) and (4)). The stable and interpretable objects are the fit comparison across blocks—the supply side alone recovers almost all of the variation, the full demand side does not—and the coefficient on $\ln \bar{Z}_r^D$, which is pinned by the model’s supply elasticity.

5.6 The empirical content of Z^D

The deforestation productivity is recovered as a residual from the inversion, which invites the concern that it is a free-floating fudge factor rather than a genuine regional supply fundamental: if it merely relabels a mobile input or unmodeled demand, the supply-driven reading of leakage would be undercut. Table 5 addresses this by regressing the inverted productivity on characteristics that are *external* to

Table 4: Correlates of observed deforestation

	Demand (no MA) (1)	Demand (all) (2)	Supply only (3)	All (4)
<i>Demand-driven</i>				
ln Temp. crops TFP (Z_{temp}^A)	-0.258 (0.266)	-0.119 (0.285)		-0.030*** (0.010)
ln Perm. crops TFP (Z_{perm}^A)	1.638*** (0.365)	1.484*** (0.344)		-0.137*** (0.023)
ln Pasture TFP (Z_{past}^A)	-1.422*** (0.420)	-1.391*** (0.407)		0.772*** (0.022)
ln Non-agric. TFP (Z^{NA})	-0.097 (0.090)	-0.098 (0.089)		-0.006 (0.004)
ln Market access		-17.943*** (5.498)		-0.169 (0.281)
<i>Supply-driven</i>				
ln Deforestation TFP ($\bar{Z}^D$)			1.462*** (0.037)	1.671*** (0.002)
ln Natural area			-0.118 (0.240)	0.586*** (0.016)
Year FE	Yes	Yes	Yes	Yes
Region FE	Yes	Yes	Yes	Yes
Observations	7,024	7,024	7,024	7,024
Pseudo R^2	0.852	0.853	0.955	0.967

Note: PPML (Poisson) regressions at the micro-region-year level; the outcome is the observed deforestation level and the regressors are model-inverted regional fundamentals (in logs). The sample is the deforestation frontier—micro-regions with a positive deforestation TFP ($\bar{Z}^D > 0$); regions with zero deforestation (and hence $\bar{Z}^D = 0$) have no frontier and are excluded. All columns include year and micro-region fixed effects. Column (1) includes the demand-side fundamentals (crop and non-agricultural TFPs) without market access; column (2) adds market access; column (3) includes the supply-side fundamentals only ($\bar{Z}^D$ and natural area); column (4) includes all fundamentals together. Standard errors clustered by micro-region in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

the inversion—climate, biomass, and geography, none of which is used to recover it. The outcome is the forest-invariant base fundamental $\overline{Z_{rt}^D}$; the effective level $Z_{rt}^D = \overline{Z_{rt}^D}(T_{rt}^F)^\psi$ that enters the land-supply curve is 0.98 correlated with it and gives the same picture. Biophysical conditions during the clearing season—the hottest-month temperature and the driest-month precipitation, aggregated to micro-regions from a ~ 10 km grid—together with annual rainfall, soil moisture, and above-ground carbon density account for about a third of the cross-regional variation in the fundamental (column 1), with the principal signs those of a clearing-cost interpretation: clearing productivity is higher where the dry season is *hotter* and *drier*—the conditions under which forest is cheapest to fell and burn—and lower where carbon density is high, denser forest being more costly to fell. The point is that roughly a third of a model residual is recovered by temperature, rainfall, and biomass, variables the model never sees, and in economically sensible directions. This is what we would expect if the object measures how costly it is to convert local forest to agricultural land, and is hard to reconcile with its standing in for a spatially mobile input.

Enforcement. Command-and-control should lower *effective* clearing productivity, so Z^D ought to fall with enforcement intensity. We measure enforcement with IBAMA flora-infraction notices (*autos de infração*). The raw count of fines is *positively* correlated with Z^D —unsurprisingly, since enforcement is deployed where deforestation, and hence Z^D , is highest. Normalizing into a detection rate (fines per unit *actually deforested*) is the natural deterrence object—the share of clearing that gets caught—but that ratio shares the same-period deforestation T_{rt}^D with Z_{rt}^D itself, so we use only the *one-year lag* of the detection rate as the regressor, breaking the same-period mechanical link. The lagged fine rate is negative and significant even here (-0.06 on clearing productivity, columns 1–2 of Table 17; -0.10 on clearing itself, column 3): within a region, a year of more intense enforcement per unit of clearing is followed by lower clearing productivity the next year—deterrence-consistent co-movement that survives region fixed effects, read with the detection-bias caveat above. Remaining forest enters positive, consistent with depletion lowering Z^D over time.

Table 5: The empirical content of the deforestation productivity $\overline{Z^D}$

	Outcome: (log) base deforestation productivity $\ln \overline{Z_{rt}^D}$			
	Year FE		Two-way FE	
	(1)	(2)	(3)	(4)
ln Precip. (annual)	-0.207** [0.095]	-0.122 [0.087]	-0.102** [0.041]	-0.097** [0.041]
ln Precip. (driest mo.)	-0.054*** [0.015]	-0.037** [0.015]	0.004 [0.006]	0.005 [0.006]
Temp. (mean)	-0.045* [0.024]	-0.012 [0.022]	0.070*** [0.021]	0.033* [0.020]
Temp. (hottest mo.)	0.226*** [0.024]	0.191*** [0.022]	-0.043*** [0.009]	-0.039*** [0.008]
ln Soil moisture	0.133*** [0.040]	0.147*** [0.037]	0.062** [0.029]	0.073** [0.030]
ln Carbon density	-0.103** [0.044]	-0.139*** [0.044]		
Fines / deforestation, $t - 1$		-0.277*** [0.032]		-0.071*** [0.019]
Year fixed effects	Yes	Yes	Yes	Yes
Micro-region fixed effects	No	No	Yes	Yes
R^2	0.361	0.405	0.810	0.843
Within R^2	0.352	0.401	0.005	0.011
Observations	7024	6156	7024	6156

Note: OLS at the micro-region-year level, 2002–2019, SE clustered by micro-region. The outcome is the (log) base deforestation productivity $\overline{Z_{rt}^D}$; all regressors are external to the inversion. Columns (1)–(2): year fixed effects; columns (3)–(4): year and micro-region fixed effects (time-invariant covariates drop out; within- R^2 is net of the fixed effects). Columns (2) and (4) add the IBAMA *Flora* detection rate (fines per unit deforested area, asinh-transformed), lagged one year only—the contemporaneous rate shares the same-period deforestation with Z^D itself and so is not reported; temperatures in °C, carbon density per unit forest area. Because fine counts co-move positively with clearing activity, a negative lagged coefficient is a lower bound on deterrence, not an overstatement of it. Fines measure infractions $\times$ detection probability—descriptive co-movement, not a causal estimate. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

6 Counterfactual Results

We now use the calibrated model to simulate counterfactual scenarios of local deforestation policy and measure leakage for the two policies of Section 3. All counterfactuals are computed in the open economy—the headline specification, with the rest of the world bidding for Brazilian output at the world market; Section 7 shows that closing the economy leaves cumulative leakage essentially unchanged.

Table 6 previews the headline results over 2003–2018. The Priority List avoided 80,051 km² of deforestation gross of general-equilibrium effects; the general-equilibrium response added back 2,935 km², so net avoided deforestation is 77,116 km² and leakage is 3.7%—the policy retains about 96% of its local gains in the aggregate. Protected Areas established from 2003 avoided 7,311 km² gross, with 467 km² leaking back: net avoided deforestation of 6,844 km², a retention of about 94% and leakage of 6.4%. In carbon terms, the two policies avoided a net 4.1 and 0.34 Gt of CO₂ respectively, and carbon leakage is *below* the area rates—3.0% and 5.5%—because the targeted regions are among the most carbon-dense. The Priority List leaks proportionally less despite being the far larger intervention—the pattern the theory of Section 4 predicts, since the Priority List targets the high- Z^D deforestation frontier and Protected Areas do not. One caveat applies to every number in this section and we state it up front: all figures are computed at the baseline calibration of the deforestation supply elasticity, $\delta = 0.41$, the all-Brazil estimate. Because leakage is increasing in δ and δ is not tightly identified, Section 7 re-solves the model across the identified range; cumulative leakage runs from 3.7% to 8.8% for the Priority List and from 6.4% to 14.7% for Protected Areas, with the ranking between the two policies invariant throughout.

6.1 Defining counterfactuals and leakage

For each of the two policies, we simulate two counterfactual scenarios for the evolution of nation-wide deforestation between the years 2003 and 2018.

Counterfactual A: No policy—a scenario in which, between 2003 and 2018, the policy was never enacted.

Counterfactual B: Policy, no leakage—a scenario in which, between 2003 and 2018, the policy was implemented, but it has no general equilibrium effects.

Table 6: Headline results: avoided deforestation and leakage, 2003–2018

Policy	Gross avoided (km ²)	GE leakage (km ²)	Net avoided (km ²)	Leakage (%)
Priority List	80,051	2,935	77,116	3.7
Protected Areas	7,311	467	6,844	6.4

Note: Cumulative results from the open-economy counterfactual simulations, 2003–2018. “Gross avoided” is the reduction in cumulative deforestation from the policy holding prices at their no-policy level (counterfactual A minus counterfactual B, defined below); “GE leakage” is the additional deforestation induced by general-equilibrium price responses (data minus counterfactual B); “Net avoided” is their difference; “Leakage” is GE leakage as a share of gross avoided deforestation, equation (13).

The difference between counterfactual A and the observed evolution of deforestation measures the overall effect of the policy, including both the direct effect via banned deforestation in targeted areas, and the indirect effect via potential leakage to non-targeted areas. The difference between the observed evolution of deforestation and counterfactual B measures the amount of leakage caused by the localized policies. The difference between counterfactuals A and B measures the total effect of the localized policies under no deforestation leakage, i.e. only considering the direct effect via banned deforestation in targeted areas.

For each of the two policies, we construct the counterfactuals as follows. First, we calibrate the baseline regional fundamentals $\{Z_{rt}^s, B_{rt}\}$ to fit the observed deforestation between the years of 2003 and 2018. We construct $(Z_{rt}^D)^{nopolicy}$, the deforestation productivity of each region in a counterfactual scenario where the policy was not implemented, as described in subsections 6.1 and 6.1 below.

For counterfactual A, we simulate total land cleared for agriculture in each year between 2003 and 2018, solving for the optimal deforestation level $(T_{rt}^D)^{nopolicy}$ according to Equation 6 using the calculated regional deforestation productivity $(Z_{rt}^D)^{nopolicy}$. Total agricultural land at each period is then given by equation 11.

$$(T_{rt+1}^A)^{nopolicy} = (1 - \rho)(T_{rt}^A)^{nopolicy} + (T_{rt}^D)^{nopolicy} \quad (11)$$

For counterfactual B, we calculate a “no leakage” scenario that has the implied prices coming from counterfactual A (i.e. no-policy prices), and baseline region de-

forestation productivities estimated through the model Z_r^D (with policy). Fixing the prices under the no-policy scenario effectively shuts down the channel through which deforestation could leak to other areas, as both labor market and goods market relocation mechanisms operate through price channels.

$$(T_r^D)^{leakage} = \left(\delta \frac{(q_r)^{nopolicy}}{(p_r)^{nopolicy}} \right)^{\frac{\delta}{1-\delta}} (Z_r^D)^{\frac{1}{1-\delta}} \quad (12)$$

Given the counterfactual levels of deforestation for counterfactuals A and B, leakage at horizon t is the share of avoided deforestation undone by general-equilibrium effects:

$$\text{Leakage}_t = \sum_r \left(\underbrace{(T_{rt}^D)^{policy}}_{\text{Data}} - \underbrace{(T_{rt}^D)^{leakage}}_{\text{Counterfactual B}} \right) / \left(\underbrace{(T_{rt}^D)^{nopolicy}}_{\text{Counterfactual A}} - \underbrace{(T_{rt}^D)^{leakage}}_{\text{Counterfactual B}} \right) \quad (13)$$

We also convert deforestation from surface area to carbon emissions using data on the heterogeneous carbon density of the natural ecosystems of the various regions.

Two clarifications on this construction. Counterfactual B is an accounting device, not an equilibrium: it prices the policy’s direct effect at the *no-policy* price system, so the decomposition is exact by construction—the total policy effect splits identically into a fixed-price direct effect (counterfactual A minus B) and a general-equilibrium term (data minus B)—but the size of each piece inherits the fixed-price convention, in the spirit of a Laspeyres decomposition. The alternative Paasche convention—pricing the direct effect at with-policy prices—would shift mass between the two pieces without changing the total policy effect; we adopt the Laspeyres convention throughout and note that the leakage *share* inherits it. Leakage as defined in (13) is therefore precisely the component of the policy’s effect that operates through equilibrium price adjustment—wherever it lands, including inside the targeted regions themselves.

No Priority List deforestation productivities

To estimate the counterfactual deforestation productivities of regions in the Priority List,²⁰ we first estimate the effect of the policy on the values of Z_{rt}^D obtained from

²⁰The Priority List is defined at the municipality level, while the model’s regions are micro-regions; we treat a micro-region as listed if any of its municipalities is (69 municipalities across 31 of the 557 micro-regions). The approximation errs only by over-inclusion: over 2008–2018, actually-

the model inversion, with a Poisson quasi-maximum likelihood model with region and time fixed effects,

$$\mathbb{E}(\bar{Z}_{rt}^D | r, t, \text{Priority}_{rt}) = \exp(\alpha_r + \delta_t + \beta \text{Priority}_{rt})$$

Having estimated β , we construct a no-priority-list $(\bar{Z}_{rt}^D)^{nopolicy}$ that equals

$$(\bar{Z}_{rt}^D)^{nopolicy} = \bar{Z}_{rt}^D / \exp(\hat{\beta}). \quad (14)$$

Using this new productivity of deforestation, we simulate the model forward from 2003. The overall deforestation productivities $(Z_{rt}^D)^{nopolicy}$ differ both because of the changes in $\bar{Z}_{rt}^D$ in Priority List regions that are external to the simulation and because of the endogenous changes in T_{rt}^F that happen in all regions:

$$(Z_{rt}^D)^{nopolicy} = (\bar{Z}_{rt}^D)^{nopolicy} ((T_{rt}^F)^{nopolicy})^\psi \quad (15)$$

Appendix Figure 28 shows these changes in deforestation productivities.

No Protected Areas deforestation productivities

To estimate the counterfactual deforestation productivities as if there were no new protected areas established from 2003, we assume that the part of deforestation productivity that changes is not $\bar{Z}_{rt}^D$ but T_{rt}^F , the area of r under unprotected forest. Given the extremely low rates of deforestation inside protected areas, we assume perfect enforcement. Let T_{rt}^P be the area of region r that gets protected in year t . To simulate the reversal of new Protected Areas, we increase the area available for deforestation by T_{rt}^P , so that

$$(Z_{rt}^D)^{nopolicy} = \bar{Z}_{rt}^D ((T_{rt}^F)^{nopolicy} + T_{rt}^P)^\psi, \quad (16)$$

where $(T_{rt}^F)^{nopolicy}$ is the remaining area of unprotected forest resulting from the simulation of counterfactuals with no new protected areas in previous periods. Appendix


listed municipalities account for 70% of the flagged micro-regions' observed deforestation (never-listed neighbors for the remaining 30%, about 8% of national deforestation) and 63% of their 2008 natural area. Because the policy shock $\hat{\beta}$ below is estimated on the same micro-region aggregates, the treatment definition and the estimated shock are mutually consistent: dilution from unlisted neighbors attenuates $\hat{\beta}$ rather than compounding the error.

Figure 29 shows the resulting changes in deforestation productivities.

6.2 Priority List

For the Priority List, panel (a) of Figure 7 plots observed deforestation against counterfactuals A (no policy) and B (policy without GE effects).


Figure 7: Counterfactual deforestation trends



Note: This figure shows the trends in yearly deforestation levels (in thousands of square kilometers) across Brazil between 2003 and 2018 under counterfactual simulations based on the Priority List (panel a) and on the establishment of Protected Areas (panel b). In each panel, the blue curve represents observed deforestation, the red curve the counterfactual scenario in which the policy was never implemented (counterfactual A), and the green curve the counterfactual in which the policy was implemented but GE effects, and hence leakage, are shut down by fixing prices at the no-policy scenario (counterfactual B).

Panel (a) of Figure 8 shows the evolution of cumulative leakage as defined in equation 13. Leakage at the start of the period is below 1%, and rises gradually to 3.7% of avoided deforestation by 2018—the endogenous stock-flow dynamics anticipated by Remark 1, to which we return in Section 6.5. The headline number must be read with its geography. About 60% of these general-equilibrium effects occur *within* Priority List regions: because the Priority List does not fully ban deforestation, part of the price response falls on the listed municipalities themselves—in-situ effects that dampen local gains but do not move clearing anywhere else, and do not bias reduced-form estimates. Leakage *outside* the Priority List—relocation in the strict sense, clearing displaced to untargeted regions—is the remaining 40%: about 1.5% of avoided deforestation, some 1,200 km² over fifteen years. In carbon terms

Figure 8: Cumulative leakage: Priority List vs Protected Areas



Note: This figure shows the evolution of cumulative leakage over time for the Priority List (panel a) and for the establishment of Protected Areas (panel b), following formula (13). Within each panel, the left plot shows leakage in terms of deforested area and the right plot in terms of carbon emissions, calculated using the micro-region-level average carbon density (above- and below-ground) of natural ecosystems in 2010. In panel (a), the blue line indicates the total % of (cumulative) avoided deforestation that is undone by increases due to GE effects (“leakage”); because the Priority List does not fully ban deforestation, some GE effects are experienced within Priority-List regions (see Appendix Figure 33), and the green line shows the fraction of leakage due to such in-situ effects, while the red line shows the fraction that is “true leakage” to non-Priority-List regions. For Protected Areas the in-situ distinction is not meaningful, so panel (b) reports total leakage only.

leakage is smaller still—3.0% for the Priority List and 5.5% for Protected Areas—since Priority List regions exhibit some of the highest carbon densities; the regions with even higher densities, in the northern Amazon basin, absorb little displaced clearing because their market access and agricultural productivities are low.

6.3 Protected Areas

Panel (b) of Figure 7 shows the corresponding trends for Protected Areas, and panel (b) of Figure 8 the cumulative leakage. In this case there is no meaningful distinction between GE effects in-situ and displacement, because we assume Protected Areas lead to zero deforestation within them. Although the total deforestation reduction caused by the new Protected Areas is an order of magnitude smaller than that of the Priority List, their leakage is *higher*: 6.4% of avoided deforestation by 2018. This is because Protected Areas do not target the regions with the highest fundamental productivity of deforestation—those regions remain available to absorb the increased

demand for land. In the language of Proposition 1, Protected Areas remove a small share b of the national supply of new land, and leave the cheap suppliers standing; the Priority List removes the cheapest suppliers themselves.

6.4 Where leakage lands

The supply-driven mechanism makes a testable spatial prediction: displaced clearing should flow down the residual supply curve, to the untargeted regions where clearing is cheapest—high Z_r^D —rather than to well-connected or agriculturally productive regions per se. The model output confirms this sharply. Cumulative GE effects on deforestation from the Priority List total 2,935 km² through 2018, of which 60% occur within Priority-List regions and 40% in the 526 untargeted micro-regions (mapped in Appendix Figure 33). Among receivers, absorption is highly concentrated—the top decile of untargeted regions absorbs 45% of displaced clearing—and Table 7 shows what predicts it.

Table 7: Where leakage lands: absorbed GE effects and pre-policy fundamentals

Outcome: log cumulative absorbed leakage, non-PL regions			
	(1)	(2)	(3)
(log) Deforestation TFP Z_r^D (2003)	1.07*** [0.07]	1.06*** [0.09]	1.31*** [0.12]
(log) Pasture TFP (2003)			0.50*** [0.09]
(log) Market access (2003)		-0.06 [0.22]	-0.51*** [0.18]
R^2	0.340	0.340	0.405
Observations	446	446	446

Note: OLS, cross-section of untargeted (non-Priority-List) micro-regions with positive absorbed leakage ($N = 446$ of 526). The outcome is the log of the region’s cumulative GE effect on deforestation, 2003–2018 (data minus counterfactual B, in km²), computed from the open-economy engine output. Regressors are pre-policy (2002/03) model-estimated fundamentals. Heteroskedasticity-robust standard errors in brackets. ***, **, *: significance at 1%, 5%, 10%.

The receiving regions’ deforestation productivity alone explains 34% of the spatial distribution of absorbed leakage, with an elasticity near one (column 1). Because the outcome is in logs, the sample conditions on the 446 of 526 untargeted regions with

positive absorbed leakage; the estimates therefore describe the intensity of absorption among receivers, not selection into receiving. Conditional on Z^D and pasture productivity, market access enters *negatively* (column 3): displaced clearing does not flow to well-connected regions—it flows to the next-cheapest clearers, which tend to be remote. A demand-driven model would instead send displaced production toward high-market-access, high-land-value regions. This prediction is also the natural point of contact with reduced-form neighbor analyses of leakage: proximity matters only insofar as it proxies for cheap clearing.

Overall, even though there is a detectable amount of leakage from localized anti-deforestation policies, both policies were not only effective in reducing deforestation locally but retained about 96% (Priority List) and 94% (Protected Areas) of that effect at the national level over 15 years at the baseline calibration—and no less than about 91% and 85% anywhere in the identified range of the supply elasticity (Section 7). Concerns that leakage undoes the majority of the gains from localized anti-deforestation policy are, in this setting, unwarranted—and the reason is the supply-driven nature of Brazilian deforestation together with where the policies aimed.

The comparison between the Priority List and Protected Areas is one contrast between two implemented policies. We now generalize it along the three margins a policymaker chooses: *whom* to target, *how hard* to enforce, and *over what horizon* to evaluate.

6.5 Targeting

By Proposition 1, leakage per unit of avoided deforestation falls with the targeted share of national deforestation b , and by Proposition 2 it falls with the supply-driven share of the targeted regions' deforestation. Both point the same way in Brazil: because deforestation is supply-driven (Section 4.5), targeting the biggest deforesters is targeting the cheapest clearers, and both propositions push leakage down. The Priority List—which selected municipalities precisely on observed deforestation—removes the cheapest suppliers of new land ($b = 0.37$ of national pre-policy deforestation among listed regions); Protected Areas—sited largely on conservation and political criteria in high-forest, low-deforestation areas—remove expensive or inframarginal suppliers and leave the frontier standing. The 3.7% vs 6.4% contrast above is thus not an idiosyncrasy of the two programs but two points on a general targeting curve.

To trace out that curve, we compute leakage when a policy of identical scope—the same number of regions, the same estimated supply shock, the same start date—is aimed by alternative targeting rules: the actual Priority List; highest observed deforestation; highest deforestation productivity Z_r^D ; and highest land value v_r (demand-driven targeting).²¹ The results line up exactly as the theory predicts (Figure 9). Supply-side targeting both delivers conservation at scale and minimizes leakage: the actual Priority List and the highest-deforestation rule both leak 3.7%, and targeting the highest- Z^D regions directly leaks least of all (3.3%). Demand-driven targeting is in a different regime altogether: aiming the same policy at the regions with the most valuable land produces leakage of 36%—an order of magnitude more—with almost all of it displaced outside the targeted set, while avoiding almost no deforestation to begin with (732 km² net, against 79,635 for the actual policy in the same exercise). The sufficient statistic of Corollary 1 organizes the whole pattern: the land-value rule selects regions whose share of the market *value* of clearing is eleven times their share of cleared *area*, precisely the configuration that maximizes leakage.


6.6 Policy strength

The Priority List is a partial ban: the PPML estimate above implies the policy cut the targeted regions’ deforestation supply by the fraction $f = 1 - \exp(\hat{\beta})^{1/(1-\delta)} = 0.48$ at fixed prices, so the policy removes $bf = 0.18$ of the national supply of new land. Proposition 1 implies that stronger enforcement (higher f , hence a higher effective banned share) *lowers* leakage per unit of avoided deforestation while raising gross avoided deforestation—strength and targeting are complements, not substitutes. We quantify this by varying the treatment intensity f for the Priority List and tracing leakage, gross avoided, and net avoided deforestation (Appendix Figure 34): leakage per unit falls and gross avoided rises as enforcement strengthens.

The fixed-price supply cut $f = 0.48$ also reconciles the model with the reduced form. The synthetic difference-in-differences of Section 3 estimates a 36% *equilibrium* reduction in listed municipalities’ deforestation—smaller than f , and it should be: the reduced form measures the policy’s effect after general-equilibrium adjustment,

²¹A rule replicating the Protected-Areas siting is not part of this exercise: Protected Areas operate through the removal of convertible land rather than through the supply shock the exercise holds fixed, so the two are not comparable at identical scope. The Protected-Areas point on the curve is represented by the actual policy.

Figure 9: Conservation delivered versus leaked, by targeting rule



Note: Cumulative 2003–2018 deforestation outcomes (open economy) when policies of identical scope—the Priority List’s region count and estimated supply shock, from 2008—are aimed by different targeting rules. The downward bar is gross avoided clearing in the targeted regions; the upward bars are the general-equilibrium *increases* in clearing, in targeted and in non-targeted regions (net national avoided = gross $\times (1 - L)$). Filled diamonds (right axis) give the total leakage rate L ; hollow diamonds its within-targeted-regions component only. The grey label beneath each rule is the Corollary 1 sufficient statistic—the targeted regions’ share of the market *value* of clearing divided by their share of *area* cleared; leakage rises monotonically in it. Demand-side targeting (highest land value) removes almost no clearing in the first place and relocates most of what it removes. The actual-Priority-List column re-imposes the estimated supply shock on the reconstructed no-policy world, so its totals differ slightly from the calibrated policy counterfactual of Table 6 (79,635 vs. 77,116 km² net).

which claws back part of the fixed-price reduction inside the treated municipalities themselves (the in-situ component of leakage above). The two numbers are not competing estimates of one object but the fixed-price and equilibrium ends of the same shock, and their ordering—equilibrium response below fixed-price response—is itself a qualitative confirmation of general-equilibrium dampening.

6.7 Horizon

Leakage is not a single number but a path: 0.7% in the first policy year, rising to 3.7% by 2018 for the Priority List. Remark 1 explains why: the demand relevant

for leakage is the demand for *newly cleared* land—a flow—while farmers demand the *stock* of agricultural land, and Brazil’s stock-flow ratio is roughly 53. At short horizons, removing supply mostly re-times the adjustment of an almost-fixed stock, so the effective demand elasticity is enormous and leakage negligible; as the horizon lengthens, the elasticity falls and leakage rises toward its long-run level.

The design lesson for policy is the paper’s bottom line: *target the economic margin that generates deforestation*. Where deforestation is supply-driven—as in Brazil—targeting the observed deforestation frontier removes the cheapest suppliers of new land and minimizes leakage; policies sited away from that frontier leak more per unit of conservation; and evaluations must specify their horizon, because leakage compounds as the land market adjusts. Two cautions apply to all three margins. A leakage *rate* is not a welfare statement: a policy can leak little simply by avoiding almost nothing, so rates must be read jointly with the conservation delivered—which is why our exhibits report net avoided deforestation alongside L . And low leakage is not optimality: L prices the general-equilibrium tax on a policy’s gross delivery, not the policy’s benefits against its costs, which we take up in Section 7.

7 Mechanisms, Robustness, and Welfare

This section verifies that the supply-driven mechanism—not an auxiliary modeling choice—drives the headline results, and reports what the policies did to the rest of the economy.

7.1 What drives the result

The theory converts the paper’s findings from numbers into economics, and identifies exactly what the model would need for leakage to be large: demand-driven deforestation (conservation targeting productive, well-connected, expensive land), high substitutability across origins, and inelastic aggregate land demand (homothetic preferences, a single agricultural good, no factor substitution). Brazil’s data reject each in turn: the fundamentals correlates of Section 4.5 reject demand-driven variation; the reduced-form intensification of Section 3.1 rejects inelastic demand; and the where-leakage-lands result of Section 6.4 confirms that displaced clearing follows clearing costs, not market access.

Two structural exercises sharpen this account. First, a *swap counterfactual*: re-solving the calibrated model in constructed economies that reproduce Brazil’s baseline deforestation vector with its source reassigned. The supply-driven reassignment (all variation onto Z^D) calibrates and behaves like the actual economy. The demand-driven reassignment delivers a stronger result than the exercise was designed for: with Z^D homogenized, *no* configuration of agricultural productivities can reproduce the observed dispersion of deforestation—the required land-rent dispersion exceeds what the model’s trade and migration structure permits. Brazil’s deforestation pattern is not merely better explained by the supply side; it is *infeasible* for the demand side.²²

Second, *channel shutdowns*: re-solving the Priority-List counterfactual with one structural ingredient changed at a time (the parameter sweep of Table 20). Eliminating internal migration entirely leaves cumulative leakage essentially unchanged (3.7% vs 3.7%): the labor-market response to conservation is absorbed by the *occupational* margin, with workers displaced from agriculture switching sectors locally rather than relocating spatially. Conversely, making occupational switching nearly impossible ($\chi = 1.05$) raises leakage to 3.9%, while more elastic switching lowers it (3.6% at $\chi = 2$)—the sectoral labor-supply elasticity is one of the demand margins in ε_D .²³

7.2 Direct tests of the mediating channels

The mechanism also leaves a reduced-form footprint in observed data. For each micro-region we build a market-access exposure to conservation—the avoided clearing in the regions it is linked to through crop-and-market overlap, $\text{Exposure}_i = \sum_{j \neq i} m_{ij} \text{Cons}_j$, a Donaldson–Hornbeck object constructed from observed trade costs, crop revenues and regional incomes (Appendix G)—alongside each region’s own avoided clearing Cons_i . Two facts follow on land values. Being Priority-Listed raises the observed, revenue-based land-rent proxy by about 25%, and the effect scales with own avoided


²²In stylized versions of the model where both sources are feasible, demand-driven leakage exceeds supply-driven leakage by a factor of three or more and can exceed 100%, so demand-driven conservation can raise total deforested area.

²³Each variant is a full re-inversion and re-resolution of the model (Table 20). The world-demand margin is assessed directly by the open-versus-closed comparison below, which shuts the foreign block off entirely. Shutdowns of non-homothetic preferences and of multiple agricultural goods change the demand system the inversion conditions on—they define a different model rather than a parameter deviation—and are therefore outside the scope of the sweep; the sweep instead spans the parameters governing those margins (η enters through α , and the across-good substitution through θ).

clearing (significant at 1%); since the proxy is one the calibration never targets, this is direct, non-circular evidence that conservation makes farmland scarcer and dearer. The demand exposure enters positively but small and imprecise in rents—exactly as the model implies, because when goods are close substitutes ($\sigma = 9.6$) a supply cut is absorbed with little price movement, so the incidence on *rents* is muted.

The demand channel instead reveals itself in *quantities*. Across untargeted regions the exposure is a strong, robust predictor of *where* the model’s displaced clearing lands, surviving a control for deforestation productivity Z^D (which itself enters with near-unit elasticity): leakage concentrates where regions are both demand-linked to the conserved areas and able to bring land into use cheaply, and the supply-weighted exposure (exposure $\times$ baseline clearing) tracks the leakage map far more tightly than raw demand exposure (Figure 10). Consistent with intensification rather than extensive expansion, exposed regions shift the composition of land toward crops rather than pasture. Appendix G details the exposure construction and reports all regressions.

Figure 10: Where leakage lands, demand exposure, and supply-adjusted exposure



Note: Micro-regions shaded by percentile rank (over non-Priority-List regions) of: (1) cumulative absorbed leakage $L_i = \sum_t \max(T_r^D - T_{r,ap}^D, 0)$; (2) the market-access demand exposure $\sum_{j \neq i} m_{ij} \text{Cons}_j$; (3) the same exposure $\times$ baseline (no-policy) deforestation. Priority-List regions in grey. Details in Appendix G.

7.3 Robustness

Open vs closed economy Because Brazil’s deforestation-linked commodities are largely exported, a natural concern is that leakage computed in a closed economy understates relocation through world markets. It does not: computing all coun-

terfactuals with the rest-of-world block switched off leaves the small-leakage finding intact—cumulative leakage is 3.7% open versus 3.2% closed for the Priority List, and 6.4% versus 5.7% for Protected Areas (Appendix E, Table 18). The reason is instructive: world demand raises ε_D (an additional substitution margin, pushing leakage down) but also makes demand for Brazilian land less satiable (pushing leakage up); at the calibrated ζ the two effects nearly cancel. The small-leakage finding is therefore not an artifact of treating demand as domestic.

The deforestation elasticity δ Leakage is increasing in $\varepsilon_S = \delta/(1 - \delta)$ (Proposition 1), and δ is the one parameter whose uncertainty moves the headline numbers materially, so we are explicit about both the choice of baseline and what the identified range implies. The model is calibrated at $\delta = 0.41$, the pooled all-Brazil estimate with the depletion control ($\varepsilon_S \approx 0.7$; Table 3). We use the all-Brazil estimate as the headline because the model spans all of Brazil: in the model, regional differences in the activity of the land margin arise endogenously through forest availability (the cap $T^D \leq T^{\max}$ and the channel $Z^D \propto (T^{\max})^\psi$) rather than through a region-specific δ , so the structural parameter should be estimated on the same population the model covers, with the depletion control absorbing the post-frontier margin. The identified range nonetheless runs higher: the Legal-Amazon frontier estimate is $\delta = 0.62$ ($\varepsilon_S \approx 1.6$)—reproduced by geography-predicted (Fiszbein) exposure shares (Appendix D.1.5)—and persistent-price specifications reach $\delta \approx 0.75$. Re-inverting the full model at each value (Appendix Table 20) shows what the range implies: cumulative Priority-List leakage rises monotonically from 3.7% at the calibrated baseline to 8.8% at the frontier $\delta = 0.62$, and Protected-Areas leakage from 6.4% to 14.7%. We accordingly report headline leakage as the range 3.7–8.8% (Priority List) and 6.4–14.7% (Protected Areas), treating the baseline as the point most consistent with the model’s structure and the frontier estimate as the upper bound the data admit, and we carry both values wherever a quantitative conclusion depends on δ . Only well above the identified range does leakage grow large: about 12% and 20% at $\delta = 0.70$, and 20% and 32% at $\delta = 0.80$ ($\varepsilon_S = 4$, beyond anything the data support)—still far below the one-for-one prior.²⁴ The same sweep re-inverts and re-solves the model

²⁴Full re-inversions of the model at each parameter value (parameter sweep over δ). The sufficient-statistic reading of equation (3)—evaluated at the model-implied $\varepsilon_D \approx 15$ at the 2018 horizon, recomputing ε_S and the effective banned share bf at each δ —is within about a percentage point of the full re-resolution near the baseline and understates it by a few points at the top of the δ range. The

across the empirically defensible ranges of every other key parameter— σ , θ , ζ , χ , α , β , κ , and a no-migration economy—one deviation at a time: a star design of one baseline plus 55 single-deviation cells, of which 54 converge (the $\delta = 0.75$ re-solution does not; its value is bracketed by the adjacent cells; Appendix Table 20). Outside of δ , no parameter moves cumulative Priority-List leakage by more than 0.6 percentage points— α is the largest mover at 0.6, σ manages 0.3, and every other parameter 0.2 or less; for Protected Areas the largest non- δ mover is α at 1.2 percentage points—and shutting internal migration off entirely leaves leakage at 3.7% (6.4% for Protected Areas). Even far beyond the identified range the level stays bounded: 40% and 54% at an implausible $\delta = 0.90$ ($\varepsilon_S = 9$)—below the one-for-one prior everywhere the sweep reaches. Most importantly, the theory’s qualitative prediction is invariant: the Priority List leaks less than Protected Areas in every one of the 55 converged cells. The *ranking*—the paper’s mechanism result—survives the entire parameter space; only the level moves, and within the identified range it stays small.

Substitution elasticities The baseline $\sigma = 9.6$ is already at the high end of the literature, giving relocation ample scope; Propositions 1–2 imply leakage rises with σ and falls with ε_D , so the supply-driven mechanism—not a low substitution elasticity—is what keeps leakage small. The sweep confirms this quantitatively: re-inverting at each $\sigma \in [6, 12]$ moves Priority-List leakage by less than one percentage point in total (Appendix Table 20). Indeed, at the calibrated baseline the non- δ comparative statics are so flat—no parameter moves leakage by more than 0.6 percentage points across its swept range—that their signs are not empirically resolved; the theory’s directional predictions are properties of the mechanism, while the sweep’s message is that only the supply elasticity matters quantitatively.

SUTVA bias in the reduced form Small leakage also implies that the local estimates of Section 3 are approximately valid, and we verify this directly rather than assert it: using the model as the data-generating process, we run the paper’s own synthetic difference-in-differences estimator on model-generated panels with and without general-equilibrium spillovers. The spillovers attenuate the recovered treatment effect by +0.013 log points—about 3% of the estimate, in line with the 3.7% leakage

$\delta = 0.75$ cell does not converge—the model’s fixed point becomes ill-conditioned this far above the calibrated baseline—and is bounded by its converged neighbors (12% at $\delta = 0.70$, 20% at $\delta = 0.80$) together with the monotonicity of the δ -path.

share—and the full-GE model analogue (-0.43) closely reproduces the reduced-form estimate on actual data (-0.44), an out-of-sample validation of the model against quasi-experimental evidence it was never calibrated to fit (Appendix A.4, Table 11).

7.4 Structural transformation and welfare

Conservation policy operates on the land margin of a developing economy, and the model prices its consequences beyond deforestation. Table 8 collects the general-equilibrium effects of the two policies over 2003–2018, alongside their implications for structural transformation.

The effects on development are real but modest, and they are the flip side of small leakage: because the targeted land was cheap, the economy loses little by not clearing it. At the national level the Priority List raises the agricultural employment share slightly ($+0.10$ pp by 2018, about 96,000 workers) and leaves real income per worker essentially unchanged (weakly positive, below a hundredth of a percent). This aggregate masks an informative spatial reallocation—the same displacement that generates leakage. In the targeted municipalities, scarcer land pushes labor *out* of agriculture; in the non-targeted regions that absorb displaced clearing, labor is pulled *into* agriculture. Rather than accelerating structural transformation, the policy’s intensification response—more labor per hectare on less land—retains workers in farming on net. Protected Areas, which bind on land that was not being farmed, have effects an order of magnitude smaller.

Figure 11 maps the spatial reallocation of the agricultural employment share for the Priority List, making the mechanism visible: the targeted regions shed agricultural labor while the untargeted regions that absorb displaced clearing gain it.

Because the targeted land is cheap, conservation is also inexpensive in welfare terms. We compute the economic cost of each policy as the general-equilibrium change in aggregate real income between the with-policy baseline and the no-policy counterfactual, and divide it by the carbon content of the net avoided deforestation—the region-specific biomass released per hectare cleared, converted to CO_2 —to obtain a cost per tonne of CO_2 avoided (Table 9).²⁵ The implied cost is essentially zero and,


²⁵Real-income costs are cumulated over 2003–2018 and dollarized to 2018 USD; avoided CO_2 applies the same region carbon densities used in the counterfactuals, with the standard 44/12 carbon-to- CO_2 factor. We benchmark against a *range* of social costs of carbon rather than a single value—roughly \$50–\$190 per tonne, spanning the 2021 U.S. Interagency Working Group central estimate

Table 8: General-equilibrium effects of the two policies, 2003–2018 (open economy)

	Priority List	Protected Areas
<i>Deforestation</i>		
Gross avoided (km ²)	80,051	7,311
GE leakage (km ²)	2,935	467
Net avoided (km ²)	77,116	6,844
Leakage (%)	3.7	6.4
<i>Structural transformation & incomes (Δ vs. no policy, 2018)</i>		
Ag. employment share (pp)	+0.10	−0.01
Ag. employment (000s workers)	+95.9	−7.3
Real income per worker (%)	+0.00	+0.01

Note: All quantities are cumulative changes relative to the no-policy counterfactual (counterfactual A) by 2018, from the open-economy engine output. Deforestation rows follow the definitions of Table 6. Structural-transformation rows report the change in the national agricultural employment share (percentage points), agricultural employment (thousands of workers), and real income per worker (%).

Figure 11: Priority List: spatial reallocation of the agricultural employment share



Note: Change in each micro-region’s agricultural employment share between the with-policy baseline and the no-Priority-List counterfactual, 2018. Targeted (listed) regions lose agricultural employment share as land is made scarce; non-targeted regions that absorb displaced clearing gain it.

if anything, marginally negative—about $-\$0.001$ per tonne for the Priority List and $-\$0.004$ for Protected Areas—because the policies leave aggregate real income all but (about $\$51$ at a 3% discount rate) and the 2022 U.S. EPA proposed value (about $\$185$).

unchanged (indeed slightly higher) while avoiding roughly 4.1 and 0.34 billion tonnes of CO₂ respectively. Both policies thus avoid emissions at a net economic cost that is orders of magnitude below even the lowest social cost of carbon we consider, and of the opposite sign.²⁶

Table 9: Economic cost per tonne of CO₂ avoided

Policy	Net avoided deforestation (km ²)	Avoided CO ₂ (Gt)	Implied cost (\$/t CO ₂)
Priority List	77,116	4.14	-0.001
Protected Areas	6,844	0.34	-0.004

Note: Net avoided deforestation (2003–2018, open economy) converted to avoided CO₂ using region-specific carbon densities. “Implied cost” is the cumulative general-equilibrium real-income change divided by avoided CO₂, in 2018 USD; a *negative* cost means the policy marginally raises aggregate real income. Social-cost-of-carbon benchmarks span roughly \$50–\$190 per tonne (see text); the implied costs are well under 1% of the lowest of these.

This should not be read as conservation being costless, still less as free money left on the table. The figure is an *aggregate output* cost, and two wedges separate it from the decision problem agents actually face. First, it excludes the fiscal cost of enforcement—placing the Amazon under these instruments runs to at least \$1.7 billion per year (da Silva et al. 2022)—along with the administrative and monitoring capacity they require. Second, the near-zero *aggregate* nets winners against losers: making land scarcer raises land rents for incumbent landowners and shifts agricultural activity and employment toward the untargeted regions that absorb displaced clearing (Figure 11), while frontier would-be deforesters forgo the private clearing profits that drive deforestation in the first place and consumers face marginally higher food prices. The near-zero aggregate cost is the same phenomenon as small leakage—forgoing cheap land sacrifices little output because intensification and reallocation absorb the adjustment—and the gap between a trivial national cost and a large global carbon benefit (billions of tonnes of CO₂ valued at any plausible social cost) is precisely the externality these policies internalize. What limits their extension, then, is the

²⁶This is the general-equilibrium efficiency cost of each policy, not its fiscal cost. Enforcement spending adds little per tonne: extrapolating the minimum command-and-control cost of da Silva et al. (2022)—about \$1.7 billion per year to place 80% of the Legal Amazon (some four million km²) under enforcement—to the Priority List’s roughly one million km² over 2008–2018 implies on the order of \$1 per tonne of avoided CO₂, still two orders of magnitude below the social costs of carbon we benchmark against.

misalignment between who bears the cost—Brazil’s fisc and frontier profiteers—and who reaps the carbon benefit, not a large national output sacrifice. A full welfare accounting of this incidence, and the design of efficient rather than merely observed targeting, we leave to future work.

8 Conclusion

To what extent are spatially targeted conservation policies globally effective, once we account for the possible geographical displacement of environmental damage? The answer has deep implications for policy design, given the global public good nature of ecosystem conservation. We address this question in the context of tropical deforestation—an activity responsible for one-fifth of global CO₂-equivalent emissions in the past two decades—in Brazil, home to a third of the world’s remaining rainforests.

Conservation leakage is intrinsically a general equilibrium problem in space: a policy crackdown on deforestation in a delimited region changes economic incentives through the price of deforested land, generating sectoral and spatial reallocation that reduced-form evaluations cannot separate from genuine conservation. We build a spatial general equilibrium model of the Brazilian economy in which agricultural land is the endogenous output of a costly deforestation sector, and use it to quantify the leakage of the two flagship spatially targeted policies of the past decades. Over 2003–2018, general-equilibrium adjustment undoes 3.7% of the deforestation avoided by the Priority List and 6.4% of that avoided by Protected Areas at the baseline all-Brazil calibration of the deforestation supply elasticity, and 8.8% and 14.7% at the frontier estimate that bounds the identified range from above: the policies retain roughly 96% and 94% of their local gains at the national level at baseline, and no less than about 91% and 85% anywhere in the range.

The paper’s central lesson is why. Leakage is governed by whether targeted deforestation is supply-driven or demand-driven, and Brazilian deforestation is overwhelmingly supply-driven: the biggest deforesters are poor, remote, pasture-intensive places where clearing is cheap, not places where land is valuable. A policy that targets them—as the Priority List, by construction, does—removes the economy’s cheapest suppliers of new land, leaving a steep residual supply curve; the resulting land-price increase is absorbed by intensification and substitution toward less land-intensive

economic activity rather than by relocated clearing. Protected Areas leak more per unit of avoided deforestation precisely because they are not targeted to that frontier. For policy design, the implication travels beyond Brazil: target the economic margin that generates deforestation, and diagnose that margin—the relative land price of candidate target regions is an observable sufficient statistic for leakage risk.

This macro approach and the micro evaluation literature are complements. Our results largely vindicate reduced-form designs in this setting—running the paper’s own estimator on model-generated data with and without spillovers shows the SUTVA bias in local estimates of the Priority List is about 3% of the estimated effect (Appendix [A.4](#))—but that conclusion could not have been reached without a general-equilibrium model, and it would reverse in a demand-driven setting. Our framework provides a starting point for other questions requiring a spatial model of deforestation and agriculture: the welfare cost per tonne of carbon and efficient targeting, the consequences of conservation for structural transformation, the effects of deforestation on rainfall in nearby areas ([Leite-Filho et al. 2021](#)), or the role of agricultural technological change in environmental conservation.

References

- Adão, Rodrigo, Michal Kolesár, and Eduardo Morales. Shift-Share Designs: Theory and Inference*. *The Quarterly Journal of Economics*, 134(4):1949–2010, 08 2019. ISSN 0033-5533. doi: 10.1093/qje/qjz025. URL <https://doi.org/10.1093/qje/qjz025>.
- Akerman, Ariel. Deforestation and the demographic transition: Lessons from Brazil. Unpublished manuscript, 2025.
- Alix-Garcia, Jennifer M., Elizabeth N. Shapiro, and Katharine R. E. Sims. Forest conservation and slippage: Evidence from mexico’s national payments for ecosystem services program. *Land Economics*, 88(4):613–638, 2012. ISSN 0023-7639. doi: 10.3368/le.88.4.613. URL <https://le.uwpress.org/content/88/4/613>.
- Alston, Lee J., Gary D. Libecap, and Bernardo Mueller. *Titles, Conflict, and Land Use: The Development of Property Rights and Land Reform on the Brazilian Amazon Frontier*. University of Michigan Press, 1999. ISBN 9780472110063. URL <http://www.jstor.org/stable/10.3998/mpub.16208>.
- Alvarez, Jorge A. The agricultural wage gap: Evidence from brazilian micro-data. *American Economic Journal: Macroeconomics*, 12(1):153–73, January 2020. doi: 10.1257/mac.20170436. URL <https://www.aeaweb.org/articles?id=10.1257/mac.20170436>.
- Araujo, Rafael, Francisco Costa, and Marcelo Sant’Anna. Efficient conservation of the Brazilian Amazon: Estimates from a dynamic model. *The Review of Economic Studies*, 93(1):72–105, 2026. doi: 10.1093/restud/rdaf031.
- Arkhangelsky, Dmitry, Susan Athey, David A. Hirshberg, Guido W. Imbens, and Stefan Wager. Synthetic difference-in-differences. *American Economic Review*, 111(12):4088–4118, December 2021. doi: 10.1257/aer.20190159. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20190159>.
- Assunção, Juliano and Romero Rocha. Getting greener by going black: the effect of blacklisting municipalities on amazon deforestation. *Environment and Development Economics*, 24(2):115–137, 2019. doi: 10.1017/S1355770X18000499.

- Assunção, Juliano, Clarissa Gandour, and Rudi Rocha. Deforestation slowdown in the brazilian amazon: prices or policies? *Environment and Development Economics*, 20(6):697–722, 2015. doi: 10.1017/S1355770X15000078.
- Assunção, Juliano, Robert McMillan, Joshua Murphy, and Eduardo Souza-Rodrigues. Optimal environmental targeting in the amazon rainforest. *The Review of Economic Studies*, 90:1608–1641, Jul 2023. URL <https://doi-org.libproxy.ucl.ac.uk/10.1093/restud/rdac06>.
- Balboni, Clare, Aaron Berman, Robin Burgess, and Benjamin A. Olken. The economics of tropical deforestation. *Annual Review of Economics*, 15(1):723–754, 2023. doi: 10.1146/annurev-economics-090622-024705.
- Baragwanath, Kathryn and Ella Bayi. Collective property rights reduce deforestation in the Brazilian Amazon. *Proceedings of the National Academy of Sciences*, 117(34):20495–20502, 2020. doi: 10.1073/pnas.1917874117.
- BenYishay, Ariel, Silke Heuser, Daniel Runfola, and Rachel Trichler. Indigenous land rights and deforestation: Evidence from the Brazilian Amazon. *Journal of Environmental Economics and Management*, 86:29–47, 2017. doi: 10.1016/j.jeem.2017.07.008.
- Berman, Nicolas, Mathieu Couttenier, Antoine Leblois, and Raphael Soubeyran. Crop prices and deforestation in the tropics. *Journal of environmental economics and management*, 119:102819–, 2023. ISSN 0095-0696.
- Boppart, Timo. Structural change and the kaldor facts in a growth model with relative price effects and non-gorman preferences. *Econometrica*, 82(6):2167–2196, 2014. doi: <https://doi.org/10.3982/ECTA11354>. URL <https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA11354>.
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel. Quasi-experimental shift-share research designs. *The Review of Economic Studies*, 89(1):181–213, 2022.
- Buggle, Johannes, Thierry Mayer, Seyhun Orcan Sakalli, and Mathias Thoenig. The refugee’s dilemma: Evidence from jewish migration out of nazi germany. *The Quarterly journal of economics*, 138(2):1273–1345, 2023. ISSN 0033-5533.

- Burgess, Robin, Francisco J. M. Costa, and Benjamin A. Olken. The Brazilian Amazon’s double reversal of fortune. SocArXiv 67xg5, Center for Open Science, 2019.
- Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli. Agricultural productivity and structural transformation: Evidence from brazil. *American Economic Review*, 106(6):1320–65, June 2016. doi: 10.1257/aer.20131061. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20131061>.
- Cisneros, Elías, Sophie L. Zhou, and Jan Börner. Naming and shaming for conservation: Evidence from the brazilian amazon. *PLoS ONE*, 10(9):e0136402, 2015. doi: 10.1371/journal.pone.0136402.
- Clarke, Damian, Daniel PailaÑir, Susan Athey, and Guido Imbens. On synthetic difference-in-differences and related estimation methods in Stata. *The Stata Journal*, 24(4):557–598, 2024. doi: 10.1177/1536867X241297914.
- Costinot, Arnaud, Dave Donaldson, and Cory Smith. Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world. *Journal of Political Economy*, 124(1):205–248, 2016. doi: 10.1086/684719.
- Da Mata, Daniel and Mario Dotta. Commodity booms and the environment, 2021. Working paper, SSRN No. 3900793.
- da Silva, José Maria Cardoso, Luís Claudio Fernandes Barbosa, Julie Topf, Ima Célia G Vieira, and Fabio R Scarano. Minimum costs to conserve 80% of the brazilian amazon. *Perspectives in Ecology and Conservation*, 20(3):216–222, 2022.
- de Vasconcelos, José Romeu. Matriz do fluxo de comércio interestadual de bens e serviços no brasil – 1999. Texto para discussão, Instituto de Pesquisa Econômica Aplicada (IPEA), August 2001.
- Dekle, Robert, Jonathan Eaton, and Samuel Kortum. Unbalanced trade. *American Economic Review*, 97(2):351–355, 2007.
- Desmet, Klaus and Esteban Rossi-Hansberg. Spatial development. *American Economic Review*, 104(4):1211–1243, 2014.

- Dix-Carneiro, Rafael and Brian K Kovak. Trade liberalization and regional dynamics. *American Economic Review*, 107(10):2908–2946, 2017.
- Domínguez-Iino, Tomás. Efficiency and redistribution in environmental policy: An equilibrium analysis of agricultural supply chains. *Journal of Political Economy*, 134(4), 2026.
- Donaldson, Dave and Richard Hornbeck. Railroads and american economic growth: A “market access” approach. *The Quarterly Journal of Economics*, 131(2):799–858, 02 2016. ISSN 0033-5533. doi: 10.1093/qje/qjw002. URL <https://doi.org/10.1093/qje/qjw002>.
- DOU. Lei nº 6.321, de 21 de dezembro de 2007. *Diário Oficial [da] República Federativa do Brasil*, 2007.
- DOU. Lei nº 12.651, de 25 de maio de 2012. *Diário Oficial [da] República Federativa do Brasil*, 2012.
- Eckert, Fabian and Michael Peters. Spatial structural change. Working Paper 30489, National Bureau of Economic Research, September 2022. URL <http://www.nber.org/papers/w30489>.
- Farrokhi, Farid and Heitor S Pellegrina. Trade, technology, and agricultural productivity. *Journal of Political Economy*, 131(9):2509–2555, 2023. doi: 10.1086/724319.
- Farrokhi, Farid, Elliot Kang, Sebastian Sotelo, and Heitor S Pellegrina. Deforestation: A global and dynamic perspective. Working Paper 34150, National Bureau of Economic Research, August 2025. URL <http://www.nber.org/papers/w34150>.
- Fiszbein, Martin. Agricultural diversity, structural change, and long-run development: Evidence from the united states. *American Economic Journal: Macroeconomics*, 14(2):1–43, 2022.
- Fuller, Carley, Stefania Ondeï, Barry W. Brook, and Jessie C. Buettel. First, do no harm: A systematic review of deforestation spillovers from protected areas. *Global Ecology and Conservation*, 18:e00591, 2019. ISSN 2351-9894. doi: <https://doi.org/10.1016/j.gecco.2019.e00591>. URL <https://www.sciencedirect.com/science/article/pii/S2351989419301143>.

- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624, 2020.
- Gollin, Douglas. Getting income shares right. *Journal of Political Economy*, 110(2): 458–474, 2002.
- Gollin, Douglas and Julien Wolfersberger. Agricultural trade and deforestation: The role of new roads. *The Economic Journal*, 2025. doi: 10.1093/ej/ueaf125.
- Guerriero, Marta. The labor share of income around the world: Evidence from a panel of countries and industries. In Fields, Gary and Saumik Paul, editors, *Labor Income Share in Asia*, pages 39–79. Springer, Singapore, 2019.
- Hansen, Casper Worm and Asger Mose Wingender. National and global impacts of genetically modified crops. *American Economic Review: Insights*, 5(2):224–240, 2023. doi: 10.1257/aeri.20220144.
- Hsiao, Allan. Coordination and commitment in international climate action: Evidence from palm oil. *Econometrica*, 94(1):1–33, 2026. doi: 10.3982/ECTA20608.
- IPCC. Ipcc, 2019: Climate change and land: an ipcc special report on climate change, desertification, land degradation, sustainable land management, food security, and greenhouse gas fluxes in terrestrial ecosystems. 2019.
- Lagakos, David and Michael E. Waugh. Selection, agriculture, and cross-country productivity differences. *American Economic Review*, 103(2):948–80, April 2013. doi: 10.1257/aer.103.2.948. URL <https://www.aeaweb.org/articles?id=10.1257/aer.103.2.948>.
- Leite-Filho, Argemiro Teixeira, Britaldo Silveira Soares-Filho, Juliana Leroy Davis, Gabriel Medeiros Abrahão, and Jan Börner. Deforestation reduces rainfall and agricultural revenues in the brazilian amazon. *Nature Communications*, 12(1): 2591, 2021. doi: 10.1038/s41467-021-22840-7. URL <https://doi.org/10.1038/s41467-021-22840-7>.
- Mishra, Prakash. The global allocative efficiency of deforestation. Working paper, 2026. URL https://mishrap42.github.io/files/PrakashMishra_JMP.pdf.

- Moretz-Sohn, Caio. Protected area erasure accelerates deforestation in the Brazilian Amazon, 2026. URL https://osf.io/preprints/socarxiv/wgkty_v1. SocArXiv preprint wgkty. R&R, Journal of the Association of Environmental and Resource Economists.
- Morten, Melanie and Jaqueline Oliveira. The effects of roads on trade and migration: Evidence from a planned capital city. *American Economic Journal: Applied Economics*, 16(2):389–421, 2024. doi: 10.1257/app.20180487.
- Nolte, Christoph, Arun Agrawal, Kirsten M. Silvius, and Britaldo S. Soares-Filho. Governance regime and location influence avoided deforestation success of protected areas in the Brazilian Amazon. *Proceedings of the National Academy of Sciences*, 110(13):4956–4961, 2013. doi: 10.1073/pnas.1214786110.
- Pellegrina, Heitor S. Trade, productivity, and the spatial organization of agriculture: Evidence from Brazil. *Journal of Development Economics*, 156:102816, 2022. doi: 10.1016/j.jdeveco.2022.102816.
- Pellegrina, Heitor S and Sebastian Sotelo. Migration, specialization, and trade: Evidence from Brazil’s March to the West. *Journal of Political Economy*, 133(12): 3993–4049, 2025. doi: 10.1086/738147.
- Pfaff, Alexander and Juan Robalino. Spillovers from conservation programs. *Annual Review of Resource Economics*, 9(1):299–315, 2017. doi: 10.1146/annurev-resource-100516-053543. URL <https://doi.org/10.1146/annurev-resource-100516-053543>.
- Qin, Yuanwei, Xiangming Xiao, Fang Liu, Fabio de Sa e Silva, Yosio Shimabukuro, Egidio Arai, and Philip Martin Fearnside. Forest conservation in Indigenous territories and protected areas in the Brazilian Amazon. *Nature Sustainability*, 6: 295–305, 2023. doi: 10.1038/s41893-022-01018-z.
- Redding, Stephen J. Goods trade, factor mobility and welfare. *Journal of International Economics*, 101:148–167, 2016.
- Redding, Stephen J. and Esteban Rossi-Hansberg. Quantitative spatial economics. *Annual Review of Economics*, 9:21–58, 2017.

- Robalino, Juan, Alexander Pfaff, and Laura Villalobos. Heterogeneous local spillovers from protected areas in costa rica. *Journal of the Association of Environmental and Resource Economists*, 4(3):pp. 795–820, 2017. ISSN 23335955, 23335963. URL <https://www.jstor.org/stable/26544527>.
- Roberts, Michael J. and Wolfram Schlenker. Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate. *American Economic Review*, 103(6):2265–2295, 2013. doi: 10.1257/aer.103.6.2265.
- Scott, Paul T. Dynamic discrete choice estimation of agricultural land use. Working paper, Toulouse School of Economics, 2013.
- Soares-Filho, Britaldo, Paulo Moutinho, Daniel Nepstad, Anthony Anderson, Hermann Rodrigues, Ricardo Garcia, Laura Dietzsch, Frank Merry, Maria Bowman, Letícia Hissa, Rafaella Silvestrini, and Cláudio Maretti. Role of Brazilian Amazon protected areas in climate change mitigation. *Proceedings of the National Academy of Sciences*, 107(24):10821–10826, 2010. doi: 10.1073/pnas.0913048107.
- Sotelo, Sebastian. Domestic trade frictions and agriculture. *Journal of Political Economy*, 128(7):2690–2738, 2020. doi: 10.1086/706859.
- Souza-Rodrigues, Eduardo. Deforestation in the Amazon: A unified framework for estimation and policy analysis. *The Review of Economic Studies*, 86(6):2713–2744, 2019. doi: 10.1093/restud/rdy070.
- Tsuda, Shunsuke, Yoshito Takasaki, and Mari Tanaka. Human and nature: Economies of density and conservation in the Amazon rainforest. Working paper, 2025. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4652170.

A Reduced-form analysis: additional results and robustness


A.1 Additional descriptives

Figure 12 shows the evolution over time of the area under each spatially targeted policy. Protected Areas have been established gradually since redemocratization in

1985. The Priority List started in 2008, when 36 municipalities covering approximately 800,000 square kilometers were added; the total number of listed municipalities has since remained around 40, with total area increasing to around 1,000,000 square kilometers. Total deforestation in the five preceding years is a good predictor of Priority List status, although it does not exactly determine inclusion. Since 2008, 20 municipalities have exited the list—acquiring the status of “monitored”—and 3 of those have returned. There have been two periods of significant change: 2008–2012, with most of the action concentrated in 2008, and 2017–2020, with most of the action concentrated in 2017, when 9 municipalities covering over 200,000 square kilometers were added, reflecting a shift of deforestation further north, toward the Amazon river.


Figure 13 shows trends in natural vegetation cover by 2020 protection status. Between 1985 and 2020, the bulk of forest loss happened outside protected areas: forest cover in Indigenous Territories and Conservation Units remained above 90% and 80% respectively, while cover outside them fell from above 60% to under 50%. This comparison alone is not evidence of policy effectiveness—placement is endogenous (the “additionality” question) and the comparison is silent on displacement (the “leakage” question)—which is what motivates the designs in Section 3 and, ultimately, the model.

Figure 12: Evolution of spatially targeted conservation policies



Note: This figure illustrates the evolution over time of the total area under different types of locally targeted policies by the Brazilian government. The yellow line represents the total area of demarcated Indigenous Territories, the orange line represents the total area under Conservation Units, and the green line represents the total area of all municipalities included in the Priority List policy.

Figure 13: Area covered in natural vegetation under different conservation policies



Note: This figure illustrates the evolution over time of the territorial coverage in natural vegetation, as a share of total area, for regions under the different types of locally targeted policies implemented by the Brazilian government. The yellow line represents the evolution of natural vegetation coverage in Indigenous Territories, the orange line represents the evolution of natural vegetation coverage in Conservation Units, and the green line represents the evolution of natural vegetation cover in municipalities from the Priority List.

A.2 Priority List: crop-level agricultural outcomes

Table 10 reports synthetic difference-in-differences estimates, as in equation (1), of the effect of the Priority List on agricultural value added and on crop-specific revenue, harvested area, farm-gate prices, and yields. The expansion of cropland documented in Section 3.1 is concentrated in maize and soy, whose revenues and harvested areas increase significantly, while beans—a smallholder staple—contract.

A.3 Protected Areas: specification, estimates, and placebo

Specification. We work on a fine 1 km grid, recording for each cell whether its centroid is forested, and estimate a regression discontinuity design around Protected-Area borders:

$$\text{Forest}_{ht} = \delta_t + \gamma_{m(h)} + f(\text{Distance}_{ht}) + \beta D_{ht} + \epsilon_{ht}, \quad (17)$$

where $\text{Forest}_{ht} = 1$ if the centroid of cell h is forested at year t (a linear probability model), δ_t are year fixed effects, $\gamma_{m(h)}$ are municipality fixed effects, $f(\cdot)$ is a continuous function (a quadratic spline split at the border), Distance_{ht} is the running variable

Table 10: Effects of Priority List on Crop-specific Agricultural Outcomes

Outcome	ATT	Std. err.	Obs.
(log) soy yield	-0.000	(0.069)	4,954
(log) maize yield	0.150*	(0.079)	15,527
(log) rice yield	0.072	(0.070)	14,160
(log) cassava yield	-0.020	(0.050)	15,787
(log) beans yield	-0.107	(0.090)	12,906
(log) cattle per ha	0.184*	(0.101)	16,063

Note: This table shows the results of the municipality-year-level synthetic differences in differences regressions of agricultural value added and crop-specific outcomes (specifically: revenue from crop, area harvested, average farm-gate price, and average yields) in logarithms on the onset of the Priority List. The first row is a regression as (1) but with outcome the log of agricultural value added as obtained from the system of regional accounts. The second row onwards are regressions as (1) with outcomes calculated from the yearly municipal agricultural survey PAM. Standard errors are calculated via bootstrap with 50 repetitions. We only include municipalities in the Brazilian Legal Amazon, as they are more likely to resemble Priority List municipalities. Years in which the crop is not grown (according to the data) are dropped. Subsequently, municipalities that have not grown a crop (in a year that was not dropped in the previous step). Standard errors in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

measuring the distance in km to the border of the nearest protected area—negative (positive) values mean the cell falls inside (outside)—and $D_{ht} = 1\{\text{Distance}_{ht} \leq 0\}$ indicates falling within a protected area. The coefficient β is thus the discontinuity in the probability of forest cover at the border. We use forest cover rather than deforestation as the outcome because protected areas are typically established in regions with very low ongoing deforestation, so the interpretation of the effect is clearer. The fine 1 km grid lets us examine an arbitrarily small neighborhood of the cut-off; the placebo exercise below probes the continuity assumption directly.

We estimate β under three sets of fixed effects—year only, municipality, and municipality-year, the last absorbing the political process of establishment. The discontinuity is stable across all three: forest cover is roughly 13 to 15 percentage points more likely just inside a protected area than just outside.

Placebo: borders of future Protected Areas. To probe the smoothness assumption, we examine what happens at the border of a future protected area before

it is established:

$$\text{Forest}_{ht} = \delta_t + \gamma_{m(h)} + f(\text{MinDistance}_h) + \tilde{\beta}\tilde{D}_h + \epsilon_{ht}, \quad (18)$$

where $\text{MinDistance}_h \equiv \min_t \text{Distance}_{ht}$ is the time-invariant minimum distance that a hexagon h ever has from a Protected Area—which is also the final one, because protected areas never get dismantled and hence Distance_{ht} can only decrease—and $\tilde{D}_h = 1\{\text{MinDistance}_h \leq 0\}$. For each hexagon we define the event time $T(h)$ from which the distance does not further decrease, and restrict the sample to observations (h, t) such that $t < T(h) + k$, where k is the number of years before establishment considered. Figure 14 shows the border contrast by event time, both in the raw bin means (panel a) and after residualizing on municipality-year fixed effects (panel b). The two give the same reading: no discontinuity is visible more than ten years before establishment; a smaller one emerges in the years immediately preceding it; and the gap widens once protection is formal. A discontinuity around the border of a future protected area can arise from (i) lengthy legal disputes over demarcation, during which deforestation decreases before final settlement,²⁷ or (ii) historical occupation of the area by indigenous people whose livelihood depends on the preservation of the forest before legal demarcation. This anticipatory divergence is why the border contrast in Section 3.2 is presented as descriptive evidence that protection binds, rather than as a causal estimate.


A.4 Bounding the SUTVA bias: the model as data-generating process

The identification of the local effects in Section 3 relies on the stable unit treatment value assumption: control regions must be unaffected by treatment elsewhere. Leakage violates it directly—displaced clearing raises deforestation in the donor pool, biasing the estimated effect toward zero. Rather than asserting from the aggregate leakage number that this bias is small, we measure it, using the model as the data-generating process and the paper’s own estimator as the instrument.


We generate two synthetic deforestation panels at the model’s native micro-region

²⁷This is particularly likely in the case of Indigenous Territories. Their establishment follows a judicial process that includes several stages: study, delimitation, declaration, homologation, and regularization, that can take decades.

Figure 14: Forest cover at Protected-Area borders, by period relative to establishment



(a) Raw bin means



(b) Residualized on municipality-year fixed effects

Note: Mean forest cover in distance bins around Protected-Area borders, by period relative to establishment (cells classified by the minimum—and final—distance they ever reach from a border). Panel (a) plots the raw bin means, so the levels are directly interpretable as forest-cover shares; panel (b) plots the same bins after residualizing on municipality-year fixed effects, so that only variation in distance to the border within a municipality and year identifies the contrast. The two panels give the same reading: the inside–outside gap predates formal protection—consistent with lengthy demarcation disputes and prior occupation by forest-preserving indigenous communities—but it is visibly smaller before establishment and widens after it: protection binds.

level, identical except for the presence of general-equilibrium spillovers. Panel A is the full-GE outcome: log deforestation from the with-policy baseline against the no-Priority-List counterfactual, embedding leakage. Panel B shuts the spillover channel down: each region’s deforestation responds to its own policy shock holding all equilibrium prices at their no-policy values, so untreated regions do not move by construction. We then run the staggered synthetic difference-in-differences estimator of Section 3.1—the same estimator, cohort structure, and treatment definition, with the 31 Priority-List micro-regions as treated units—on each panel. The difference between the two recovered ATTs is the SUTVA/spillover bias in the reduced-form

design.

Table 11 reports the result. A first observation validates the exercise before the headline: the full-GE model analogue (-0.43) is close to the reduced-form estimate on actual data (≈ -0.44 ; Table 1)—an out-of-sample check of the model against the quasi-experimental evidence it was never calibrated to fit, since the Priority-List shock enters the model only through the inverted fundamentals, not through the estimated treatment effect.

The bias itself is $+0.013$ log points (bootstrap s.e. 0.002): general-equilibrium spillovers attenuate the estimated treatment effect by about 3% of its absolute value—closely in line with the model’s cumulative leakage share of 3.7%, as the theory predicts, and an order of magnitude too small to overturn any conclusion drawn from the local estimates. The exercise turns the paper’s SUTVA caveat into a result: in this supply-driven setting the reduced-form toolkit is approximately valid, precisely because leakage is small. In a demand-driven setting the same exercise would deliver a large bias, so the validation is a property of the economics, not of the method. The bias is computed at the micro-region level, the model’s unit of analysis; it is the magnitude of the bias, not the point estimate of the ATT, that transfers to the municipality-level design of Section 3.1.

B Proofs of the results in Section 4

B.1 Proposition 1

Let the baseline equilibrium clear at price v_0 with national deforestation $T_0 = S(v_0) = D(v_0)$, where $S(v) = \sum_r (Z_r^D)^{\frac{1}{1-\delta}} (\delta v)^{\varepsilon_S}$ has constant elasticity ε_S and D has elasticity $-\varepsilon_D$ at v_0 . The ban removes suppliers accounting for share b of $S(v_0)$; residual supply is $S_{nb}(v) = (1 - b) T_0 (v/v_0)^{\varepsilon_S}$. Writing $\hat{v} = d \ln v$, market clearing to first order requires the increase in residual supply plus the demand contraction to absorb the removed quantity $b T_0$:

$$(1 - b) \varepsilon_S \hat{v} T_0 + \varepsilon_D \hat{v} T_0 = b T_0 \implies \hat{v} = \frac{b}{(1 - b) \varepsilon_S + \varepsilon_D}.$$

Table 11: SUTVA/spillover bias in the reduced-form design, bounded with the model as DGP

Priority List $\rightarrow$ (log) deforestation	Estimate
Reduced-form micro SDID, ATT (data)	-0.44
Model analogue, full GE (panel A)	-0.425 [0.121]
Model, no spillovers (panel B)	-0.438 [0.120]
SUTVA / spillover bias ($A - B$)	0.013*** [0.002]
as % of ATT	2.9%
Model cumulative leakage share	3.7%

Note: Synthetic difference-in-differences estimates of the Priority-List effect on log deforestation in model-generated panels at the micro-region level (31 treated regions, staggered entry as in the data). “Full GE” (panel A) is the model outcome including general-equilibrium spillovers; “no spillovers” (panel B) evaluates each region’s deforestation response at no-policy equilibrium prices, eliminating leakage by construction. The SUTVA/spillover bias is the difference $A - B$; bootstrap standard errors in brackets. The first row is the estimate on actual data from Table 1, shown for comparison. *p<0.1, **p<0.05, *** p<0.01.

Leakage is the induced increase in non-banned deforestation per unit removed:

$$L = \frac{(1 - b) \varepsilon_S \hat{v} T_0}{b T_0} = \frac{(1 - b) \varepsilon_S}{(1 - b) \varepsilon_S + \varepsilon_D}.$$

Comparative statics (i)–(iii) are immediate since $L = \chi/(\chi + \varepsilon_D)$ with $\chi \equiv (1 - b) \varepsilon_S$ increasing in ε_S and decreasing in b . For the exact solution with iso-elastic demand $D(v) = T_0(v/v_0)^{-\varepsilon_D}$: market clearing $(1 - b)(v/v_0)^{\varepsilon_S} = (v/v_0)^{-\varepsilon_D}$ gives $v/v_0 = (1 - b)^{-\frac{1}{\varepsilon_S + \varepsilon_D}}$ and

$$L^{exact} = \frac{(1 - b) \left[(1 - b)^{-\frac{\varepsilon_S}{\varepsilon_S + \varepsilon_D}} - 1 \right]}{b},$$

which converges to (3) as $b \rightarrow 0$ and exceeds it moderately at large b (e.g. at $\varepsilon_S = \varepsilon_D = 1$, $b = 0.4$: 0.436 vs 0.375). $\square$

B.2 Lemma 1

Write the two-region economy of Section 4.3 with $\kappa_r \equiv (Z_r^D)^{\frac{1}{1-\delta}}$, supply $T_r = \kappa_r v_r^{\varepsilon_S}$, output $Y_r = Z_r^A T_r$, producer price $p_r = v_r / Z_r^A$ (zero profits, land-only), and CES expenditure shares $\lambda_r \propto p_r^{1-\sigma}$. Market clearing $v_r T_r = \lambda_r X$ in ratio form gives

$$\frac{\lambda_1}{\lambda_2} = \omega x = \left(\frac{\omega}{z}\right)^{1-\sigma}, \quad z \equiv \frac{Z_1^A}{Z_2^A}, \quad \omega \equiv \frac{v_1}{v_2}, \quad x \equiv \frac{T_1}{T_2}, \quad (19)$$

so $x = \omega^{-\sigma} z^{\sigma-1}$, while supply gives $x = (\kappa_1 / \kappa_2) \omega^{\varepsilon_S}$. Fixing x : given any $z \geq 1$, (19) pins $\omega = z^{\frac{\sigma-1}{\sigma}} x^{-\frac{1}{\sigma}}$ (strictly increasing in z), and the supply equation then pins $\kappa_1 / \kappa_2 = x \omega^{-\varepsilon_S}$ uniquely. At $z = 1$ (pure supply): $\omega = x^{-1/\sigma} < 1$. The pure demand economy has $\kappa_1 / \kappa_2 = 1$, i.e. $\omega = x^{1/\varepsilon_S}$, reached at $z = x^{\frac{\sigma+\varepsilon_S}{\varepsilon_S(\sigma-1)}}$. Equilibrium log deforestation is linear in log fundamentals, $\ln T_r = \text{const} + \frac{\sigma}{\varepsilon_S + \sigma} \ln \kappa_r + \frac{\varepsilon_S(\sigma-1)}{\varepsilon_S + \sigma} \ln Z_r^A$, so the demand-driven share of $\ln x$ is $\frac{\varepsilon_S(\sigma-1)}{\varepsilon_S + \sigma} \ln z / \ln x$, strictly increasing in z , hence in ω . $\square$

B.3 Proposition 2

Log-linearize around the baseline. Let aggregate expenditure respond to the CES price index P with elasticity $1 - \varepsilon_D$ in value (so ε_D is the aggregate quantity elasticity), and hats denote log deviations. With $\hat{p}_r = \hat{v}_r$ (fundamentals Z^A unshocked) and $\hat{P} = \sum_r \lambda_r \hat{v}_r$, market clearing $v_r T_r = \lambda_r X(P)$ gives, for each region,

$$\hat{v}_r + \hat{T}_r = (1 - \sigma)(\hat{v}_r - \hat{P}) + (1 - \varepsilon_D)\hat{P},$$

and supply gives $\hat{T}_r = \hat{\kappa}_r + \varepsilon_S \hat{v}_r$. For the unshocked region ($\hat{\kappa}_2 = 0$): $(\varepsilon_S + \sigma)\hat{v}_2 = (\sigma - \varepsilon_D)\hat{P}$. Substituting into $\hat{P} = \lambda_1 \hat{v}_1 + \lambda_2 \hat{v}_2$ and solving:

$$\hat{P} = \frac{\lambda_1(\varepsilon_S + \sigma)}{\varepsilon_S + \sigma\lambda_1 + \varepsilon_D(1 - \lambda_1)} \hat{v}_1.$$

Then

$$\hat{T}_2 = \varepsilon_S \hat{v}_2 = \frac{\varepsilon_S(\sigma - \varepsilon_D)\lambda_1}{\varepsilon_S + \sigma\lambda_1 + \varepsilon_D(1 - \lambda_1)} \hat{v}_1, \quad \hat{T}_1 = \frac{\varepsilon_S\lambda_1(\sigma - \varepsilon_D) - \sigma(\varepsilon_S + \varepsilon_D)}{\varepsilon_S + \sigma\lambda_1 + \varepsilon_D(1 - \lambda_1)} \hat{v}_1,$$

where the second expression uses $\hat{T}_1 = \hat{\kappa}_1 + \varepsilon_S \hat{v}_1$ and $\hat{\kappa}_1 = -(\varepsilon_S + \sigma)\hat{v}_1 + (\sigma - \varepsilon_D)\hat{P}$. A negative supply shock ($\hat{\kappa}_1 < 0$) raises v_1 , and $\hat{T}_1 < 0$ always (the bracket is negative for $\lambda_1 \leq 1$, $\sigma > \varepsilon_D$). Leakage:

$$L = -\frac{dT_2}{dT_1} = -\frac{T_2}{T_1} \cdot \frac{\hat{T}_2}{\hat{T}_1} = \frac{1}{x} \cdot \frac{\varepsilon_S(\sigma - \varepsilon_D)\lambda_1}{\sigma(\varepsilon_S + \varepsilon_D) - \varepsilon_S(\sigma - \varepsilon_D)\lambda_1}.$$

Monotonicity: L has the form $c\lambda_1/(d - e\lambda_1)$ with $c, d, e > 0$ (for $\sigma > \varepsilon_D$), strictly increasing in λ_1 ; by Lemma 1, $\lambda_1 = \omega x/(1 + \omega x)$ is strictly increasing in ω , which is strictly increasing in the demand-driven share, holding x fixed. N -region version: for a marginal shock to region j , all unshocked regions share the same $\hat{v}_r$, so the derivation is identical with $\lambda_1 \rightarrow \lambda_j$ and $\frac{1}{x} \rightarrow \sum_{r \neq j} T_r/T_j$. Frictionless nesting: as $\sigma \rightarrow \infty$ with the pure supply-driven baseline, $\lambda_1 \rightarrow x/(1 + x)$ and (4) $\rightarrow \varepsilon_S/(\varepsilon_S + \varepsilon_D(1 + x))$, the marginal-incidence formula with residual supply mass T_2 and demand mass $T_1 + T_2$, i.e. Proposition 1 as $b \rightarrow 0$. All statements are verified against exact numerical solution over a 243-point parameter grid (maximum error $\sim 10^{-6}$), including the N -region generalization and the $\sigma \rightarrow \infty$ nesting. $\square$

C Model details


C.1 Agricultural Land Accumulation with Forest Regrowth

C.2 Deriving local expenditure by sector

Wages in equilibrium assuming Frechet Frank-copula. Let the joint distribution of pairs of individual productivities (z_{ir}^A, z_{ir}^{NA}) is given by the Frank copula as in LW with parameters $(\chi^A, \chi^{NA}, \rho)$. Workers choose the sector with the higher *amenity-weighted* effective wage $G_r^s w_r^s z_{ir}^s$, where G_r^s is a region–sector taste term normalized to $G_r^A \equiv 1$; the selection therefore depends on the amenity-adjusted return ratio $\tilde{w}_r \equiv G_r^A w_r^A / (G_r^{NA} w_r^{NA}) = w_r^A / (G_r^{NA} w_r^{NA})$. Amenities are a pure taste term: realized labor income still equals $w_r^s z_{ir}^s$, so only the selection boundary carries G_r^{NA} . There are no simple closed-form expressions for the share of employment in agriculture and the labor income in each sector. Instead,

$$s_r^A = Pr[z_{ir}^{NA}/z_{ir}^A \leq \tilde{w}_r] = \int_0^\infty \int_0^{z^A \tilde{w}_r} f(z^A, z^{NA}) dz^{NA} dz^A$$

Figure 15: Agricultural Land Accumulation with Forest Regrowth



Note: This figure illustrates the transition dynamics in the market for agricultural land under a constant forest regrowth rate. The graph on the left shows supply and demand curves for land, and the equation for the equilibrium in the market for the equilibrium quantity. The graph on the right shows the transition curve for agricultural land: with forest regrowth, deforestation reaches a steady state, where the blue curve intercepts the 45-degree line.

$$\bar{y}_r^{LA} = w_r^A \mathbb{E}[z_{ir}^A | z_{ir}^{NA}/z_{ir}^A \leq \tilde{w}_r] = w_r^A \frac{1}{s_r^A} \int_0^\infty \int_0^{z^A \tilde{w}_r} z^A f(z^A, z^{NA}) dz^{NA} dz^A$$

$$\bar{y}_r^{LNA} = w_r^{NA} \mathbb{E}[z_{ir}^{NA} | z_{ir}^{NA}/z_{ir}^A > \tilde{w}_r] = w_r^{NA} \frac{1}{s_r^{NA}} \int_0^\infty \int_{z^A \tilde{w}_r}^\infty z^{NA} f(z^A, z^{NA}) dz^{NA} dz^A$$

Setting $G_r^{NA} \equiv 1$ recovers the pure-selection (Roy) model, in which $\tilde{w}_r = w_r^A/w_r^{NA}$.

C.3 Model inversion

Given all production and preference parameters, the distribution of population, initial land endowments, and $K+3$ vectors of observable endogenous quantities

$$\{s_r^A, T_r^D, p^{NAr} Y_r^{NA}, \{T_r^{Ak}, p^{Akr} Y_r^{Ak}\}_k\}_r,$$

below we formulate the system of equations that needs to be solved in order to find the regional TFPs and amenities.

The Cobb-Douglas form of agricultural income of a each commodity k in region

r yields

$$p^{Akr}Y_r^{Ak} = \frac{v_r T_r^{Ak}}{\alpha_k} \quad (20)$$

Adding up over all k ,

$$v_r = \frac{1}{T_r^A} \sum_k \alpha_k p^{Akr} Y_r^{Ak}. \quad (21)$$

GDP accounting

$$\sum_k p^{Akr} Y_r^{Ak} + p_r^{NA} Y_r^{NA} = L_r \bar{y}_r^L + v_r T_r^A \quad (22)$$

so

$$\bar{y}_r^L = \frac{1}{L_r} \left(\sum_k p^{Akr} Y_r^{Ak} + p_r^{NA} Y_r^{NA} - v_r T_r^A \right) \quad (23)$$

Having the regional prices at origin (start with a guess), we can get goods prices at destination

$$p_d^s = \left(\sum_{r=1}^R (\tau_d^r p^{sr})^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \quad (24)$$

$$p_r^A = \left(\sum_{k=1}^K (p^{Akr})^{1-\theta} \right)^{\frac{1}{1-\theta}} \quad (25)$$

$$p_r = (p_r^A)^\phi (p_r^{NA})^{1-\phi} \quad (26)$$

Land prices

$$q_r = \frac{1}{1 - \beta(1 - \rho)} v_r \quad (27)$$

Deforestation investment

$$p_r I_r^D = \delta q_r T_r^D \quad (28)$$

Local household expenditure

$$E_r = v_r T_r^A + \bar{y}_r^L L_r - p_r I_r^D \quad (29)$$

Local expenditure shares in agriculture

$$\vartheta_r^A = \phi + \nu \left(\frac{E_r}{p_r L_r} \right)^{-\eta} \lambda_r \quad (30)$$

where $\lambda_r \equiv \frac{\mathbb{E}((y_{ir})^{1-\eta})}{(\mathbb{E}(y_{ir}))^{1-\eta}}$ Regional expenditures in each sector

$$X_d^A \equiv \vartheta_d^A E_d + \phi p_d I_d^D, \quad X_d^{NA} \equiv (1 - \vartheta_d^A) E_d + (1 - \phi) p_d I_d^D$$

$$X_d^{Ak} = X_d^A \left(\frac{p_d^{Ak}}{p_d^A} \right)^{1-\theta}$$

Rewrite trade shares as

$$\pi_d^{so} = (\tau_d^o p^{so})^{1-\sigma} (p_d^s)^{\sigma-1}$$

in order to solve for updated origin prices (for contraction mapping) in the market clearing equations

$$T_r^{Ak} v_r = \sum_{d=1}^R \alpha_k (\tau_d^r p^{Akr})^{1-\sigma} (p_d^{Ak})^{\sigma-1} X_d^{Ak} \quad (31)$$

$$= (p^{Akr})^{1-\sigma} \sum_{d=1}^R \alpha_k (\tau_d^r)^{1-\sigma} (p_d^{Ak})^{\sigma-1} X_d^{Ak}. \quad (32)$$

Solving for p^{Akr} ,

$$p^{Akr} = \left(\frac{T_r^{Ak} v_r}{\alpha_k} \right)^{\frac{1}{1-\sigma}} \left(\sum_{d=1}^R (\tau_d^r)^{1-\sigma} (p_d^{Ak})^{\sigma-1} X_d^{Ak} \right)^{\frac{1}{\sigma-1}}. \quad (33)$$

Similarly, for the non-agricultural sector,

$$p^{NAr} = (\bar{y}_r^L (1 - s_r^A) L_r)^{\frac{1}{1-\sigma}} \left(\sum_{d=1}^R (\tau_d^r)^{1-\sigma} (p_d^{NA})^{\sigma-1} X_d^{NA} \right)^{\frac{1}{\sigma-1}}. \quad (34)$$

To back out the Z_r^s , first solve for w_r^A and w_r^{NA} ,

$$w_r^A = \lambda^{-1} \bar{y}_r^L (s_r^A)^{1/\chi}, \quad w_r^{NA} = \lambda^{-1} \bar{y}_r^L (1 - s_r^A)^{1/\chi} \quad (35)$$

Then

$$Z_r^{Ak} = \frac{1}{p_r^{Ak}} \left(\frac{w_r^A}{1 - \alpha_k} \right)^{1-\alpha_k} \left(\frac{v_r^A}{\alpha_k} \right)^{\alpha_k} \quad (36)$$

$$Z_r^{NA} = \frac{1}{p_r^{NA}} w_r^{NA} \quad (37)$$

Non-agricultural value added, wage, and amenity. With the non-agricultural technology $Y_r^{NA} = Z_r^{NA}(H_r)^{1-\gamma_{NA}}(L_r^{NA})^{\gamma_{NA}}$, non-agricultural value added exceeds non-agricultural labor income by the factor $1/\gamma_{NA}$: $p^{NAr}Y_r^{NA} = \bar{y}_r^{LNA}(1-s_r^A)L_r/\gamma_{NA}$. We use this the other way around in the inversion. The observed non-agricultural value added $p^{NAr}Y_r^{NA}$ pins per-worker non-agricultural labor income, and hence the non-agricultural wage,

$$\bar{y}_r^{LNA} = \frac{\gamma_{NA} p^{NAr} Y_r^{NA}}{(1-s_r^A)L_r}, \quad w_r^{NA} = \frac{\bar{y}_r^{LNA}}{\mathbb{E}[z_{ir}^{NA} \mid \text{non-ag}]}, \quad (38)$$

so that the model reproduces the observed spatial distribution of non-agricultural income. The occupational-choice share then identifies the sectoral amenity as a residual: the observed s_r^A inverts to the amenity-adjusted return $\tilde{w}_r$, and with $G_r^A \equiv 1$,

$$G_r^{NA} = \frac{w_r^A}{\tilde{w}_r w_r^{NA}}. \quad (39)$$

The complementary fixed factor earns $(1-\gamma_{NA})p^{NAr}Y_r^{NA}$; these rents accrue to local residents,²⁸ so local household expenditure (29) is augmented to

$$E_r = v_r T_r^A + \bar{y}_r^L L_r + (1-\gamma_{NA})p^{NAr}Y_r^{NA} - p_r I_r^D. \quad (40)$$

Setting $\gamma_{NA} = 1$ and $G_r^{NA} = 1$ recovers the pure-selection model, in which the non-agricultural wage is instead forced off the agricultural wage through the inverted ratio $\tilde{w}_r = w_r^A/w_r^{NA}$ and there are no non-agricultural fixed-factor rents.

The migration value with a sectoral amenity. The amenity is a consumption-equivalent taste shifter: a worker in sector s with income $y_i^s = w_r^s z_{ir}^s$ has indirect utility $u_i^s = \frac{1}{\eta} (G_r^s y_i^s / p_r)^\eta - \nu \log(p_r^A / p_r^{NA}) + \nu$, so choosing the higher $G_r^s w_r^s z_{ir}^s$ is exactly utility-maximizing and amenity G_r^s is worth the same as income scaled by G_r^s .

²⁸We adopt local ownership of the fixed factor, as in Redding (2016): the rents enter the income of the region that earns them, so they feed back into local demand and welfare. In the taxonomy of Redding and Rossi-Hansberg (2017), the alternatives are a national portfolio that pools all fixed-factor rents and pays an equal per-capita dividend (as in Desmet and Rossi-Hansberg 2014) and absentee ownership that removes the rents from the general-equilibrium income effect altogether. We verified that a national portfolio leaves the leakage results essentially unchanged; local ownership is the more conservative choice for the welfare interpretation because it keeps the rents spatially concentrated.

Regional welfare is the expected maximum,

$$V_r = \frac{1}{\eta p_r^\eta} \mathbb{E} \left[\left(\max_s G_r^s w_r^s z_{ir}^s \right)^\eta \right] - \nu \log \frac{p_r^A}{p_r^{NA}} + \nu, \quad (41)$$

where the expectation is over the Frank–Fréchet draws (its closed form needs the same one-dimensional quadrature as the labor supplies). Consistent with the representative-agent evaluation used elsewhere for V_r , we evaluate the expected maximum at the mean selected incomes, which yields a region multiplier on consumption-equivalent income

$$\Omega_r = \frac{s_r^A \bar{y}_r^{LA} + (1 - s_r^A) G_r^{NA} \bar{y}_r^{LNA}}{s_r^A \bar{y}_r^{LA} + (1 - s_r^A) \bar{y}_r^{LNA}}, \quad V_r = \frac{1}{\eta} \left(\frac{\Omega_r e_r}{p_r} \right)^\eta - \nu \log \frac{p_r^A}{p_r^{NA}} + \nu, \quad (42)$$

with $e_r = E_r/L_r$ inheriting the local rents of (40). As an income-weighted average of 1 (agricultural workers) and G_r^{NA} (non-agricultural workers), Ω_r lies between 1 and G_r^{NA} . Where $G_r^{NA} < 1$ (a compensating agricultural amenity, about three-quarters of micro-regions in our estimates) $\Omega_r < 1$ discounts welfare; where $G_r^{NA} > 1$ (a net non-agricultural amenity, the remaining quarter) $\Omega_r > 1$. Because G_r^{NA} now carries the sectoral taste, the inverted residential amenity $B_r = \text{destFE}_r^{1/\epsilon}/V_r$ no longer absorbs it; setting $G_r^{NA} = 1$ gives $\Omega_r = 1$ and returns (42) to the pure-selection migration value.

Aggregate market clearing. Because of market clearing and the fact that the trade shares from all origins add up to 1 for each destination, it follows that

$$\sum_{r=1}^R L_r (1 - s_r^A) y_r^{LNA} = \sum_{r=1}^R \sum_{d=1}^R \pi_d^{NAr} X_d^{NA} = \sum_{d=1}^R X_d^{NA} \quad (43)$$

$$= \sum_{d=1}^R \left[(1 - \phi) p_d Y_d - \nu E_d^{1-\eta} \left(\frac{p_d^A}{p_d^{NA}} \right) (p_d L_d)^\eta \frac{\mathbb{E}((y_{id})^{1-\eta})}{(\mathbb{E}(y_{id}))^{1-\eta}} \right] \quad (44)$$

And if we decompose region-sector consumer price indices p_d^s as the product of a Brazil-wide sectoral price index $p^s \equiv \left(\sum_{o=1}^R (p^{so})^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$ and a region specific component $\tilde{p}_d^s \equiv p_d^s/p^s$, then we can solve for a Brazil-wide relative price index of agricul-

ture:

$$\frac{p^A}{p^{NA}} = \left(\left[\sum_{d=1}^R (1-\phi) p_d Y_d - \sum_{r=1}^R L_r (1-s_r^A) y_r^{LNA} \right] / \left[\sum_{d=1}^R \nu E_d^{1-\eta} \left(\frac{\tilde{p}_d^A}{\tilde{p}_d^{NA}} \right) (p_d L_d)^\eta \frac{\mathbb{E}((y_{id})^{1-\eta})}{(\mathbb{E}(y_{id}))^{1-\eta}} \right] \right) \quad (45)$$

Aggregate agricultural expenditure share in each region The share of consumer expenditure in agriculture of a household with expenditure e_i equals

$$\vartheta_{ir}^A = \phi + \nu \left(e_{ir} \right)^{-\eta} p_r^\eta. \quad (46)$$

Thus, the total agricultural expenditure in region r equals

$$X_r^A = \phi p_r I_r^D + \int \vartheta_{ir}^A e_{ir} dG(i) = \phi (E_r + p_r I_r^D) + \nu p_r^\eta \int (e_{ir})^{1-\eta} dG(i),$$

Assume that people with higher labor incomes get proportionately higher land rents and also spend proportionately more on deforestation. Define the ratio between land rents and labor income from the aggregates, a_r , and the ratio of the total income spent on deforestation investments, b_r .

$$a_r \equiv \frac{v_r T_r^A}{\int_i y_{ir}^L dG(i)}, \quad b_r \equiv \frac{p_r I_r^D}{\int_i y_{ir}^L dG(i) + v_r T_r^A}.$$

Then the assumption of proportional land rents and deforestation expenditures can be expressed formally as:

$$(1) y_{ir}^T = a_r y_{ir}^L, \quad (2) x_{ir}^d = b_r y_{ir} \implies (3) \frac{e_{ir}}{y_{ir}^L} = (1 + a_r)(1 - b_r) \equiv m_r$$

$$e_{ir} = \underbrace{\max\{w_r^{NA} z_{ir}^{NA}, w_r^A z_{ir}^A\}}_{y_{ir}^L} m_r,$$

where

$$m_r = \frac{E_r}{L_r \bar{y}_r^L}, \quad \bar{y}_r^L = s_r^A y_r^{LA} + s_r^{NA} y_r^{LNA} = \mathbb{E}(y_{ir}^L)$$

For the second term of the total agricultural expenditure equation, we compute the following integral

$$\begin{aligned} \int (e_{ir})^{1-\eta} dG(i) &= (m_r)^{1-\eta} L_r \mathbb{E}((y_{ir}^L)^{1-\eta}) \\ \mathbb{E}((y_{ir}^L)^{1-\eta}) &= \left(\int_0^\infty \int_0^{z^A w_r^A / w_r^{NA}} (w_r^A z^A)^{1-\eta} f(z^A, z^{NA}) dz^{NA} dz^A \right. \\ &\quad \left. + \int_0^\infty \int_{z^A w_r^A / w_r^{NA}}^\infty (w_r^{NA} z^{NA})^{1-\eta} f(z^A, z^{NA}) dz^{NA} dz^A \right) \end{aligned}$$

Thus, the aggregate local share of consumer expenditure in agriculture equals

$$\vartheta_r^A = \phi + \nu \left(\frac{E_r}{p_r L_r} \right)^{-\eta} \frac{\mathbb{E}((y_{ir})^{1-\eta})}{(\mathbb{E}(y_{ir}))^{1-\eta}}. \quad (47)$$

C.4 Assumptions behind the inversion

The inversion above is a sequence of closed-form back-outs interrupted by a single fixed point. It is useful to make explicit the structural assumptions that each main step relies on, both to clarify what is and is not disciplined by the data and to delimit the scope of the international-trade extension in Appendix E.

Cobb–Douglas agriculture with constant returns and perfect competition.

The recovery of land rents and revenues in equations (20)–(21) rests on agricultural production being Cobb–Douglas in land and labor with constant returns, supplied under perfect competition. Zero profits then make land’s share of crop revenue exactly α_k , so land payments equal α_k times observed revenue and the regional land rent v_r is identified from observed land and revenue *without any goods prices*. The same zero-profit, constant-returns logic underlies the GDP-accounting identity (22), in which total revenue is exhausted by payments to the two factors and average labor income $\bar{y}_r^L$ is read off directly. This is the load-bearing assumption: it is what lets the entire factor side—land rents, wages, incomes, land prices and deforestation investment—be computed in closed form, prior to and independently of the price fixed point.

Labor heterogeneity with self-selection. The mapping from the observed agricultural employment share s_r^A to sectoral wages (equation (35)) requires that workers

draw a pair of sectoral productivities from the Frank copula and sort into the higher-paying sector. This generates a monotone relationship between s_r^A and the wage ratio w_r^A/w_r^{NA} that the inversion simply inverts; without a parametric joint distribution, an employment share alone would not pin down relative wages.

Land as a capitalized asset and deforestation as investment. Equations (27)–(28) treat the price of land as the present value of its rental stream, $q_r = v_r/[1 - \beta(1 - \rho)]$, and deforestation as an investment technology with returns-to-scale parameter δ , so that $p_r I_r^D = \delta q_r T_r^D$. These are what make δ the structural land-supply elasticity and let the deforestation TFP be recovered from the investment first-order condition.

Non-homothetic (PIGL) preferences. The agricultural expenditure share (46) is income-dependent through the Engel elasticity η ; aggregating over the within-region income distribution generates the correction term λ_r . This is the channel through which falling real income raises the food share—the source of the model’s structural-transformation content, and the reason the demand side responds non-trivially to any closure that moves regional expenditure.

Armington/CES trade with iceberg costs. The trade shares and price indices (equations (8)–(24)) treat each region’s output of each good as a distinct variety, combined with elasticity σ across regions and θ across crops, shipped subject to iceberg costs τ_d^o . This is the *only* block that requires goods prices, and hence the only place a fixed point is needed: market clearing for each origin pins the vector of origin prices, from which the productivities Z follow by zero-profit pricing.

Region-level balanced trade. Finally, household expenditure (29) equals local factor income net of deforestation investment, $E_r = v_r T_r^A + \bar{y}_r^L L_r - p_r I_r^D$. Combined with the accounting identity that income equals the value of local production, this imposes that each micro-region’s total absorption (consumption plus investment) equals the value of its output: every region’s trade is *balanced in aggregate value*, with no cross-region asset market through which a region could run a current-account surplus or deficit. We emphasize the scope of this restriction, because it is narrower than it may first appear. It constrains only the *total* value a region trades, not its

sectoral composition: a region may freely run an agricultural deficit financed by a non-agricultural surplus, or the reverse. Equivalently, $X_r^A \neq p_r^A Y_r^A$ and $X_r^{NA} \neq p_r^{NA} Y_r^{NA}$ are both permitted—a city specializing in non-agricultural production ships it to the rest of Brazil and imports its food, while an agricultural-frontier region does the opposite—so cross-region specialization and structural-transformation geography are fully present in the closed economy. The single thing balanced trade rules out is a region producing more total value than it absorbs. This is the standard closure for a static spatial model without an asset market, and it is precisely the margin that the open-economy block in Appendix E relaxes at the national level, where the relevant aggregate (Brazil’s export share of output) is reliably measured.

D Parameter estimation

This appendix details the estimation and calibration of the model’s structural parameters. We begin with the deforestation elasticity δ and its robustness battery, and then describe the production, preference, migration, and trade parameters, followed by the regrowth rate and the model fit.

D.1 Deforestation elasticity δ

This subsection details the estimation of the deforestation elasticity that governs the returns to scale of the deforestation technology, δ , and documents its robustness. Throughout, the object we estimate is the *supply elasticity of new agricultural land*—the elasticity of deforestation to agricultural land rents—which in the model equals $\frac{\delta}{1-\delta}$. We recover it with a shift-share (Bartik) instrument and then map it to the structural parameter via

$$\delta = \frac{b}{1+b}, \quad b \equiv \widehat{\frac{\delta}{1-\delta}}. \quad (48)$$

D.1.1 Estimating equation and instrument

As derived in Section 5.4 (equation (10)), the optimal deforestation policy implies the log-linear relationship

$$\log(T_{rt}^D) = \alpha_r + \alpha_t + \underbrace{\frac{\delta}{1-\delta}}_b \log v_{rt} + \theta \log T_{r,t-1}^F + \varepsilon_{rt}, \quad (49)$$

where T_{rt}^D is deforestation (gross natural-vegetation loss) in municipality r and year t , v_{rt} is the model-implied agricultural land-rent index, and α_r, α_t are municipality and year fixed effects. The standard specification conditions on $\log T_{r,t-1}^F$, the lagged remaining stock of natural vegetation—a predetermined *depletion control*, not an instrument. It absorbs the frontier-depletion margin (a municipality with little forest left has a mechanically weak rent-clearing response), so b is the rent response *holding forest availability fixed*—the reduced-form counterpart of the model’s forest-availability channel (the cap $T^D \leq T^{\max}$ and $Z^D \propto (T^{\max})^\psi$). With the depletion margin modelled, the estimation sample is *all* Brazilian municipalities, with the Legal-Amazon frontier subsample reported alongside throughout. Land rents v_{rt} are endogenous to local deforestation shocks, so we construct a shift-share instrument that isolates demand-driven variation in land rents coming from world commodity prices interacted with a municipality’s pre-determined crop mix.

For each municipality we freeze the base-year exposure shares $s_{r,c,0}$ —the share of agricultural area in crop c in the first sample year—and combine them with time-varying world “pink-sheet” commodity prices p_{ct} (real, deflated) to form

$$Z_{rt} = \sum_c s_{r,c,0} \log p_{ct}, \quad (50)$$

covering 19 priced commodities plus a residual “missing-sector” share with price zero. We instrument $\log v_{rt}$ with Z_{rt} in equation (49), weighting by agricultural area and clustering by municipality. The land-rent index $v_{rt} = \sum_c \alpha_c \text{rev}_{rc} / \sum_c \text{area}_{rc}$ uses national crop-level land-rent weights α_c recovered from the 2006 Agricultural Census (see the production parameters in Section D.3).

D.1.2 Baseline estimate

The baseline specification, for all Brazilian municipalities over 2008–2018 with the depletion control, is reported in the main text (Table 3). The conventional first stage is strong (the price-exposure index raises land rents with the expected sign; unit-clustered first-stage $F \approx 148$; but see the inference discussion below), the depletion control enters strongly with the expected sign (more remaining forest, more clearing), and the OLS coefficient is sharply attenuated relative to IV, the pattern expected when the land-rent regressor is measured with error. The IV land-supply elasticity on

the pooled sample is $b \approx 0.68$ (s.e. 0.28), implying a structural $\delta = 0.68/1.68 \approx 0.41$ (s.e. 0.10). On the Legal-Amazon frontier subsample—where the land margin is active—the same specification gives $b \approx 1.66$ (s.e. 0.43), i.e.

$$\delta = \frac{1.66}{1 + 1.66} \approx 0.62,$$

essentially the pre-control frontier estimate (0.61); at the micro-region level the pooled estimate is 0.48 (s.e. 0.14). The pooled all-Brazil estimate is the value at which the model is calibrated ($\delta = 0.41$, matching the inversion); the frontier estimate marks the high end of the identified range, and because leakage is increasing in δ , the sweep’s frontier cells bound the headline conclusions from above (the full-re-inversion sweep of Table 20 spans the identified range).

D.1.3 Robustness

Table 12 collects the robustness battery at both spatial levels (municipality and micro-region), reporting the structural $\delta = \hat{\beta}_{IV}/(1 + \hat{\beta}_{IV})$ throughout; every row deviates from the all-Brazil, depletion-controlled baseline in one dimension. We organize it to parallel the appendix of the companion suitability-based study, which subjects the same economic object to the analogous checks.

Sample start year (Panel A). Because the Bartik exposure shares are frozen at the first sample year, the start year is a natural source of researcher discretion. On the pooled sample the start year matters more than it did on the frontier: δ is 0.56 (2003 start), 0.21 (2005), and 0.41 (2008, baseline) at the municipality level, with the micro-region level shifted upward (0.64/0.34/0.48). The 2005 start is the outlier in both; the frontier subsample, by contrast, stays in a tight 0.5–0.6 band across start years. Figure 16 plots the path. The baseline 2008 start coincides with the onset of the priority-list enforcement regime.

Price persistence and dynamics (Panel B). The deforestation response to prices is dynamic: clearing decisions respond to *persistent* rather than transitory price movements (as in Araujo et al. 2026). Because the calibrated model is static, the relevant object is the response to a price change perceived as permanent. We isolate it by replacing annual prices in (50) with 3- and 5-year moving averages:

δ rises monotonically from 0.41 (annual) to 0.58 (3-year) to 0.70 (5-year) on the pooled sample ($0.48 \rightarrow 0.63 \rightarrow 0.75$ at the micro-region level). The permanent-shock elasticity is thus larger than the annual one—consistent with short-run $<$ long-run responses—and is the value most appropriate for a static, steady-state interpretation of the model.

Exposure weighting (Panel C). The area-share design is dominated by a single shock: pasture (beef) carries $\approx 83\%$ of the shock-importance mass. Re-weighting exposure by output tonnage spreads the mass across commodities while raising δ to 0.56; a micro-region leave-one-out exposure design (which purges a municipality’s own shares; cf. [Adão et al. 2019](#)) gives $\delta \approx 0.52$; revenue weighting yields 0.24 with a weak first stage and is not preferred. The fitted-share (Fiszbein) row is discussed in its own subsection below. Across the designs with first-stage power, δ stays within the identified range discussed below.

Sample (Panel D). The all-Brazil baseline gives $\delta = 0.41$; the Legal-Amazon frontier subsample gives 0.62. The gap is economically sensible: outside the Amazon most native vegetation has already been cleared, so the rent-clearing margin is thinner even conditional on the depletion control—indeed, estimated *separately* outside the Legal Amazon the elasticity is wrong-signed and unstable, a symptom of both the post-frontier margin and an Amazon-calibrated land-rent index. In the model, these regional differences arise endogenously through forest availability in the deforestation block (the cap $T^D \leq T^{\max}$ and the channel $Z^D \propto (T^{\max})^\psi$) rather than a region-specific δ .

Depletion control (Panel E). Removing the lagged remaining-forest control barely moves the estimates (0.41 vs. 0.38 pooled; 0.62 vs. 0.61 on the frontier), so the control is not doing mechanical work; its role is to license the pooled sample by absorbing the post-frontier depletion margin, and its strongly positive coefficient is itself evidence for the forest-availability mechanism the model builds in through ψ .

D.1.4 Identification and inference with few shocks


Two distinct assumptions can license a shift-share design, and they have different implications for inference ([Goldsmith-Pinkham et al. 2020](#), [Borusyak et al. 2022](#)). Our

Table 12: Robustness of the deforestation elasticity $\delta = \hat{\beta}_{IV}/(1 + \hat{\beta}_{IV})$

Panel	Specification	$\delta = \hat{\beta}_{IV}/(1 + \hat{\beta}_{IV})$		first-stage F	
		muni	micro	muni	micro
Baseline	all Brazil 2008-2018	0.41	0.48	148	31
A. Start year	2003-2018	0.56	0.64	201	44
	2005-2018	0.21	0.34	182	42
	2008-2018 (baseline)	0.41	0.48	148	31
B. Price persistence	annual (baseline)	0.41	0.48	148	31
	3-yr MA	0.58	0.63	98	22
	5-yr MA	0.70	0.75	78	16
C. Exposure weighting	area (baseline)	0.41	0.48	148	31
	tons	0.56	0.67	94	33
	revenue	0.24	0.39	30	9
	area, leave-one-out	0.52		85	
	fitted (Fiszbein) shares	0.63	0.56	2	2
D. Sample	Legal Amazon	0.62	0.38	104	18
	all Brazil (baseline)	0.41	0.48	148	31
E. Remaining-forest control	with lag control (baseline)	0.41	0.48	148	31
	without control	0.38	0.46	146	31

Note: Each row reports the structural $\delta = \hat{\beta}_{IV}/(1 + \hat{\beta}_{IV})$ where $\hat{\beta}_{IV} = \delta/\widehat{(1 - \delta)}$ is the IV land-supply elasticity from equation (49), at the municipality (“muni”, unit-clustered) and micro-region (“micro”, micro-clustered) level. The baseline is all Brazilian municipalities 2008–2018 (area-share exposure, log outcome, area weights, lagged log remaining-forest control); each row deviates in one dimension. The model is calibrated at the pooled all-Brazil estimate $\delta = 0.41$; the frontier (Legal-Amazon) estimate 0.62 is the high end of the identified range (leakage is increasing in δ , so the frontier cells of the sweep bound the headline from above). **Zeros.** The log outcome drops municipality–years with zero recorded forest loss; these matter mainly for post-frontier municipalities with little remaining forest. A Poisson (PPML) estimator that retains the zeros yields a comparable frontier elasticity. Estimated separately outside the Legal Amazon, the elasticity is wrong-signed and unstable, consistent with those regions being post-frontier (and with an Amazon-calibrated land-rent index); in the model, regional differences in deforestation arise endogenously through forest availability (the cap $T^D \leq T^{\max}$ and the deforestation–TFP channel $Z^D \propto (T^{\max})^\psi$) rather than a region-specific δ . PPML and region-specific specifications are reported as robustness.

Figure 16: δ is stable across sample start years



Note: Structural $\delta = \hat{\beta}_{IV} / (1 + \hat{\beta}_{IV})$ by sample start year (end year fixed at 2018), all-Brazil baseline with the depletion control, municipality and micro-region levels. Shaded band is the 95% confidence interval (delta method).

primary reading is the share-exogeneity (difference-in-differences) license of [Goldsmith-Pinkham et al. \(2020\)](#): conditional on municipality and year fixed effects, the base-year crop mix is as-good-as-random with respect to the de-trended deforestation error, the realized price path is conditioned on, and the effective sample is the cross-section of exposure units. Under that license the unit-clustered standard errors and first-stage F statistics reported throughout are the appropriate objects, and the fitted-share exercise below directly tests the license’s key threat.

We disclose, however, that the design rests on few shocks, and that inference which treats the shocks as the random object is far weaker. Pasture (beef) carries roughly 83% of the Rotemberg (shock-importance) weight ([Goldsmith-Pinkham et al. 2020](#)), and the crop-level effective number of shocks is ≈ 1.4 ; no exposure re-weighting raises it materially, because the number of deforestation-relevant commodities is intrinsically small—frontier clearing is overwhelmingly a cattle-margin phenomenon, and commodity-price booms are its well-documented proximate driver ([Da Mata and Dotta 2021](#), whose exposure design independently corroborates the relevance of price shocks for clearing and fires). Following [Adão et al. \(2019\)](#) and [Borusyak et al. \(2022\)](#), exposure-robust (shock-clustered) inference on the frontier design yields a first-stage $F \approx 1.5$, re-clustering the conventional first stage by year collapses it from ≈ 133 to ≈ 8 , and the 95% Anderson–Rubin confidence set for δ is unbounded. Read con-

structively, these diagnostics are the case *for* the share-exogeneity license rather than merely a caveat to it: a shock-exogeneity (BHJ) design requires many quasi-random shocks for its asymptotics to have content, and with effectively one-and-a-half shocks that route is simply unavailable in this setting. Treating the shares as the identifying variation is therefore not a modeling convenience but the only coherent reading of the design—which is why we defend its key assumption directly, with the fitted-share exercise below, instead of leaning on shock-level asymptotics. The residual-dependence critique of [Adão et al. \(2019\)](#) still applies *regardless* of the license one adopts: any aggregate disturbance that moves deforestation differentially by crop mix makes errors correlated across municipalities with similar exposures, so the unit-clustered precision overstates what the price variation alone can deliver.

We therefore use the shift-share for what it can credibly do—establish the sign and magnitude range of the land-supply elasticity from a transparent demand shifter—and do not rest any headline claim on its reduced-form precision. Point identification of the calibrated δ is anchored instead on the model inversion (which reproduces observed deforestation exactly at the calibrated value), on the companion suitability-based design (a different instrument yielding 0.56–0.67), and on the corroborating pattern across estimators (dynamic-panel estimates on the frontier are as low as 0.07–0.12, reflecting the depletion feedback the model captures through ψ); the counterfactual sweep of Table 20 then spans the full identified range, and the paper’s qualitative conclusions are invariant across it.

D.1.5 Exposure-share exogeneity: fitted (Fiszbein) shares

The main threat to the share-exogeneity license is specific: on the Amazon frontier, a high initial *pasture* share—the share that carries the shock-importance mass—is close to “frontier status,” which is directly correlated with deforestation trends. To test this margin we replace the realized base-year shares $s_{r,c,0}$ in (50) with *fitted* shares $\hat{s}_{r,c,0} = \mathbb{E}[s_{r,c,0} | \text{fundamentals}]$, predicted from pre-determined agro-climatic and market-access fundamentals only—GAEZ potential yields and pasture suitability, pre-sample climate means, and distances to ports—in the spirit of [Fiszbein \(2022\)](#)’s predicted specialization (equivalently, the $\hat{s}$ serve as exogenous share instruments in the sense of [Goldsmith-Pinkham et al. 2020](#)). Identification then comes from price shocks interacted with the *predictable-from-geography* part of the crop mix, purg-

ing the endogenous frontier component. Geography predicts the crucial share well (weighted R^2 for the pasture share ≈ 0.5 on the frontier), and the Rotemberg weights of the fitted design remain concentrated on beef—identification still flows through pasture, but now through its geography-predictable component. On the frontier sample, where the design has first-stage power ($F \approx 24$), the fitted-share instrument reproduces the frontier estimate: $\delta = 0.71$ (s.e. 0.06) against the realized-share 0.62 (s.e. 0.06; both with the depletion control), inside the fitted estimate’s 95% confidence interval—so the endogenous component of the exposure shares is not what identifies δ . On the pooled all-Brazil sample the fitted instrument loses its first stage ($F \approx 2$): geography-predicted crop mix interacted with prices moves land rents only where the land margin is active, which is itself consistent with the frontier interpretation; the (weakly identified) pooled point estimate, 0.63 (s.e. 0.24), sits between the pooled baseline 0.41 and the frontier value. This exercise defends the *share* margin of the design; it cannot, by construction, relax the few-shock inference problem above, since predicting the shares from geography creates no new independent price variation.

D.1.6 Relation to the literature

Our calibrated value, $\delta = 0.41$ (the pooled all-Brazil estimate; land-supply elasticity $b \approx 0.7$, with 0.62 on the Legal-Amazon frontier and 0.70–0.75 under the permanent-price design), is consistent with the structural land-use literature once the estimands are aligned. [Farrokhi et al. \(2025\)](#) calibrate a global deforestation elasticity of 0.9, and [Berman et al. \(2023\)](#) report comparable crop-price elasticities of pan-tropical deforestation; both bracket our flow elasticity. [Souza-Rodrigues \(2019\)](#) recovers extensive-margin deforestation responses from transport-cost variation, and [Araujo et al. \(2026\)](#) emphasize the dynamic short-run/long-run wedge that Panel B and the dynamic-panel estimates trace out.

A key subtlety is the *stock-versus-flow* distinction. The companion suitability-based study estimates the elasticity of the agricultural-land *stock* to returns and reports a mean of 0.56–0.67 across exposure designs; deforestation is the *flow*, a small fraction of the stock, so a modest stock elasticity implies a larger flow elasticity. Our flow δ and that paper’s stock elasticity therefore agree once placed on the same footing—and, reassuringly, the structural $\delta = \hat{\beta}_{IV}/(1 + \hat{\beta}_{IV})$ itself lands squarely inside the 0.56–0.67 band. The two strategies are complementary: the price-Bartik

design here uses a transparent world-price shock and a direct deforestation outcome, while the suitability design buys exogeneity of the exposure shares (in the spirit of [Farrokhi and Pellegrina 2023](#), [Domínguez-lino 2026](#)).

D.2 Returns to remaining forest: ψ

This subsection derives and reports the steady-state regression used to calibrate ψ , the parameter dictating how the productivity of aggregate deforestation declines as less area is left in unprotected natural ecosystems. While $\delta < 1$ (yearly decreasing returns to scale) makes it optimal to space deforestation out over time, $\psi > 0$ makes the cost of deforesting a given area higher when there is less forest left—either because the forest that is easier to cut is chosen first, or because the value of standing forest rises as it becomes scarcer. We calibrate ψ from the steady-state condition equating forest loss and regrowth, $T_{r,ss}^D = \rho T_{r,ss}^A$. Substituting the optimal level of deforestation (6) and $Z_{r,ss}^D = \bar{Z}_{r,ss}^D (T_{r,ss}^F)^\psi$ yields

$$\frac{T_{r,ss}^F}{T_{r,ss}^A} = \rho (\bar{Z}_{r,ss}^D)^{\frac{1}{\delta-1}} \delta^{\frac{\delta}{\delta-1}} \left(\frac{q_{r,ss}}{p_{r,ss}} \right)^{\frac{\delta}{\delta-1}} (T_r^F)^{1-\frac{\psi}{1-\delta}},$$

so that $1 - \psi/(1 - \delta)$ is the long-run scale elasticity of forest cover: how the total area in forest relates to the relative area in forest. Taking logs, we run

$$\log \left(\frac{T_{rt}^F}{T_{rt}^A} \right) = \alpha_{s(r)t} - \frac{\delta}{1-\delta} \log(v_{rt}) + \left(1 - \frac{\psi}{1-\delta} \right) \log(T_{rt}^F) + \epsilon_{rt}, \quad (51)$$

where $s(r)$ is the state of region r , and state-year fixed effects restrict attention to within-state-year differences; the relevant variation is cross-sectional, since we are interested in long-run differences across regions of different sizes. Since regions with high deforestation are less likely to have reached their steady state, column (3) of Table 13 considers exclusively regions where deforestation rates are below the regrowth rate. Given the calibrated $\delta = 0.41$, the forest-area coefficients imply $\psi = (1 - \hat{\beta}_{TF})(1 - \delta)$ between 0.26 and 0.27 across specifications; our point estimate is $\psi = 0.26$, from the steady-state sample of column (3). The regression's own land-rent coefficient would imply a smaller δ than the calibrated value—unsurprising, since the within-state-year design is not our preferred source of identification for δ —so we impose the calibrated δ throughout, which keeps ψ and δ jointly consistent with the inversion; the model-

based method (regressing inverted Z_{rt}^D on remaining forest in a log-log regression) delivers similar values.

Table 13: Estimation of ψ

Dependent Variable:	$\ln(T^F/T^A)$		
Model:	(1)	(2)	(3)
<i>Variables</i>			
log forest area T^F [= $1 - \psi/(1 - \delta)$]	0.5530*** (0.0137)	0.5431*** (0.0141)	0.5681*** (0.0144)
log land rent v [= $-\delta/(1 - \delta)$]		-0.1277*** (0.0219)	-0.0350** (0.0166)
<i>Fixed-effects</i>			
State-year	Yes	Yes	Yes
Low-deforestation sample	No	No	Yes
<i>Fit statistics</i>			
Observations	55,284	55,284	34,495
R ²	0.72474	0.73041	0.66746
Within R ²	0.54795	0.55727	0.45761
<i>Clustered (municipality) standard-errors in parentheses</i>			
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>			

Note: This table shows the results of a regression of the log of the ratio of land in forest over land in agriculture on the log of the total area in forest at the municipality-year level, following (51). The first coefficient estimate displayed is $1 - \frac{\psi}{1-\delta}$, where ψ governs the “returns to total unprotected forest left” and δ governs the returns to scale of the deforestation production function. All columns controls for state-year fixed effects and weigh municipalities by agricultural area. Column (1) does not have any additional controls. Columns (2) and (3) control for the log of the estimated rental rate of agricultural land. Column (3) restricts the sample to observations with less than 1 square kilometer of forest loss so that it is more plausibly looking at a steady state. Standard errors in parenthesis. *p<0.1, **p<0.05, *** p<0.01.

D.3 Production Parameters

To estimate α_k for each agricultural commodity k , we rely on the 2006 Agricultural Census, a census of the universe of farming establishments. Using data on land prices, land in each agricultural activity, labor in each agricultural activity, and agricultural wages, we estimate how the ratio of land income to labor income varies between

pastures for cattle grazing, temporary crops, and perennials.

Let the area planted in each type of agricultural activity T_r^{Ak} , workers in type of agricultural activity N_r^{Ak} , labor incomes per capita in agriculture $\bar{y}_r^{LA}$, and land value q_r .

The Cobb-Douglas functional form of the agricultural production function and the zero-profit condition together imply that

$$\frac{\alpha_k}{1 - \alpha_k} = \frac{\sum_{r=1}^R v_r T_r^{Ak}}{\sum_{r=1}^R \bar{y}_r^{LA} N_r^{Ak}}.$$

Doing this we find that the land share for pastures, temporary crops, and perennials vary significantly and they are, respectively: 0.71, 0.54, and 0.36. These differences make it so that there is greater substitutability between land and labor in agricultural production as a whole. This has important implications for leakage. Consider a region where the supply of agricultural workers increases. This could be due to more migration from regions where conservation policies have been enacted. Higher supply of agricultural workers increases the demand for agricultural land. In a model with multiple agricultural sectors, however, it might also shift land use away from very land intensive activities, for example cattle ranching, toward more labor intensive ones, such as perennial crops. The extent to which this happens in equilibrium would depend mainly on consumers' willingness to substitute consumption between goods and on the slope of the supply of deforestation. If deforestation is very cheap and consumers very reluctant to substitute toward less land-intensive goods, cattle-ranching will remain preferable and the inflow of workers will lead to much deforestation. By contrast, in a region where deforestation is very costly and consumers readily substitute beef with soy, the increased supply of labor will mean farmers may opt to convert pastures to soy fields and demand very little deforestation.

Non-agricultural labor share γ_{NA} . The non-agricultural sector has a complementary fixed factor (capital and structures) of share $1 - \gamma_{NA}$, the analogue of the agricultural land share α_k , so that non-agricultural value added equals non-agricultural labor income divided by γ_{NA} . We *calibrate* $\gamma_{NA} = 0.66$, the standard aggregate labor share of roughly two-thirds (Gollin 2002, Guerriero 2019). We do not take the labor share directly from the Brazilian accounts because the naive ratio is downward biased in a high-informality economy: the census non-agricultural wage bill (per-worker earn-

ings times employment) over regional-accounts non-agricultural value added is only 0.39, but this omits the labor income of the self-employed and informal workers, which the accounts book as operating surplus (capital). [Gollin \(2002\)](#) shows that adjusting for self-employment raises measured labor shares to around two-thirds and removes their spurious correlation with income per capita; [Guerriero \(2019\)](#) confirms adjusted shares near $2/3$ across a large panel of countries and industries. We therefore use $\gamma_{NA} = 0.66$ and retain the 0.39 census ratio only as an unadjusted cross-check. The share is a scalar held constant across regions and years (a region-by-region ratio would mechanically make model value added match the accounts and is therefore not used, and a constant factor share avoids an income-dependent fixed point). Because γ_{NA} is pinned outside the model it does not compete with any region-level fundamental; moreover, under local ownership of the fixed factor it nets out of the equilibrium (labor earns γ_{NA} and local capital $1 - \gamma_{NA}$ of the same value added, both accruing to residents), so the leakage results are invariant to its exact value and it matters only for the reported non-agricultural value-added level and for welfare under national ownership.

Identification of the sectoral amenity. The occupational-choice block adds one region-level fundamental, the non-agricultural amenity $G_{NA,r}$, and it is exactly identified alongside the non-agricultural productivity Z_r^{NA} . Per region, two moments pin two fundamentals: the observed non-agricultural value added pins Z_r^{NA} (through the non-agricultural wage, equation 38), and the observed agricultural employment share pins $G_{NA,r}$ as the residual of the amenity-adjusted selection (equation 39). The scalar copula parameters $(\chi_A, \chi_{NA}, \iota)$ remain identified by different moments—the within-region earnings dispersion across sectors—so keeping both the Roy selection structure and the amenity does not over-identify the system. The amenity must be a pure taste term entering the choice, not the wage or output; folding it into a region-specific Fréchet scale or shape would make it collinear with Z_r^{NA} .

D.4 Preference parameters

For PIGL parameters ϕ and ν , we take the values estimated by [Eckert and Peters \(2022\)](#). The parameter governing the non-homotheticity of preferences, η , also referred to as Engel elasticity, captures the rate at which higher-income consumers

shift their budget shares away from agricultural goods (i.e. food items). We estimated by regressing the logarithm of the share of expenditure in food on the logarithm of total expenditure using the household-level data from the 2017/2018 consumer expenditure survey (POF). The table below shows the coefficients from that regression. For robustness, we consider both food expenditure as a share of total non-durables expenditure (column 1 - the preferred specification) and as a share of total household income (column 2).

Table 14: Estimation of η


Dependent Variables: Model:	Log Food Share (of expenditure) (1)	Log Food Share (of income) (2)
<i>Variables</i>		
Log(Non-durable Expenditure)	-0.5065*** (0.0036)	
Log(Income)		-0.5746*** (0.0055)
<i>Fit statistics</i>		
Observations	45,322	45,322
R ²	0.30887	0.19458
Dependent variable mean	-1.7353	-3.1687
<i>IID standard-errors in parentheses</i>		
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		

Note: This table displays the estimated Engel curve for food: the relationship between income and budget share of food, which illustrates the non-homotheticity in consumer preferences. The dots represent the average share of total household expenditure dedicated to food items for 20 bins of total household expenditure. Standard errors in parenthesis. *p<0.1, **p<0.05, *** p<0.01.

D.5 Migration parameters

Recall that the utility of a migrant moving from origin o to destination d depends on: (i) consumption utility V_d , (ii) residential amenities B_d , (iii) an origin-destination migration cost μ_{od} , and (iv) an idiosyncratic preference shock with a Frechet distribution with shape parameter ϵ . These shocks are drawn every year when making

Figure 17: Engel curve for food consumption



Note: This figure displays the estimated Engel curve for food: the relationship between income and budget share of food, which illustrates the non-homotheticity in consumer preferences. The dots represent the average share of total household expenditure dedicated to food items for 20 bins of total household expenditure.

migration decisions so that each year, the share of migrants from o that pick d as destination equals

$$\text{Share of Migrants}_d^o = \frac{(V_d B_d \mu_d^o)^\epsilon}{\sum_j (V_j B_j \mu_j^o)^\epsilon}.$$

The above is referred to as the migration gravity equation. To estimate it, first we take logs so that it becomes the following equation with origin and destination fixed effects

$$\log(\text{Share of Migrants}_d^o) = \delta_o + \delta_d + \epsilon \log(\mu_d^o). \quad (52)$$

Using bilateral migration flows between the 557 micro-regions according to the 2010 census, we use the equation above to estimate $\epsilon \mu_d^o$ and the correlation of (log) income and destination fixed effects to estimate ϵ . Our preferred approach, used in the calibration, estimates the full matrix of bilateral migration utilities μ_d^o from the residuals of the regression of (log) migration rates on origin and destination fixed effects. A more parsimonious alternative models μ_d^o as a function of the linear distance between micro-regions, $\mu_d^o = (\text{dist}_d^o + 1)^v$, estimated from a regression of migration shares on

log distance plus 1 with origin and destination fixed effects, as reported in table 15, column (1). We add one so that $\mu_o^o = 1$ and $\mu_d^o(\text{dist}_d^o)$ is a decreasing function of distance as long as $v < 0$. The distance elasticity of migration is estimated to be -1.48.

To estimate ϵ , note that the destination fixed effect, according to the structure of the model, depends on income at destination and amenities as follows,

$$\delta_d = \epsilon \log(V_d) + \epsilon \log(B_d). \quad (53)$$

And $\log(V_d)$, up to an approximation of equation 7, is $\log(V_d) \approx \eta \log(e_d/p_d) - \log(\eta)$. Thus, following Bugge et al. (2023), we estimate ϵ from a regression of the fixed effect on the estimated fixed effects on the logarithm of income in the 2010 census. The income used in this regression is the income of those people who were in-migrants to destination d . Although average income and real expenditure (e) are not the same thing, as long as they are proportional, a regression of destination fixed effects on income can help us estimate ϵ . Column (2) of table 15 shows the results of this PPML regression, where we obtain $\epsilon\eta = 0.615$, which implies $\epsilon = 1.22$ for the estimated $\eta = 0.506$. Two concerns with this method are (i) endogeneity: income may be correlated with amenities, which would bias the estimation, and (ii) measurement error: we use income instead of expenditure and we do not divide incomes by their regional price index, which is unobserved. A panel of migration flows from older and newer census data with micro-region and year fixed effects, as in Bugge et al. (2023), would address both; we instead assess how much the calibrated ϵ can matter through the parameter sweep of Appendix Table 20, where varying the migration block—up to shutting internal migration down entirely—moves cumulative leakage by at most 0.1 percentage points, so the leakage results are insensitive to the residual uncertainty in ϵ .

D.6 Trade parameters

Iceberg trade costs are estimated to fit trade costs in 1999 and 2017 from the residuals in the regression equation of (log) trade flows with origin and destination fixed effects. Taking logarithms from equation 8, and sending the volume of expenditure of consumers in region d to the right hand side so that the data on trade flow rather

Table 15: Estimation of migration parameters

Dependent Variables:	Share of Migrants	Destination F.E.	
Model:	(1)	(2)	(3)
	Poisson	OLS	OLS
<i>Variables</i>			
log(Distance+1)	-1.483*** (0.0071)		
log(income)		0.6153*** (0.0345)	0.4905*** (0.0320)
<i>Fixed-effects</i>			
dest	Yes		
orig	Yes		
<i>Fit statistics</i>			
Observations	311,364	558	557
Pseudo R ²	0.80453	0.36162	0.28347

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: This table displays the results of the regressions used for the estimation of migration parameters. Column (1) is a Poisson Pseudo Maximum Likelihood (PPML) regression of bilateral migration shares in the past 5 years at the micro-region level on the log of their distance plus 1. Migration data comes from the 2010 population census. The regression includes origin and destination fixed effects. Column (2) is also a PPML regression. It has as dependent variable the destination fixed effect estimated from a PPML as in column (1) but without the fixed effects only, and not the log of distance plus 1, as specified in (52). The coefficient displayed is the coefficient on the log of the average income of migrants according to the 2010 population census. Column (3) is as column (2) but considers the average income of all people living in the destination micro-region. Standard errors in parenthesis. *p<0.1, **p<0.05, *** p<0.01.

than shares can be used, yields

$$\log X_d^o = \log p^o - \log p_d X_d + (1 - \sigma) \log(\tau_d^o)$$

Notice that these data do not include the observations for which the origin and the destination are the same, as it is data from inter-state customs. Thus, if we assume that $\tau_d^o = (\text{dist}_{od} + 1)^\kappa$, then the coefficient of the log-log regression of trade flows on distance (+1) with origin and destination fixed effects equals $\kappa(1 - \sigma)$. Given the results from the regressions reported in the table below, $\kappa \approx \frac{-0.9}{1-9.6} \approx 0.105$ (at $\sigma = 9.6$; the gravity coefficient itself is independent of σ , so κ is simply re-derived when σ changes).

Table 16: Estimation of trade parameters

Dependent Variable:	Bilateral trade flows	
Year:	1999	2017
Model:	(1)	(2)
<i>Variables</i>		
log(Travel cost)	-0.8831*** (0.0949)	-0.9217*** (0.0551)
<i>Fixed-effects</i>		
state_origin	Yes	Yes
state_destination	Yes	Yes
<i>Fit statistics</i>		
Observations	702	702
Pseudo R ²	0.96789	0.95761
<i>Clustered (state of origin) standard-errors in parentheses</i>		
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		

Note: This table displays the results of the regressions used for the estimation of trade parameter κ that governs the distance-elasticity of trade. They are OLS regressions of the bilateral trade flows between different states on the log of their distance plus with origin and destination fixed effects. The structural interpretation of the coefficient is that it equals $(1 - \sigma)\kappa$. Column (1) uses 1999 state-level trade flow data and column (2) uses 2017 trade flow data. Standard errors in parenthesis. *p<0.1, **p<0.05, *** p<0.01.

We use this same distance elasticity κ for the international iceberg cost to the rest of the world, applying it to each micro-region's distance to its relevant export port

(Appendix E).

D.7 Regrowth rate

The distribution of regrowth rates is shown in Figure 18, panel (a) for four different years in the model period. We can see that the mode of the regrowth rate distribution is around 0.3%, and there is a long right tail, with the 99th percentile having 13% and the largest regrowth rate being 60%. For simplicity, we abstract away from this heterogeneity and take one regrowth rate for all of Brazil. To inform where this heterogeneity comes from and the extent to which regrowth rates are changing over time, in panel (b) we plot the regrowth rates by Brazil’s large regions: North (N), North-East (NE), South (S), Center-West (CW), and South-East (SE).


D.8 Maps of inverted fundamentals

The headline deforestation-TFP maps appear in the main text (Figure 6). The remaining inverted fundamentals—the agricultural and non-agricultural TFPs and the residential amenities—are mapped here.


D.9 Model fit

This subsection reports the model’s fit on moments the inversion never targets. The inversion of Appendix C.3 rationalizes observed deforestation, land use, sectoral employment shares, and agricultural revenues exactly, so fit on those series is uninformative by construction; the validation content lies in the objects the calibrated model produces as *outputs*. Figure 22 compares the model-implied land price q_r —the capitalized value of the endogenous land rent—with observed land values from the agricultural census: the model reproduces the cross-regional distribution of land prices it was never asked to match, which is the observable most directly tied to the leakage mechanism, since it is relative land values that the supply-versus-demand decomposition turns on. Figures 23 and 24 compare model-implied average labor income, overall and in agriculture, with their census counterparts, and Figure 25 does the same for agricultural value added from the regional accounts. In each case micro-regions line up along the 45-degree line, with the fit tightest where agricultural activity (circle size) is largest.

Figure 18: Forest regrowth rates




(a) Distribution of forest regrowth rates of micro-regions



(b) Trends of forest regrowth rates by 5 regions


Note: The figures above illustrate the regional and temporal variation of forest regrowth rates across Brazil. Forest regrowth is defined as the percentage of non-forest area that becomes forest in a given year. Panel (a) above shows the distribution of regrowth rates by micro-region for 4 different years: 2003, 2008, 2013, and 2018. Panel (b) shows the yearly regrowth rates of Brazil's five large macro-regions: the North East (NE), the North (N), the South (S), and the Center-West (CW) for every year in the study period of the counterfactual analysis (2003-2020).

Figure 19: Estimated Agricultural TFPs (composite)



Note: These maps show the spatial distribution of the composite agricultural productivity across Brazil for the years 2003, 2010, and 2019. The estimation of each year's productivities is done through an exact inversion of the model in that year, assuming equilibrium conditions.

Figure 20: Estimated Non Agricultural TFPs




Note: These maps show the spatial distribution of the non-agricultural productivity across Brazil for the years 2003, 2010, and 2019. The estimation of each year's productivities is done through an exact inversion of the model in that year that assumes equilibrium conditions.

D.10 Within-region determinants of Z^D

Table 17 complements the cross-sectional correlates of Section 5.6 by adding micro-region *and* year fixed effects, so identification comes purely from within-region variation over time. Only time-varying characteristics survive: the clearing-cost fundamentals that carry the cross-section (biomass, geography) are absorbed by the region effects, so the within- R^2 is mechanically small and the informative content is the coefficients. Climate is flat within region—as expected of a slow-moving fundamental—while the depletion of remaining forest is an active within-region margin for both Z^D

Figure 21: Estimated Amenities



Note: This map shows the spatial distribution of the residential amenities across Brazil. Amenities are time-invariant in the inversion (backed out from 2010), and are recovered by rationalizing observed migration flows given estimated bilateral migration costs and region-specific yearly utilities.

and deforestation.

E The open economy: international trade

Brazil’s deforestation-linked commodities are overwhelmingly export-oriented—roughly two-thirds of soybeans and a quarter of beef are sold abroad—so the demand that bids for newly cleared land originates substantially at the world market rather than in domestic final consumption ([Hansen and Wingender 2023](#)). The model is therefore an *open* economy: the rest of the world (RoW) is an explicit trading partner. This is the headline specification; the closed economy of a model in which leakage operated only through domestic goods and labor markets is recovered as the special case in which the foreign block is switched off, and we use it as a robustness benchmark in Appendix [E.4](#). This appendix details the foreign block, its calibration, and the open-vs-closed comparison.

Table 17: Within-region determinants of Z^D and deforestation (time-varying characteristics)

	Level $\ln Z^D$ (1)	Fund. $\ln \overline{Z^D}$ (2)	(log) deforestation (3)
Temp. (mean)	0.001 [0.018]	0.001 [0.018]	0.017 [0.030]
$\ln$ Precip. (annual)	-0.079* [0.041]	-0.079* [0.041]	-0.120* [0.066]
$\ln$ Soil moisture	0.051* [0.030]	0.051* [0.030]	0.084 [0.051]
Fines / deforestation, $t - 1$	-0.063*** [0.019]	-0.063*** [0.019]	-0.102*** [0.029]
$\ln$ Remaining forest T^F	4.423*** [0.648]	4.167*** [0.648]	7.243*** [1.046]
Micro-region FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Within R^2	0.027	0.025	0.027
Observations	6156	6156	6156

Note: Two-way (micro-region + year) fixed-effects OLS, micro-region-year, 2002–2019, SE clustered by micro-region. Identification is purely within-region over time, so only time-varying characteristics enter; the time-invariant clearing-cost fundamentals of Table 5 (biomass, geography) are absorbed by the region effects, which is why within- R^2 is mechanically small. The enforcement regressor is the IBAMA *Flora* detection rate (fines per unit deforested area, asinh-transformed), lagged one year only, for the same reason as Table 5: the contemporaneous rate shares the same-period deforestation with both Z^D and, in column (3), the outcome variable itself. It is detection-biased (fines = infractions $\times$ detection probability), so a negative coefficient is a lower bound on deterrence. Z^D is $\approx 90\%$ between-region, deforestation $\approx 78\%$; this table reads the residual within-region variation. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.




Figure 22: Model fit: comparing predicted and observed land prices


Note: This compares the predicted land prices from model calibration q_r with observed land value. Each circle is a micro-region, and the size of the circle corresponds to the total agricultural area

E.1 The rest-of-the-world block

The economy is the R Brazilian micro-regions of the main text plus a single RoW node that trades goods but supplies no labor or land (consistent with the no-international-migration assumption maintained throughout). RoW enters as one additional *destination* for each tradeable good—the K agricultural crops and the non-agricultural composite—sourcing it across micro-regions through the same CES/Armington demand system that governs domestic trade. For good g the RoW expenditure share on origin o takes the same form as the domestic shares in equation (8),

$$\pi_{\text{RoW}}^{go} = \frac{(\tau_{\text{RoW}}^o)^{1-\sigma} (p^{go})^{1-\sigma}}{\sum_{r=1}^R (\tau_{\text{RoW}}^r)^{1-\sigma} (p^{gr})^{1-\sigma}}, \quad \tau_{\text{RoW}}^o = (\text{dist-to-port}_o)^\kappa, \quad (54)$$

Figure 23: Model fit: labor income



Note: This compares the predicted average labor income from model calibration y_r^L with observed average labor income. Each circle is a micro-region, and the size of the circle corresponds to the total agricultural area


where the international iceberg τ_{RoW}^o depends on the micro-region's distance to its relevant export gateway—the soy port for agricultural goods, the nearest general port for non-agricultural goods—through the same distance elasticity κ estimated from interstate trade flows (Appendix D.6).

World demand is price-elastic. RoW demand for the Brazilian composite of good g is a constant-elasticity, downward-sloping curve in the price index it faces,

$$Q_{\text{RoW}}^g = A^g (p_{\text{RoW}}^g)^{-\zeta}, \quad p_{\text{RoW}}^g = \left(\sum_o (\tau_{\text{RoW}}^o)^{1-\sigma} (p^{go})^{1-\sigma} \right)^{1/(1-\sigma)}, \quad (55)$$

so foreign expenditure is $X_{\text{RoW}}^g = p_{\text{RoW}}^g Q_{\text{RoW}}^g = A^g (p_{\text{RoW}}^g)^{1-\zeta}$. Market clearing for each origin is then equation (9): the value of origin o 's output equals domestic absorption

Figure 24: Model fit: labor income in agriculture



Note: This compares the predicted average labor income from model calibration y_r^{LA} with observed average labor income from the agricultural sector. Each circle is a micro-region, and the size of the circle corresponds to the total agricultural area

plus the foreign term $\pi_{RoW}^{go} X_{RoW}^g$. The residual export demand Brazil faces for good g has elasticity ζ , which summarizes how the world's other suppliers and buyers respond when Brazilian prices move; the fixed-quantity case is the inelastic corner $\zeta \rightarrow 0$.

Two properties matter for leakage. First, the regional split of any given foreign demand is entirely model-determined and price-elastic: if protection makes a region's land scarcer and raises its costs, its share π_{RoW}^{go} falls and the export pull relocates to lower-cost regions. Second, with $\zeta > 0$ the *aggregate* foreign pull itself responds to price, so a conservation policy that raises Brazilian costs also shrinks total exports rather than merely reallocating them. The leakage channel therefore operates at the world market as well as within Brazil.




Figure 25: Model fit: comparing predicted and observed agricultural value added

Note: This compares the predicted total GDP in the agricultural sector from model calibration with observed value added from the agricultural sector. Each circle is a micro-region, and the size of the circle corresponds to the total GDP

E.2 Calibration

The block needs three objects. (i) The demand level A^g is pinned so the curve passes through the data: $A^g = Q_{\text{obs}}^g (p_{\text{RoW,obs}}^g)^\zeta$, i.e. at observed world prices RoW buys Brazil's observed export quantity, set by the national export share of output for each model good. (Soybeans and beef individually are $\approx$ two-thirds and one-quarter export-oriented; because the model aggregates crops into three groups, the calibrated group shares are lower—temporary crops $\approx$ 24%, pasture $\approx$ 5%, and non-agriculture $\approx$ 6%, the last computed against non-agricultural value added.) This is a national aggregate, measured reliably and independent of the cross-region attribution of customs records—the model never uses the *municipal* export geography, which COMEX records at the port or exporter's fiscal municipality rather than at the production

region, so the spatial pattern of exports is an *output* of the model rather than an input. (ii) The international iceberg uses each micro-region’s distance to its export gateway with the interstate distance elasticity κ (Appendix D.6). (iii) The two trade elasticities are taken from Domínguez-Iino (2026): the across-region Armington elasticity $\sigma = 9.6$ (his across-county estimate) and the foreign-demand elasticity $\zeta = 6$ (his across-nation estimate, which is precisely the residual-demand elasticity Brazil faces). That $\zeta < \sigma$ encodes Brazil’s modest market power in soy and beef.

E.3 Closure: balanced trade

Export earnings raise each region’s factor income through market clearing and are spent, so the household budget is identical in form to the closed model,

$$E_r = \underbrace{\bar{y}_r^L L_r + v_r T_r^A}_{\text{factor income}} - p_r I_r^D, \quad (56)$$

with net exports $NX_r = \sum_g \pi_{\text{RoW}}^{gr} X_{\text{RoW}}^g$ reported as a diagnostic but not drained from the budget. This balances trade *in aggregate value per region*, not sector by sector: a region may run an agricultural surplus financed by a non-agricultural deficit (or vice versa), so cross-region specialization is preserved; only a region’s total absorption is tied to its total output. We adopt balanced trade as the headline because Brazil’s *overall* goods balance is small and mean-reverting (within roughly $\pm 2.5\%$ of GDP), even though the agricultural surplus is large and rising and the non-agricultural balance swings widely (Appendix E.4); the volatile component is fuels and manufactures, orthogonal to the land mechanism. The exogenous-deficit closure of Dekle et al. (2007), which subtracts NX_r from (56), could be adopted as a robustness variant, but we do not pursue it: draining the agricultural surplus interacts poorly with non-homothetic demand and a non-traded non-agricultural sector, pushing non-agricultural expenditure negative in the lowest-income regions. A final note: the fixed-quantity ($\zeta \rightarrow 0$) closure is not only economically extreme—it asserts the world buys a fixed quantity at any price—but also numerically unstable, because perfectly inelastic foreign demand makes foreign expenditure rise with the price level and breaks the contraction that solves the model. Regularity of the nested demand system requires $\zeta > \theta$ (the foreign nest must be more substitutable than the across-crop nest, $\theta = 2$ here, consistent with the ordering $\theta < \zeta < \sigma$); in practice the forward

fixed point also becomes stiff as ζ falls toward this floor, so the empirically-relevant range $\zeta \in [6, \sigma]$ is both valid and well-conditioned.

E.4 Open versus closed economy

Switching the foreign block off ($Q_{\text{RoW}}^g = 0$) recovers the closed economy exactly, which we use as a benchmark. The headline counterfactual results (Section 6) are computed in the open economy; here we report the closed-economy benchmark alongside them. The closed benchmark is the *same* model with the foreign block switched off in both the inversion and the forward simulation, so the comparison controls for all other parameters (including σ). Opening the economy leaves cumulative deforestation leakage essentially unchanged: Table 18 reports the closed and open figures side by side and they differ by at most a few tenths of a percentage point, and the decomposition into gross avoided deforestation, the share lost to general-equilibrium leakage, and the net amount averted (Table 19) is likewise nearly identical. This near-identity reflects that, under balanced trade, the open economy re-routes demand without changing the supply-side determinants of leakage (the land-supply elasticity, migration, and within-Brazil substitution are common to both). This is the central message of the trade extension: allowing conservation policy to relocate export supply at the world market does not undo the small-leakage result.

Table 18: Cumulative deforestation leakage: closed vs open economy

Policy	Closed economy cumulative leakage (%)	Open economy cumulative leakage (%)
Priority List	3.2	3.7
Protected Areas (territories)	5.7	6.4

Table 19: Policy effect on cumulative deforestation: gross avoided, leakage, net averted

Policy	Economy	Gross avoided	GE leakage	Net avoided	Leakage (%)
Priority List	Open	80,051	2,935	77,116	3.7
Priority List	Closed	75,803	2,424	73,379	3.2
Protected Areas (territories)	Open	7,311	467	6,844	6.4
Protected Areas (territories)	Closed	6,959	398	6,561	5.7




Figure 26: Gross policy effect split into net deforestation averted and the part lost to general-equilibrium leakage, closed vs open economy.

E.5 Trade imbalances in the data

Figure 27 motivates the balanced-trade closure. Agriculture runs a small, smooth, rising surplus (~ 1.3 – 2.5% of GDP); non-agriculture swings into a deep boom-era deficit (the 2013–14 trough is a fuel-import episode); and the overall goods balance is comparatively small and mean-reverting. Balanced trade is thus a poor approximation sector by sector but a good one in aggregate—exactly the margin the open block leaves to the RoW node.

F Additional counterfactual exhibits

This appendix collects additional exhibits from the counterfactual analysis of Section 6: the changes in deforestation productivities that define the no-policy counterfactuals, and the full set of counterfactual maps.

F.1 Counterfactual trends, spatial incidence, and the policy-design exercises

The following exhibits are referenced from Section 6 and moved here for space: the spatial map of where displaced clearing lands, the policy-strength curve, and the




Figure 27: Brazil's net goods exports by broad sector and overall, % of GDP (national value added from IBGE municipal accounts).

by-horizon sufficient-statistic table.


G Direct tests of the mediating channels

This appendix constructs the market-access exposure to conservation used in Section 7.2 and reports the reduced-form regressions. All exhibits are produced by `code/10_mechanism_exposure/` (scripts 01–02).

G.1 The exposure


A conservation restriction in region j removes Cons_j km² of agricultural land (its avoided clearing, from the Priority-List counterfactual). This reaches region i two ways. Directly, through a *supply* margin on the listed region itself: locking up Cons_i of i 's own farmland makes its remaining land scarcer. Indirectly, through a *demand* margin on everyone else: conservation in j diverts crop demand toward regions that sell the same crops to the same markets. We summarize the latter with a market-access exposure in the spirit of [Donaldson and Hornbeck \(2016\)](#), who show agricultural

Figure 28: Changes in deforestation productivity: Priority List




Note: This figure illustrates the difference in estimated deforestation TFPs between (i) the estimated value from the inversion of the model using observed data, and (ii) the simulated scenario in which there is no Priority List following equations (14) and (15). Panel (a) on top, shows the trends in deforestation TFP Z_{rt}^D for both scenarios over the study period 2003-2019. Panel (b) at the bottom show the spatial distribution of the differences in Z_{rt}^D between scenarios (i) and (ii).

Figure 29: Changes in deforestation productivity: Protected Areas




Note: This figure illustrates the difference in estimated deforestation TFPs between (i) the estimated value from the inversion of the model using observed data, and (ii) the simulated scenario in which there are no new Protected Areas from 2003 following equation (16). Panel (a) on top, shows the trends in deforestation TFP Z_{rt}^D for both scenarios over the study period 2003-2019. Panel (b) at the bottom show the spatial distribution of the differences in Z_{rt}^D between scenarios (i) and (ii).

Figure 30: Counterfactual deforestation maps: Priority List




Note: This figure shows the Priority-List-driven changes in cumulative deforestation (in thousands of square kilometers) across the various micro-regions of Brazil for years 2008, 2013, and 2019. Specifically, it maps the differences between deforestation in the data (which includes the effect Priority List policy) and the no-policy counterfactual simulation. The regions where the Priority List policy is implemented have some large decreases in deforestation, while the rest of the country has increases, albeit much smaller in magnitude.

Figure 31: Counterfactual deforestation maps: Protected Areas



Note: This figure shows the new-Protected-Area-driven changes in cumulative deforestation (in thousands of square kilometers) across the various micro-regions of Brazil for years 2003, 2008, 2013, and 2019. Specifically, it maps the differences between deforestation in the data (which includes the effect Protected Area policy) and the no-policy counterfactual simulation. The establishment of a Protected Area in part of a municipality leads to a direct reduction in deforestation by decreasing the area available for deforestation (T_{rt}^F), which mechanically reduces $Z_{rt}^D = \overline{Z}_{rt}^D (T_{rt}^F)^\psi$. However, it also leads to an increase in deforestation via GE effects. Increases in deforestation seem to be of a much smaller magnitude overall.

Figure 32: General Equilibrium Effects maps: Protected Areas




Note: This figure shows the spatial distribution of the part of new-Protected-Area-driven changes in cumulative deforestation (in thousands of square kilometers) that is due to GE effects. Namely, it is the difference between the scenario with new Protected Areas established (data) and the scenario with TFPs as in the data, but prices as in the no-new-Protected-Areas simulation. The spatial units are the various micro-regions of Brazil for years 2003, 2008, 2013, and 2019.

Table 20: Sensitivity of cumulative leakage to model parameters

parameter	value	leak PL (%)	Δ PL (pp)	leak PA (%)	Δ PA (pp)
baseline	—	3.67	—	6.39	—
alpha	0.80	4.25	+0.59	7.58	+1.19
alpha	0.90	3.94	+0.28	6.93	+0.54
alpha	1.10	3.42	-0.24	6.00	-0.39
alpha	1.15	3.32	-0.34	5.92	-0.46
beta	0.85	3.62	-0.04	6.36	-0.03
beta	0.90	3.63	-0.03	6.06	-0.33
beta	0.93	3.65	-0.01	6.36	-0.03
beta	0.97	3.72	+0.05	6.49	+0.10
beta	0.99	3.76	+0.10	6.50	+0.11
chi	1.05	3.87	+0.21	6.73	+0.34
chi	1.20	3.80	+0.13	6.66	+0.27
chi	1.35	3.73	+0.07	6.54	+0.16
chi	1.65	3.65	-0.01	6.37	-0.02
chi	1.80	3.62	-0.05	6.31	-0.08
chi	2.00	3.59	-0.07	6.24	-0.15
delta	0.10	0.59	-3.08	1.05	-5.34
delta	0.20	1.33	-2.33	2.25	-4.14
delta	0.30	2.30	-1.37	4.18	-2.21
delta	0.40	3.59	-0.07	6.29	-0.10
delta	0.45	4.41	+0.75	7.73	+1.34
delta	0.50	5.41	+1.74	9.12	+2.73
delta	0.55	6.61	+2.94	11.13	+4.74
delta	0.60	8.10	+4.44	13.57	+7.19
delta	0.62	8.79	+5.12	14.74	+8.35
delta	0.64	9.58	+5.91	15.69	+9.30
delta	0.66	10.43	+6.76	17.28	+10.90
delta	0.68	11.34	+7.67	18.61	+12.22
delta	0.70	12.41	+8.75	20.34	+13.96
delta	0.80	20.42	+16.76	31.86	+25.47
delta	0.90	39.67	+36.01	53.90	+47.51
kappa	1.05	3.67	+0.00	6.39	+0.00
kappa	5.00	3.67	+0.00	6.39	+0.00
kappa	10.00	3.67	+0.00	6.39	+0.00
kappa	100.00	3.67	+0.00	6.39	+0.00
no migration	—	3.66	-0.00	6.43	+0.04
sigma	6	3.99	+0.32	6.91	+0.53
sigma	7	3.86	+0.19	6.76	+0.37
sigma	8	3.77	+0.10	6.64	+0.25
sigma	9	3.70	+0.03	6.61	+0.22
sigma	10	3.66	-0.01	6.36	-0.03
sigma	11	3.62	-0.04	6.33	-0.05
sigma	12	3.60	-0.06	6.33	-0.06
theta	1.2	3.68	+0.02	6.35	-0.03
theta	1.5	3.68	+0.01	6.25	-0.14
theta	2.5	3.66	-0.00	6.48	+0.10
theta	3.0	3.66	-0.01	6.42	+0.04
theta	4.0	3.64	-0.03	6.29	-0.10
theta	5.0	3.63	-0.04	6.39	+0.00
zeta	3	3.56	-0.10	6.16	-0.23
zeta	4	3.60	-0.06	6.32	-0.06
zeta	5	3.64	-0.03	6.31	-0.08
zeta	7	3.69	+0.03	6.52	+0.14
zeta	8	3.73	+0.07	6.38	-0.01
zeta	9	3.77	+0.11	6.44	+0.05

Note: Cumulative leakage (2003–2018, open economy) for the Priority List and Protected Areas when a single model parameter is varied from its baseline, one at a time. Each row is a full re-inversion and re-resolution of the model at the stated value; the “ Δ ” columns give the change from the baseline in percentage points. Parameter symbols follow the calibration of Section 5: δ is the deforestation-supply elasticity, σ the Armington elasticity across regional varieties, θ the across-good substitution, ζ the world-demand elasticity, and χ the occupational-switching elasticity; the remaining rows vary preference and mobility-cost parameters (α, β, κ) and, in the last row, shut internal migration off entirely. Leakage is monotone and roughly convex in δ and mild in the substitution elasticities, consistent with the supply-driven mechanism.

Figure 33: Where leakage lands: general-equilibrium effects in space (Priority List)




Note: This map shows *where displaced clearing lands*: each micro-region’s cumulative general-equilibrium effect on deforestation, 2003–2018 (data minus counterfactual B, in km²; the difference between the world with the Priority List and one with the same deforestation-TFP fundamentals but prices as in the no-policy simulation). Color is on a log₁₀ scale, which isolates the heavy-tailed leakage signal; only regions that *absorb* displaced clearing (positive GE effect) are colored. Priority-List micro-regions are shaded grey. GE effects are concentrated in specific untargeted regions—the next-cheapest clearers—and, in-situ, within the Priority List itself.

land values capitalize a region’s access to markets:

$$\underbrace{\text{Cons}_i}_{\text{supply (own)}} \quad \text{and} \quad \underbrace{\text{Exposure}_i = \sum_{j \neq i} m_{ij} \text{Cons}_j}_{\text{demand}}, \quad m_{ij} = \sum_k s_{ik} \left(\sum_d \omega_{id}^k \pi_{jd}^k + \omega_{i,\text{RoW}}^k \pi_{j,\text{RoW}}^k \right), \quad (57)$$

where m_{ij} is a crop-and-market overlap: s_{ik} is i ’s crop- k revenue share, π_{jd}^k is j ’s share of destination d ’s crop- k spending, and ω_{id}^k is the share of i ’s crop- k sales going to d . Destinations include all domestic micro-regions and the Rest of World; destination market size is proxied by total value added (regional income). Bilateral shares follow the standard CES gravity in observed trade costs τ (the Araujo 2017 travel-cost network) with $\sigma = 9.6$ and pre-policy (2005–07) crop revenues; the engine is solved

Figure 34: Leakage and avoided deforestation as policy strength varies




Note: Cumulative leakage (top) and gross vs. net avoided deforestation (bottom), 2003–2018, open economy, as the Priority-List enforcement intensity f varies (the targeted set held fixed at the actual Priority List). The dashed line marks the estimated baseline strength $f = 0.48$. The multiplier applied to targeted regions' $\bar{Z}^D$ is $(1 - f)^{1-\delta}$.

on the same τ , so Cons_j and m_{ij} share a trade-cost concept. Exposure_i is thus large when i sells the crops the conserved regions sold, into the same markets. It uses only observed geography and the model's land shock Cons_j , not any endogenous outcome.

A multi-crop Armington model rationalizes the *form* of m_{ij} —to first order $\partial \ln v_i / \partial (-\hat{T}_j) = \frac{\sigma-1}{\sigma^2} m_{ij}$ —but we use the levels exposure as a descriptive market-access regressor, not as an estimate of that structural elasticity (the derivation is marginal, whereas conservation is a large cumulative shock, and the outcome is a revenue-based rent proxy rather than the model's v_i). For each outcome y_{it} we estimate

$$\ln y_{it} = \gamma_{\text{sup}} \text{Cons}_{it} + \gamma_{\text{dem}} \text{Exposure}_{it} + \mu_i + \lambda_t + \varepsilon_{it}, \quad (58)$$

Figure 35: Sensitivity of cumulative leakage to model parameters: tornado



Note: Each bar spans the range of cumulative Priority-List leakage (2003–2018, open economy) induced by varying the indicated parameter across its swept values, one at a time, with every cell a full re-inversion and re-resolution of the model (Table 20); bars are anchored to include the baseline (dashed line) and sorted by the size of the swing. Dots mark the individual sweep cells; the highlighted dot on the δ bar is the Legal-Amazon (frontier) estimate $\delta = 0.62$, the top of the identified range, which maps to 8.8% leakage. The deforestation-supply elasticity δ dominates; no other parameter moves leakage by more than about a percentage point.

with micro-region and year fixed effects, standard errors clustered by micro-region, and each regressor standardized to unit standard deviation so the coefficients are comparable across outcomes.

G.2 Land rents

The land-rent outcome is the observed, revenue-based proxy $v_i = \sum_c \alpha_c \text{revenue}_c / \text{area}$ (α_c the national land-rent share per crop; built from PAM revenues, independent of the model inversion). Table 21 reports a long difference (mean 2015–18 minus 2008–10). Panel a is a benchmark with no GE content: being on the Priority List raises own land value, and the continuous own-conservation intensity enters positively and significantly, in the Full sample and the Legal Amazon alike (non-BLA is omitted—

both are absent outside the frontier). Panel b enters the supply (Cons_i) and demand ($\sum_{j \neq i} m_{ij} \text{Cons}_j$) margins jointly: the own-supply term is positive and significant throughout and stronger on the frontier, while the demand term is positive but small and imprecise, consistent with a muted rent incidence at high σ .

G.3 The mediating channels, in quantities

Applying the same exposure to the mediating margins (Table 22): labor reallocation (workers per hectare), the intensity of land use (agricultural value added per hectare), and the *composition* of land—the share of agricultural land in crops rather than pasture. If exposed regions absorb diverted demand by intensifying and shifting toward crops rather than expanding at the extensive margin, the composition margin loads on the demand exposure.

G.4 Where leakage lands

That the demand exposure predicts the location of the model’s leakage is quantified in Table 23 and mapped in Figure 10: across untargeted regions, log absorbed leakage rises strongly with the exposure, and the coefficient survives controlling for deforestation productivity $\log Z^D$ (which enters with near-unit elasticity). Where displaced clearing lands is governed first by residual land supply and second, robustly, by demand substitutability with the conserved regions—the two are complements.

G.5 Robustness

Table 24 re-estimates the land-rent regression with the theory-consistent fractional shock Cons_j/T_j (own Cons_i/T_i) in place of levels, and with the demand market access split into a crop-market and a pasture-market overlap. Both leave the reading unchanged.

Table 21: Land rents rise near conservation: own supply and demand market access

Dependent Variable:	$\Delta \log \text{ land rent}$			
Specification	PL dummy		Own-conservation intensity	
Sample	Full	Legal Amazon	Full	Legal Amazon
Model:	(1)	(2)	(3)	(4)
<i>Variables</i>				
On Priority List	0.2266*** (0.0873)	0.2129** (0.0987)		
Own conservation, Cons_i			0.0517*** (0.0159)	0.0987*** (0.0346)
<i>Fit statistics</i>				
Observations	543	102	543	102
R^2	0.01621	0.05004	0.01586	0.05169
<i>Heteroskedasticity-robust standard-errors in parentheses</i>				
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>				

Dependent Variable:	$\Delta \log \text{ land rent}$	
	Full	Legal Amazon
Model:	(1)	(2)
<i>Variables</i>		
Own conservation (supply)	0.0510*** (0.0162)	0.0998*** (0.0337)
Market overlap (demand)	0.0037 (0.0176)	0.0207 (0.0289)
<i>Fit statistics</i>		
Observations	543	102
R^2	0.01594	0.05396
<i>Heteroskedasticity-robust standard-errors in parentheses</i>		
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		

Note: Long difference in the log observed land-rent proxy. *Panel a:* on-Priority-List dummy and own-conservation intensity Cons_i , Full sample and Legal Amazon. *Panel b:* supply (Cons_i) and demand market access ($\sum_{j \neq i} m_{ij} \text{Cons}_j$) of (57) entered jointly. Regressors standardized to unit SD; robust standard errors.

Table 22: Direct tests of the mediating channels

Model:	Workers / ha (1)	Value added / ha (2)	Cropland / ag area (3)
<i>Variables</i>			
Own conservation (supply)	0.0102*** (0.0031)	0.0185*** (0.0066)	0.0985*** (0.0248)
Market overlap (demand)	0.0032 (0.0041)	0.0635*** (0.0133)	0.0033 (0.0310)
<i>Fixed-effects</i>			
Microregion	Yes	Yes	Yes
Year	Yes	Yes	Yes
<i>Fit statistics</i>			
Observations	8,912	8,912	8,657
R ²	0.99926	0.95111	0.96247

Clustered (Microregion) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Equation (58) on the micro-region-year panel; micro-region and year fixed effects; standard errors clustered by micro-region. Each column is a channel; rows report the supply (Cons_i) and demand market-access ($\sum_{j \neq i} m_{ij} \text{Cons}_j$) coefficients, standardized to unit SD. Cropland/ag-area is a composition share.

Table 23: What governs where leakage lands: supply capacity and demand access

Dependent Variable:	log absorbed leakage	
Model:	(1)	(2)
<i>Variables</i>		
Market overlap (demand)	0.5921*** (0.0594)	0.2620*** (0.0329)
log Z^D (defor. TFP)		1.233*** (0.0505)
<i>Fit statistics</i>		
Observations	491	491
R ²	0.14980	0.70051
<i>Heteroskedasticity-robust standard-errors in parentheses</i>		
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		

Note: Cross-section of untargeted (non-Priority-List) micro-regions; outcome is log cumulative absorbed leakage $L_i = \sum_t \max(T_r^D - T_{r,ap}^D, 0)$. Regressors: the demand exposure $\sum_{j \neq i} m_{ij} \text{Cons}_j$ (unit SD) and log Z^D . Robust standard errors. This regresses the model's own GE leakage on an exposure built from observed geography, so it is a coherence check that the exposure summarizes where the structural model routes leakage.

Table 24: Land-rent regression: robustness of the demand term

Dependent Variable:	$\Delta \log$ land rent	
	Full	Legal Amazon
Model:	(1)	(2)
<i>Variables</i>		
Own conservation (supply)	0.0664* (0.0362)	0.1204 (0.0818)
Market overlap (demand)	-0.0238 (0.0221)	-0.0766 (0.0470)
<i>Fit statistics</i>		
Observations	543	102
R ²	0.02571	0.10867
<i>Heteroskedasticity-robust standard-errors in parentheses</i>		
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		

Dependent Variable:	$\Delta \log$ land rent	
	Full	Legal Amazon
Model:	(1)	(2)
<i>Variables</i>		
Own conservation (supply)	0.0525*** (0.0163)	0.1003*** (0.0338)
Crop-market overlap (demand)	-0.0017 (0.0160)	0.0434 (0.0523)
Pasture overlap (demand)	-0.0197 (0.0180)	0.0389 (0.0598)
<i>Fit statistics</i>		
Observations	543	102
R ²	0.01805	0.05639
<i>Heteroskedasticity-robust standard-errors in parentheses</i>		
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		

Note: As Table 21b. *Top:* fractional shock Cons_j/T_j (own Cons_i/T_i), whose theory-implied coefficients are $(\sigma - 1)/\sigma^2$ and $1/\sigma$. *Bottom:* demand market access split into crop-market and pasture-market overlap. Regressors standardized to unit SD; robust standard errors.