Efficiency and Redistribution in Environmental Policy: An Equilibrium Analysis of Agricultural Supply Chains

Tomás Domínguez-Iino

University of Chicago

I build an empirical model of the South American agricultural sector to show how environmental policy is transmitted along a supply chain when regulation at the externality's source is infeasible. Given obstacles to a first-best carbon tax on farmers, I show how second-best alternatives—downstream agribusiness taxes—reduce upstream emissions, but their effectiveness is limited by poor targeting while also being regressive. Agribusiness monopsony power worsens targeting by lowering pass-through to upstream farmers in uncompetitive and emissions-intense regions, thus eroding the Pigouvian signal where social cost is highest. By contrast, small-scale but well-targeted upstream interventions perform robustly when markets face preexisting distortions.

I. Introduction

In many of the major industries contributing to climate change, the environmental externality often coexists with other preexisting distortions—market power being a case in point. How can we regulate such industries

This paper was the first chapter of my PhD dissertation at New York University. I thank Alessandro Lizzeri, Guillaume Fréchette, Elena Manresa, and Paul Scott for their unconditional support during the early stages of this project. For helpful feedback, I am grateful to Milena Almagro, Tim Christensen, Christopher Flinn, Martin Rotemberg, Sharon Traiberman, and Daniel Waldinger. I have also benefited from conference discussions by Bouwe Dijkstra, Joel Ferguson, Allan Hsiao, Ander Iraizoz-Olaetxea, Pierre Magontier, and Nicolas Morales. Finally, I thank the editor and six anonymous reviewers for constructive suggestions

Electronically published February 25, 2026

Journal of Political Economy, volume 134, number 4, April 2026.
© 2026 The University of Chicago. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published by The University of Chicago Press.
https://doi.org/10.1086/739325

efficiently, and what are the distributional consequences of regulation? This paper provides an empirical framework to answer this question in the context of the South American agricultural sector, a global agricultural powerhouse with a major environmental impact, where the supply chain connecting farmers to consumers is intermediated by a concentrated agribusiness sector. A key feature of this setting is that agricultural emissions are mostly generated at the atomistic stage of the supply chain rather than at the concentrated stage. It is the millions of upstream farmers who make the environmentally relevant decisions, mostly through their land use choices, and not the large agribusiness firms further downstream. Given that environmental policies are easier to implement and enforce at the concentrated end of the supply chain, this raises the question of how much of their Pigouvian signal is eroded before reaching the upstream farmers whose incentives they ultimately aim to correct.

The goal of this paper is to evaluate how the transmission of environmental policy occurs along a supply chain, in particular when preexisting distortions lie between the stage where emissions are generated and the stage where regulation is feasible. The research challenge involves measuring the correlation between the two distortions of interest—the environmental externality and market power—and developing an equilibrium framework to study how they interact under different policy tools. To address this, I combine a variety of data sources to build a county-level panel of agricultural supply and demand, which I use to estimate an equilibrium model of the South American agricultural sector. On the model's supply side, I incorporate key margins determining emissions: how much land farmers deforest, which commodity they produce, and rich heterogeneity across the geographic locations where deforestation and production take place. On the demand side, I incorporate the funnel-like structure of agricultural supply chains: atomistic farmers at the upstream stage sell their output to a concentrated sector of downstream agribusiness firms with monopsony power.

Despite being a crucial item on the sustainable development agenda, environmental policy in developing-world agriculture faces multiple obstacles. First, the distributional effects of an agricultural carbon tax are regressive on both demand and supply: poor households spend a larger share of their income on food, and farmers often lie at the bottom of the income distribution. Second, agricultural commodities are traded in highly integrated global markets, resulting in substantial "leakage" risk: if one country unilaterally sanctions its imports from an emissions-intense producer, the goods are diverted to nonregulated consumer markets, and the externality remains uncorrected. Third, agricultural supply chains in developing countries are often fragmented, funnel shaped, and subject to preexisting

that substantially improved the paper. Any errors or omissions are my own. This paper was edited by Esteban Rossi-Hansberg.

distortions that may interfere with the performance of market-based policies. Beyond providing an ideal setting to study second-best environmental policy, agriculture is important in and of itself, accounting for 26% of global anthropogenic emissions (Poore and Nemecek 2018). In particular, South American agricultural emissions are 27% of world agricultural emissions (FAOSTAT), exceeding those of major sectors of the US economy such as transport or electricity generation.[1]

On the measurement side, the key empirical result is that I estimate lower supply elasticities in frontier agricultural regions near densely forested areas. I obtain this result by modeling farmer decisions with a nested discrete choice model of land use that incorporates an extensive margin of converting forested land into new agricultural land, as well as the choice of which specific agricultural commodity to produce on existing agricultural land. Hence, the model incorporates the two key margins driving agricultural emissions within a single framework while delivering estimates that are consistent with prior work estimating each margin separately. While I do not fully disentangle the microfoundations for why supply is less elastic in remote regions, most potential mechanisms are related with these regions being less developed overall. One leading explanation for why farmers in remote regions are relatively sluggish in their production responses is that they are more constrained in their access to the inputs required to adjust production.

For the model's demand side, I use granular data on domestic trade flows to document concentration among agribusiness intermediaries, a feature that I embed into my model with a layer of oligopsonistic intermediaries between farmers and final consumers. The farm-gate prices farmers receive are therefore marked down from the marginal revenue they generate for the intermediary, with the size of the markdown depending on two key objects: the supply elasticities of farmers and the degree of agribusiness concentration. Given that the market power of intermediaries is stronger when facing inelastic farmers, the geographic heterogeneity in supply elasticities is inherited by the geographic distribution of monopsony power. Because remote regions are the most emissions intense and have the least elastic supply, I find a positive spatial correlation between the degree of market power and the environmental externality's intensity.

The first part of my counterfactual analysis shows how a feasible environmental policy is transmitted along the supply chain. While a carbon tax at the emissions source (i.e., on upstream farmers) theoretically attains the first best through the textbook Pigouvian mechanism, this type of policy is mostly absent in developing-world agriculture because of logistical enforcement challenges as well as political infeasibility. Motivated

[1] South American agricultural emissions have hovered around 3 Gt CO_2-equivalent units (CO_2e) since 1990. US emissions from industry, electricity, and transport in 2018 were 1.5 Gt, 1.8 Gt, and 1.9 Gt, respectively (EPA 2018; FAOSTAT, http://www.fao.org/faostat/en/).

by these constraints, I evaluate a policy that is feasible but second best: a uniform tax (based on the average carbon content of each commodity across upstream producers) levied at the downstream stage on agribusiness firms. An international version of this is a carbon border adjustment levied at the port, while a domestic version is a flat carbon tax at the retail stage. Given the spatial pattern of supply elasticities, the downstream tax ends up being spatially mistargeted because it causes production to drop least in remote upstream locations where the environmental cost is highest. Because of inelastic supply, these remote regions are also where farmgate prices drop most. Since these locations are among the poorest, the distributional effects of a downstream tax are regressive on the supply side: the income of poor farmers is implicitly taxed at a higher rate than that of rich farmers. Finally, because the tax needs to be passed through the supply chain to farmers to shift their production incentives, agribusiness monopsony power plays a role by reducing pass-through and eroding the upstream transmission of the Pigouvian signal.[2]

The second part of my counterfactual analysis shows how a regulatory agency's optimal choice of policy instrument depends on the previously mentioned spatial correlation between market power and emissions intensity. Downstream policies, such as the taxes from the first part of my analysis, are mistargeted because they do not take into account the spatial heterogeneity in emissions intensities across upstream farmers. Moreover, market power worsens targeting by lowering pass-through most to the least competitive upstream locations, which are also the most emissions intense. By contrast, policies that are directly implemented upstream, such as a forest subsidy, are better targeted and robust to market structure because they avoid the pass-through distortions the downstream tax is subject to. However, these policies come with an enforcement cost that limits their scalability. Thus, the regulator faces a trade-off between targeting and scale when choosing between the two types of tools, with the starkness of the trade-off depending on the correlation between market power and the environmental externality. If the correlation is positive, as I find in my setting, market power worsens the targeting of the downstream tax, thus favoring direct implementation upstream (even if at small scale). If the correlation were negative, as it could well be in other empirical contexts, then market power would improve targeting by raising pass-through where environmental cost is highest. These qualitative results on

[2] Theoretically, pass-through in an imperfectly competitive market can be incomplete, complete, or more than complete, depending on the curvature of the side of the market subject to market power (Bulow and Pfleiderer 1983; Weyl and Fabinger 2013). In a monopoly problem this is determined by the curvature of demand, while in a monopsony problem the relevant analogue is the curvature of supply. In my empirical setting the log curvature of supply is such that monopsony power delivers incomplete pass-through.

the role of market power for the transmission of environmental policy add nuance to the classic intuition that by depressing quantities "the monopolist is the conservationist's friend" (Solow 1974).

In terms of related literature, this paper contributes to the intersection of industrial organization and environmental economics (Buchanan 1969; Ryan 2012; Fowlie, Reguant, and Ryan 2016) by showing that market power can be ambiguous for the transmission of environmental policy. The sufficient statistic that resolves this ambiguity is the correlation between the two distortions—the degree of market power and the environmental externality's intensity—since it determines where the pass-through of the Pigouvian signal is strongest. Moreover, by focusing on how firms exercise market power on their upstream suppliers, I open a distributional channel on the supply side that contrasts with most work on environmental policy incidence on consumers (Bento et al. 2009; Fabra and Reguant 2014). Second, my paper connects the agricultural trade literature (Costinot, Donaldson, and Smith 2016; Sotelo 2020; Pellegrina 2022) to the land use change literature from agricultural economics (Roberts and Schlenker 2013; Scott 2014; Souza-Rodrigues 2019). The trade literature uses the Ricardian framework of Eaton and Kortum (2002) to study how different commodities are allocated across existing agricultural land but abstracts from the extensive margin of land conversion. By contrast, the land use studies typically model the land use change margin as binary—land is either left in its natural forested state or used for agriculture broadly defined—but abstracts from which specific commodities are produced. My nested model incorporates both choice margins. Third, this paper relates to a growing literature in trade and the environment (Kortum and Weisbach 2017; Conte 2020; Conte et al. 2021; Farrokhi and Lashkaripour 2021; Hsiao 2021; Copeland, Shapiro and Taylor 2022; Cruz and Rossi-Hansberg 2024; Nath 2025). I contribute by focusing on the policy implications of the interaction between market power and the spatial heterogeneity in emissions intensities. Finally, this paper relates to recent studies on monopsony power in developing economies (Mitra et al. 2018; Bergquist and Dinerstein 2020; Dhingra and Tenreyro 2020; Zavala 2021; Chatterjee 2023; Rubens 2023), most of which focuses on the welfare impacts of market power per se. Instead, I incorporate market power to understand how it interferes with environmental policy.

II. Data

I construct a county-level panel of agricultural supply and demand from 1995 to 2017 using various data sources from Argentina and Brazil. The supply side consists of a county-level panel of land use, agricultural output, yields, and farm-gate prices for beef cattle, soybeans, maize, wheat, rice, sunflower, and sugarcane. On the demand side, I connect each

county's production to its nation-level destination markets using trade flow data.

Geographic unit of analysis and temporal frequency.—The Argentine data are reported at the department level (*partidos*). For Brazil I use time-consistent spatial units from Ehrl (2017): *áreas mínimas comparáveis* (AMCs). Hence, the term "counties" refers to Argentine departments and Brazilian AMCs.[3] Given that most of the data are from decadal agricultural censuses, changes over time are interpreted as long-run changes.

Land use and agricultural output.—County-level land use data are from decadal agricultural censuses of Argentina and Brazil. For Argentina, I complement the census with the Ministry of Agriculture's Datos Agroindustriales database (DA-MAGYP). Output data on crops and livestock for Argentina are from the census, DA-MAGYP, and the phytosanitary authority (Servicio Nacional de Sanidad y Calidad Agroalimentaria), while for Brazil, data are from the census and two municipal surveys: Produção Agrícola Municipal (PAM) and Pesquisa da Pecuária Municipal.

Agronomic productivity.—High-spatial-resolution agricultural productivity data are from the UN Food and Agriculture Organization global agroecological zones (FAO-GAEZ; IIASA/FAO 2012) and are reported as potential yields predicted from agroclimatic fundamentals.[4] To calibrate the model's productivity parameters, I complement FAO-GAEZ with realized yields from the output and land use datasets reported above. I use national-level data on extreme temperature from FAOSTAT to obtain temporal variation in agricultural productivity induced by weather shocks.

Trade flows.—National-level trade flows of commodities are from FAOSTAT. To determine sourcing within Argentina and Brazil, I use domestic supply chain data from Trase (https://trase.earth/), which is constructed from customs records and maps annual trade flows (in physical quantities and free-on-board [FOB] values) from source counties to national-level destination markets, as well as to the agribusiness firms intermediating the transactions. In appendix B, section B.2, I provide a detailed description of the Trase data and validate it against conventional data sources such as FAOSTAT. For the beef sector in Argentina, I complement Trase with meat-packer procurement records from DA-MAGYP.

Prices.—Farm-gate prices are obtained from production value and quantity data. For Brazil, the sources are the census and PAM. For Argentina, I

[3] Argentina has 512 "counties" with an average size of 0.5 million hectares, while Brazil has 4,298 with an average size of 0.2 million hectares. For comparison, the United States has 3,243 counties with an average size of 0.3 million hectares.

[4] To obtain "yields" for cattle, I construct a measure of cattle productivity by projecting the FAO-GAEZ pasture index on county-level cattle-stocking rates from agricultural censuses (see app. B, sec. B.1; apps. A–E are available online).

use DA-MAGYP. Destination prices are from Trase quantities and FOB values. Since values are reported as port of export FOB, destination prices reflect the price the agribusiness firms receive for delivery up to the port of export but not to the final destination market. Therefore, destination prices include domestic transport costs from farm to port but not international costs to final destination (which are paid by final destination consumers).

Emissions.—I compute land use change emissions using carbon density maps constructed from biomass data by Spawn and Gibbs (2020). For commodity-specific on-farm emissions intensities, I use life cycle assessment values from Poore and Nemecek (2018).

III. Stylized Facts

First, I provide background on agricultural emissions, emphasizing how the upstream decisions of farmers are the main driver of emissions footprints. Second, I present summary statistics highlighting the funnel shape of agricultural supply chains, introducing the notion that regulation is easier to enforce at the funnel's downstream bottleneck. Third, I show how land use patterns have evolved across time and space in my setting, highlighting how multiple commodities compete with each other and shape aggregate land use through comparative advantage. Finally, I describe how these features interact to motivate the key ingredients of my model and the policy counterfactuals I run.

A. Agricultural Emissions

FACT 1. Agricultural emissions are primarily determined by three choice margins, all of which take place at the upstream stage of the supply chain.

The first key determinant of agricultural emissions is the amount of land being cleared for production. More than 80% of agricultural emissions are generated upstream before the commodities leave the farm gate, mostly because of land use change and on-farm sources such as enteric methane (fig. 1*A*). The land use share is especially high in South America, above 70%, compared with a world average of 40% (FAOSTAT). The upstream feature of agricultural emissions contrasts with fossil fuels, where emissions are primarily released downstream at the point of consumption (i.e., the burning of fuel for energy).

The second key determinant is the specific commodity being produced. Emissions footprints vary widely across commodities, even after factoring in differing land requirements (fig. 1*A*). For example, beef contains 25 times

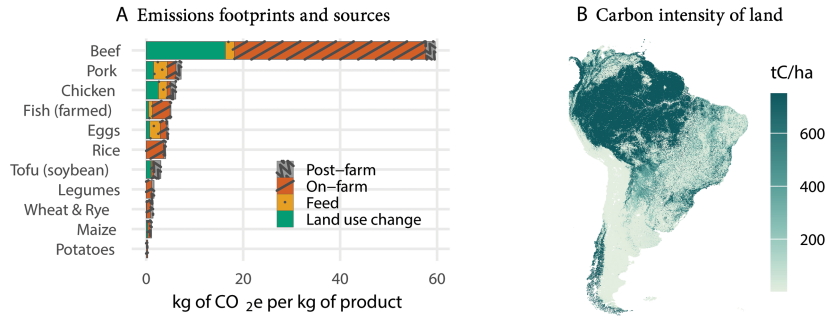

FIG. 1.—Sources of agricultural emissions along the supply chain. *A*, Emissions footprints (global averages) using publicly available data from Poore and Nemecek (2018). See appendix figure 10 (app. figs. 1–10 are available online) for footprints per kilocalorie or protein content. *B*, Aboveground carbon density constructed from biomass data from Spawn and Gibbs (2020).

more CO_2e per kilogram of protein than plant-based high-protein alternatives, largely because of enteric methane emissions unique to ruminants. The third key determinant is the location where land clearing and production take place: emissions footprints vary widely across space because of the highly uneven geographic distribution of carbon stocks (fig. 1*B*).

B. Agricultural Supply Chains

FACT 2. Agricultural supply chains are funnel shaped, as atomistic upstream farmers face a concentrated sector of downstream agribusiness buyers.

Farmers do not access consumer markets directly but rather through intermediating agribusiness firms. In Brazil there are 2.4 million upstream ranching establishments, 79% of which hold fewer than 50 head of cattle, facing a concentrated sector of downstream agribusiness firms. In the median county, the top three agribusiness firms account for 95% of sourced beef, with the top firm accounting for more than 60% (table 1).

Table 2 shows how agribusiness concentration correlates across upstream markets with a crude accounting-based markdown—the ratio of the farm-gate price with respect to the price the agribusiness firm receives at the port. Given that this gap is partly driven by transport costs, I control for an upstream location's remoteness with a standard market access measure. The correlation is negative: farm-gate prices are marked down more in upstream markets with a higher concentration of buyers, even after controlling for remoteness. Moreover, the relationship is robust to the

TABLE 1
AGRIBUSINESS CONCENTRATION IN UPSTREAM MARKETS

	Brazil			Argentina
	Beef	Maize	Soybean	Soybean
Number of agricultural establishments (sellers)	2,457,512	1,619,880	236,141	42,428
Number of agribusiness firms (buyers)	134	110	188	32
CR-1, national market	.34	.18	.15	.13
CR-3, national market	.68	.46	.42	.36
CR-1, median across local upstream markets	.66	1.00	1.00	.18
CR-3, median across local upstream markets	.94	1.00	1.00	.43
Share of upstream markets with only one agribusiness firm	.14	.86	.73	.01

NOTE.—Upstream markets are defined at the county level and for 2017, the latest year. Agricultural establishment data are from censuses. All other data are from Trase. CR-N is the top-N firm concentration ratio.

spatial granularity at which upstream markets are defined. Needless to say, these empirical patterns should not be interpreted as a causal relationship from market concentration to market outcomes. Concentration is itself a market outcome, and just like prices and markdowns, it is determined by supply and demand primitives (Bresnahan 1989). In a setting with potential monopsony power, the key primitive is the supply elasticity of farmers, which I estimate. Given that supply elasticities are shaped by land use decisions, we now turn to the drivers of land use.

TABLE 2
UPSTREAM MARKET CONCENTRATION AND ACCOUNTING-BASED MARKDOWNS

	Farm-Gate Price/Agribusiness Price Ratio (OLS)					
	(1)	(2)	(3)	(4)	(5)	(6)
CR-3	−.061*	−.062*	−.152***	−.155***	−.082***	−.082***
	(.029)	(.031)	(.022)	(.023)	(.022)	(.021)
Market definition	Mesoregion	Mesoregion	Microregion	Microregion	AMC	AMC
Control for market access	×	✓	×	✓	×	✓
Observations	260	260	899	899	4,663	4,663

NOTE.—Regressions are at the upstream market level for various degrees of spatial granularity across columns. Farm-gate prices are from censuses and production surveys. Agribusiness prices are Trase FOB prices. CR-3 is the three-firm concentration ratio. A location's market access is its mean distance to surrounding hubs (weighted by hub size), defined as $MA_i = \Sigma_j s_j d_{ij}^{-1}$, where d_{ij} is distance between upstream market i and hub j and s_j is the hub's share of national exports. Heteroskedasticity-consistent robust standard errors are reported in parenthesis.

* $p < .10$.

*** $p < .01$.

C. *Agricultural Commodities*

FACT 3. Agricultural commodities compete with each other in local land markets, shaping aggregate land use patterns through comparative advantage.

One of the key developments in South American agriculture over recent decades has been the dramatic expansion of soybean production (fig. 2*A*). Growing international demand, especially from Asia, has been a major driver behind such trends: more than 70% of soybean output is exported, and more than 50% of exports go to Asia (FAOSTAT). By crowding out other commodities, the soybean boom has resulted in a reallocation of agricultural production across land markets. Cattle grazing has

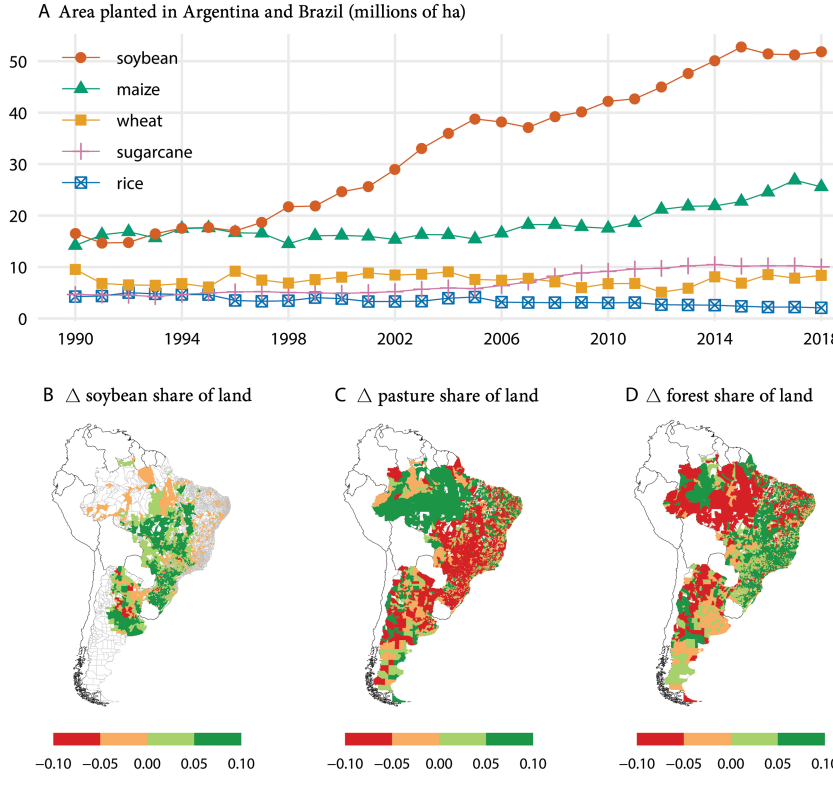

FIG. 2.—South American soybean boom across time and space. *A*, Total acreage allocated to major crops in Argentina and Brazil between 1990 and 2018. *B*, County-level changes in acreage allocated to soybeans, pasture, and forest between 1995 and 2017. Sources: Agricultural censuses, DA-MAGYP (Argentina), PAM (Brazil).

shifted from soybean-suitable areas (central-south Brazil, middle-east Argentina) to cheaper land markets in frontier agricultural regions (northern Brazil, north and west Argentina). While soybean expansion may not directly lead to deforestation, it can do so indirectly by displacing land-intensive cattle grazing to frontier regions where the forests lie, as reflected by figure 2*B*–2*D*. Accounting for interactions between commodities is therefore crucial for understanding deforestation in the South American context. In this setting, the most land-use-relevant commodities are beef cattle (pasture), soybeans, and maize—they account for 85% of all agricultural land, with pasture alone representing 70%.

D. Implications for Model Specification and Policy Counterfactuals

Fact 1 suggests that deforestation, commodity choice, and production location are key margins driving agricultural emissions. Hence, I propose a model where farmers make decisions along two emissions-relevant margins—how much land to clear and which commodity to produce—in a setting with rich spatial heterogeneity. Having different commodities compete for land allows comparative advantage to shape aggregate land use, in line with fact 3. I model agribusiness firms as oligopsonists in upstream markets given that they may plausibly hold market power as buyers, in line with fact 2. This specification nests the perfectly competitive case in order to alternate between conduct assumptions and evaluate how market power affects the transmission of environmental policy.

The stylized facts suggest that if the externality's source is upstream and atomistically heterogeneous but the downstream stage is concentrated and easier to regulate, then policymakers face a trade-off between efficient upstream targeting and ease of implementation downstream. This contrasts with sectors such as electricity, where fossil fuel emissions are concentrated downstream among a few large firms, making targeted enforcement at the emissions source logistically easier. While a carbon tax on farmers is the theoretically efficient solution, in practice it faces political barriers and requires enforcement capacity because of the atomized nature of production. In reality, policy design revolves around a range of second-best options at different stages of the supply chain.

Upstream policies often take the form of a conservation zone, which can be implemented for only a narrow set of locations because of monitoring costs.[5] By contrast, downstream tools can operate at scale because they avoid such enforcement costs: they are implemented at the supply

[5] One of the better-known policies in this category is Brazil's List of Priority Municipalities, in which the Brazilian Institute of Environment and Renewable Natural Resources increased monitoring for a subset of high-deforestation municipalities in the Amazon (Assunção et al. 2023).

chain's bottleneck (e.g., a tariff at the port) and then leverage the market mechanism to deliver a Pigouvian signal at scale to upstream farmers. However, for a downstream tax to be efficiently targeted, it requires precise tracing of the commodity's upstream origin, a condition that is not always satisfied in our setting.[6] This leaves us with a uniform downstream tax based on an average emissions footprint, regardless of the emissions intensity of the commodity's origin. Given that small-scale upstream regulation (e.g., conservation zones) and large-scale but coarsely targeted downstream taxes (e.g., tariffs or domestic retail taxes) are often proposed in this policy space, I design my policy counterfactuals to show how they are traded off against each other in terms of their targeting and scale properties and how market power shapes this trade-off.

IV. Model

A. *Upstream Supply*

Every county i contains a continuum of fields indexed ω, each of which is owned by a landowner. In each period t, the landowner selects a land use from a discrete choice set consisting of a natural use option $\mathcal{N}$ and a nest of agricultural commodities $\mathcal{C}$. If the landowner allocates her field to a commodity $c \in \mathcal{C}$, she does so by renting the land to a farmer, who combines it with labor and intermediate inputs to produce the commodity.

1. Farmer's Optimization Problem

Production technology.—I model the farmer's optimization problem over factors by following the recent literature (Sotelo 2020; Farrokhi and Pellegrina 2023). Let $L_{it}^{c}(\omega)$ denote the size of the field ω that a farmer rents from a landowner to produce commodity c. The farmer combines land with hired workers $H_{it}^{c}(\omega)$ and an intermediate input $M_{it}^{c}(\omega)$ using a Cobb-Douglas technology that delivers $Q_{it}^{c}(\omega)$ physical units of commodity c:

$$Q_{it}^{c}(\omega) = \bar{q}[A_{it}^{c}(\omega)L_{it}^{c}(\omega)]^{\gamma_L^c}[H_{it}^{c}(\omega)]^{\gamma_H^c}[M_{it}^{c}(\omega)]^{\gamma_M^c}, \quad \text{where}$$
$$A_{it}^{c}(\omega) = A_{it}^{c}\exp(\varepsilon_{it}^{c}(\omega)).$$

The term $A_{it}^{c}(\omega)$ is land productivity, consisting of a county i mean A_{it}^{c} and a field ω idiosyncratic shock $\varepsilon_{it}^{c}(\omega)$; $\bar{q} \equiv (\gamma_L^c)^{-\gamma_L^c}(\gamma_H^c)^{-\gamma_H^c}(\gamma_M^c)^{-\gamma_M^c}$ is a scalar; and $\gamma_L^c + \gamma_H^c + \gamma_M^c = 1$.

[6] For empirical evidence on the challenges that Brazilian authorities face in monitoring farmers and the challenges that meat-packers face in tracing the origin (and carbon content) of their cattle purchases, see Barreto, Brandão, and Baima (2017), Pereira et al. (2020), and Skidmore et al. (2021). Appendix E, sec. E.1, provides further details on certification challenges in the beef supply chain.

Factor demand and unit cost function.—Let p_{it}^c denote the farm-gate price a farmer receives per unit of commodity c. On the cost side, let $r_{it}^c(\omega)$ denote the rent paid by the farmer for a unit of land on field ω, and let w_{Hit} and w_{Mit} denote the wage per worker and the price of the intermediate input. Taking output and factor prices as given, the farmer's cost-minimization problem delivers demand for labor and intermediates as

$$H_{it}^c(\omega) = \frac{r_{it}^c(\omega)}{w_{Hit}} \frac{\gamma_H^c}{\gamma_L^c}$$

and

$$M_{it}^c(\omega) = \frac{r_{it}^c(\omega)}{w_{Mit}} \frac{\gamma_M^c}{\gamma_L^c}.$$

Moreover, the unit cost of production on field ω is given by

$$\kappa_{it}^c(\omega) = (w_{Hit})^{\gamma_H^c} (w_{Mit})^{\gamma_M^c} \left(\frac{r_{it}^c(\omega)}{A_{it}^c(\omega)} \right)^{\gamma_L^c}.$$

2. Landowner's Land Use Choice Problem

Landowner payoff from each commodity.—Farmer profits are zero in equilibrium, resulting in a zero-profit condition $p_{it}^c = \kappa_{it}^c(\omega)$ from which we can back out the rent per unit of land on field ω—the landowner's payoff when her field is allocated to commodity c—as $r_{it}^c(\omega) = r_{it}^c \exp(\varepsilon_{it}^c(\omega))$, where $r_{it}^c \equiv A_{it}^c (p_{it}^c)^{1/\gamma_L^c} (w_{Hit})^{-\gamma_H^c/\gamma_L^c} (w_{Mit})^{-\gamma_M^c/\gamma_L^c}$ is the county-level nonidiosyncratic component of landowner payoffs.

Landowner payoff from the natural use option.—Let $A_{it}^{\mathcal{N}}(\omega) = A_{it}^{\mathcal{N}} \exp(\varepsilon_{it}^{\mathcal{N}}(\omega))$ be the landowner's payoff per unit of land when the field is left in its natural state, consisting of a county i mean $A_{it}^{\mathcal{N}}$ and a field ω idiosyncratic shock $\varepsilon_{it}^{\mathcal{N}}(\omega)$. The term $A_{it}^{\mathcal{N}}(\omega)$ captures the dollar value of any incentives landowners have to keep part of their land forested, some of which may be nonpecuniary (e.g., ecosystem services provided by forests that are privately valued by landowners, such as prevention of soil erosion and flood risk mitigation).

Observability of payoffs.—Concretely, the discrete choice problem of landowner ω is

$$\max_{\{1,\ldots,c,\ldots,C,\mathcal{N}\}} \left[r_{it}^1 \exp(\varepsilon_{it}^1(\omega)), \ldots, r_{it}^C \exp(\varepsilon_{it}^C(\omega)), A_{it}^{\mathcal{N}} \exp(\varepsilon_{it}^{\mathcal{N}}(\omega)) \right]. \quad (1)$$

Apart from the idiosyncratic shocks, the econometrician does not observe $A_{it}^{\mathcal{N}}$, so it will be estimated. I allow county-level productivity (in the r_{it}^c term) to be decomposed as $A_{it}^c = z_{it}^c \zeta_{it}^c$, where z_{it}^c is observable (as measured by agronomic yields) while ζ_{it}^c is not.

Nesting assumption.—I assume that the idiosyncratic shocks are distributed generalized extreme value (GEV; McFadden 1977; Goldberg 1995; Train 2009). This delivers a nested structure that allows shocks to be more correlated within nests than across. The GEV distribution has two key parameters: θ governs the variance of shocks, while $\lambda \in (0, 1)$ governs their correlation within nest $\mathcal{C}$. Higher values of θ correspond to lower variance, and higher values of λ correspond to lower correlation. Appendix A, sections A.1 and A.2, provide a discussion of how the textbook formulation of the GEV distribution maps to the land use model.

Share of land allocated to each land use.—Let $\pi_{it}^{c|\mathcal{C}}$ be the probability commodity c is chosen conditional on choosing nest $\mathcal{C}$, while $\pi_{it}^{\mathcal{C}}$ is the probability of choosing nest $\mathcal{C}$. Under our distributional assumptions, we have a closed form for these objects:

$$\pi_{it}^{c|\mathcal{C}} = \frac{\left(r_{it}^{c}\right)^{\theta/\lambda}}{\sum_{c' \in \mathcal{C}} \left(r_{it}^{c'}\right)^{\theta/\lambda}} \quad \text{and} \quad \pi_{it}^{\mathcal{C}} = \frac{\left(P_{it}^{\mathcal{C}}\right)^{\lambda}}{\left(A_{it}^{\mathcal{N}}\right)^{\theta} + \left(P_{it}^{\mathcal{C}}\right)^{\lambda}} \quad \text{with} \quad P_{it}^{\mathcal{C}} \equiv \sum_{c' \in \mathcal{C}} \left(r_{it}^{c'}\right)^{\theta/\lambda}. \tag{2}$$

The term $P_{it}^{\mathcal{C}}$ is defined as the payoff of the agricultural nest. Technically, $\ln P_{it}^{\mathcal{C}}$ is the nest's inclusive value in the nested logit model. The share of land allocated to natural use is $\pi_{it}^{\mathcal{N}} = 1 - \pi_{it}^{\mathcal{C}}$.

Quantity supplied.—Unconditional choice probabilities can be written as $\pi_{it}^{c} = \pi_{it}^{c|\mathcal{C}} \pi_{it}^{\mathcal{C}}$. Given a county's total surface $\bar{L}_i$, its total acreage of commodity c is $L_{it}^{c} = \pi_{it}^{c} \bar{L}_i$. Moreover, we have a closed form for the county-level supply of commodity c:

$$Q_{it}^{c} = \int_{\omega} Q_{it}^{c}(\omega)\, d\omega = \tilde{r}_{it}^{c} \left(\pi_{it}^{c|\mathcal{C}}\right)^{-\lambda/\theta} \left(\pi_{it}^{\mathcal{C}}\right)^{-1/\theta} L_{it}^{c}, \tag{3}$$

where $\tilde{r}_{it}^{c} \equiv r_{it}^{c} / p_{it}^{c} \gamma_{L}^{c}$. As $\lambda \to 1$, correlation between commodities disappears, and the nested model collapses to a multinomial model (Costinot, Donaldson, and Smith 2016; Sotelo 2020). The nested structure is important because a multinomial model restricts substitution between commodities to be just as easy as substitution between natural and agricultural use.

Demand and supply of nonland factors.—The farmer's optimization problem delivers demand for nonland factors at the field level. Integrating over fields delivers county-level aggregates: $H_{it}^{c} = \gamma_{H}^{c}(p_{it}^{c} Q_{it}^{c} / w_{Hit})$ and $M_{it}^{c} = \gamma_{M}^{c}(p_{it}^{c} Q_{it}^{c} / w_{Mit})$. As for the supply of nonland factors, I assume that labor is mobile across sectors but not across locations. Each location has an exogenous endowment of workers $\bar{H}_{it}$ that endogenously chooses between agriculture and nonagriculture. I assume that the worker's sectoral choice problem follows a Roy specification where the share of workers supplying labor to agriculture is $s_{H|it} = w_{Hit}^{\psi} / (w_{Hit}^{\psi} + \underline{w}_{Hit}^{\psi})$, where ψ is the across-sector substitution elasticity and $\underline{w}_{Hit}$ is the nonagricultural wage. Effective labor units supply to agriculture is therefore $s_{H|it}^{(\psi-1)/\psi} \bar{H}_{it}$. As for

the intermediate input, I assume that it is imported from abroad at price $w_{Mit} = w_{Mt}\tau_i$, where w_{Mt} is the international price and τ_i is a domestic trade cost to county i from its nearest port.[7]

B. *Downstream Demand*

The demand side consists of two layers. First, agribusiness intermediaries buy commodities from upstream farmers in source counties indexed $i \in \mathcal{I}$. Second, these intermediaries sell the commodity (either in raw format or partially processed) to final consumers in destinations indexed $j \in \mathcal{J}$. Intermediaries hold market power as buyers in the upstream market but take prices as given in the downstream consumer market.

Agribusiness intermediaries.—There are N_{it}^c identical agribusiness intermediary firms, each purchasing q_{it}^c units of raw commodity c from source i. Farmers do not perceive the firms as differentiated buyers; hence, all firms buy the commodity at the same farm-gate price p_{it}^c. Apart from transporting commodities from an upstream source i to a downstream destination j, I allow firms to add value by transforming q_{it}^c units of raw commodity c into a processed product k via a technology $f(q_{it}^c; k)$. For example, from a raw commodity c = soybeans, a firm could produce k = soybean oil or k' = soybean meal. In the case where the firm does not transform the commodity and trades it raw, I set $k = c$ so that the "product" is simply the raw commodity, and therefore $f(q_{it}^c; c) = q_{it}^c$.

The firm sells product k at the port closest to source county i, obtaining an FOB price $\bar{p}_{it}^k$. Hence, transport costs from port to final destination are paid by final consumers: a destination j consumer pays $p_{ijt}^k = \bar{p}_{it}^k \tau_{ij}$, where τ_{ij} is an iceberg trade cost. Taking demand of the other firms as given, we can now pose each firm's maximization problem:

$$\max_{q_{it}^c} \quad \bar{p}_{it}^k f(q_{it}^c; k) - p_{it}^c(Q_{it}^c) q_{it}^c,$$

where q_{it}^c is an individual firm's demand for raw commodity c, Q_{it}^c is total demand, and $p_{it}^c(Q_{it}^c)$ is source i's inverse supply equation. From the first-order conditions, the farm-gate price p_{it}^c—which is the marginal cost from the intermediary's point of view—is a fraction $\mu_{it}^c < 1$ of the marginal revenue the commodity generates for the intermediary:

[7] Appendix D, sec. D.5.3, provides an extension with migration across counties and discusses implications for the main policy analysis. The assumption that the intermediate is not produced domestically is motivated by the import share of agricultural inputs being especially high in Latin America. The region has one of the world's highest import shares of fertilizers (76%) as well as machinery (68%), significantly higher than North America, Europe, and Asia and surpassed only by sub-Saharan Africa (Farrokhi and Pellegrina 2023).

$$\underbrace{p_{it}^{c}}_{\text{marginal cost}} = \underbrace{\mu_{it}^{c}}_{\text{markdown}} \times \underbrace{\bar{p}_{it}^{k} f'(q_{it}^{c}; k)}_{\text{marginal revenue}}, \quad \text{where} \quad \mu_{it}^{c} \equiv \left(1 + \frac{\partial \ln p_{it}^{c}}{\partial \ln Q_{it}^{c}} \frac{1}{N_{it}^{c}}\right)^{-1} < 1. \tag{4}$$

A farmer from source i obtains μ_{it}^{c} cents for every dollar the intermediary makes from transforming commodity c into product k. I define μ_{it}^{c}, the ratio of the raw commodity's farm-gate price to its marginal revenue, as the markdown. Markdowns follow an inverse-elasticity rule: farmers with inelastic supply are subject to wide markdowns (low μ_{it}^{c}). The setup of the intermediary problem is purposefully simple, with the goal of obtaining the smallest departure from the perfectly competitive setting typically assumed by the agricultural trade literature as well as parsimoniously nest it. In the limiting cases of perfect competition ($\mu_{it}^{c} = 1$) and raw commodity trade ($f(q_{it}^{c}, c) = q_{it}^{c}$), the farm-gate price is equal to the FOB price, and the destination market price is simply the farm-gate price adjusted by trade costs (i.e., $p_{ijt}^{c} = p_{it}^{c}\tau_{ij}$). I present model extensions with firm heterogeneity and exit/entry in appendix A, section A.4.

Consumers.—I follow the standard approach from the agricultural trade literature to model consumers: a multilevel constant elasticity of substitution (CES) demand system with differentiation across products and sources and bilateral iceberg trade costs τ_{ij}. Each destination j has a representative consumer with a three-level CES utility function. In the upper level, consumers substitute between products indexed k (e.g., maize vs. wheat). In the middle level, they substitute between source nations n of a given commodity (e.g., Brazilian maize vs. US maize). In the lower level, they substitute between counties i within a nation (e.g., maize from northern vs. southern Brazil).[8] These preferences deliver the following demand by a destination j consumer for product k from county i in nation n:

$$C_{ijt}^{k} = a_{ijt}^{k}\left(\frac{p_{ijt}^{k}}{P_{njt}^{k}}\right)^{-\eta_{\mathrm{l}}} a_{njt}^{k}\left(\frac{P_{njt}^{k}}{P_{jt}^{k}}\right)^{-\eta_{\mathrm{m}}} a_{jt}^{k}\left(\frac{P_{jt}^{k}}{P_{jt}}\right)^{-\eta_{\mathrm{u}}} \frac{X_{jt}}{P_{jt}} \quad \forall\ i \in n, \tag{5}$$

where η_{l}, η_{m}, and η_{u} are the lower-, middle-, and upper-level elasticities of substitution, p_{ijt}^{k} is the lower-level price, the a terms are preference shifters,

[8] Consumers in this model are interpreted as the first agents to receive the commodity at the destination's entry point (typically food-processing companies that transform commodities into retail food products). Hence, the different levels of the CES system should not be interpreted as the degree to which final retail consumers literally differentiate as a matter of taste—indeed, it is unlikely that they perceive significant quality differences between maize from one county vs. another. Instead, imperfect substitution in the lower level may reflect the degree to which there are frictions in sourcing from one county vs. another (e.g., unobserved switching costs across suppliers) even if the underlying product being sourced from both counties is identical. This allows a finite substitution elasticity at the lower level while allowing the interpretation that consumers do not literally differentiate products across origins.

and X_{jt} is destination j expenditure on agricultural goods (which is exogenous since this is a single-industry model). The P terms are price indexes at each level of the demand system: $P^k_{njt} \equiv (\Sigma_{i' \in n} a^k_{i'jt}(p^k_{i'jt})^{1-\eta_h})^{1/(1-\eta_h)}$, $P^k_{jt} \equiv (\Sigma_n a^k_{njt}(P^k_{njt})^{1-\eta_m})^{1/(1-\eta_m)}$, and $P_{jt} \equiv (\Sigma_c a^k_{jt}(P^k_{jt})^{1-\eta_u})^{1/(1-\eta_u)}$. Given that this is a standard CES consumer problem, its derivation is relegated to appendix A, section A.3.

C. *Equilibrium*

An equilibrium is a vector of farm-gate prices $\{p^c_{it}\}_{i,c}$ and wages $\{w_{Hit}\}_i$ such that

1. commodity supply in each county equals total demand: $Q^c_{it}(p^c_{it}) = \Sigma_j C^c_{ijt}(p^k_{ijt})\tau_{ij} \ \forall \ i, c$, where

$$p^k_{ijt} = \frac{p^c_{it}}{\mu^c_{it}} \frac{\tau_{ij}}{f'(q^c_{it}; k)}.$$ [9]

2. agricultural labor supply in each county equals local labor demand across all commodities: $s_{H|it}(w_{Hit}, \underline{w}_{Hit})^{(\psi-1)/\psi} \bar{H}_{it} = \Sigma_c H^c_{it}(w_{Hit}, w_{Mit}, p^c_{it})$.

I solve the equilibrium using an iterative algorithm that exploits the monotonicity of the excess demand function in output and factor markets (see app. D, sec. D.2 for details).

Emissions externality.—I allow emissions to originate from two sources: (i) land use change (LUC) emissions from clearing forested land and (ii) non-LUC emissions that result from the on-farm production process (e.g., fertilizer use, enteric methane from cattle). Non-LUC emissions vary by commodity; hence, switching between commodities affects total emissions even if there is no deforestation. Because agricultural production releases many types of greenhouse gases (CO_2, CH_4, etc.), I use CO_2e throughout the analysis. Let $e^{c,\text{NLUC}}_i$ denote the non-LUC emissions intensity of commodity c (tonnes of non-LUC CO_2e per tonne of output). Let e^{LUC}_i denote the LUC emissions intensity of location i (CO_2e stock per hectare of land). Values for $e^{c,\text{NLUC}}_i$ are taken from the life cycle assessment

[9] When k is not a raw commodity, demand C^k_{ijt} from eq. (5) is in units of processed product k, but supply Q^c_{it} is in units of raw commodity c. To define market clearing, we convert demand to raw commodity-equivalent units, C^c_{ijt}. Let m^{ck}_{it} denote the units of product k produced with 1 unit of commodity c ($m^{cc}_{it} = 1$) and let $\mathcal{K}(c)$ be the set of products produced from c. Then, demand expressed in units of raw commodity c is

$$C^c_{ijt}(p^k_{ijt}) = \begin{cases} C^k_{ijt}(p^k_{ijt}) & \text{if } \mathcal{K}(c) = \{c\}, \\ \sum_{k \in \mathcal{K}(c)} C^k_{ijt}(p^k_{ijt})/m^{ck}_{it} & \text{if } \mathcal{K}(c) = \{c\} \cup \{k, k+1, \ldots\}. \end{cases} \quad (6)$$

literature (Poore and Nemecek 2018), while values for e_i^{LUC} are from the carbon stock maps described in section II. I denote emissions from LUC and non-LUC sources in location i as E_i^{LUC} and $E_i^{c,\text{NLUC}}$. The change in location i's emissions between two counterfactual equilibria is computed from changes in land and output as $\Delta E_i = \Delta E_i^{\text{LUC}} + \Sigma_c \Delta E_i^{c,\text{NLUC}}$, where $\Delta E_i^{\text{LUC}} = \Delta L_i^c e_i^{\text{LUC}}$ and $\Delta E_i^{c,\text{NLUC}} = \Delta Q_i^c e_i^{c,\text{NLUC}}$. The dollar value of emissions is obtained by multiplying emissions quantities by the social cost of carbon (SCC), for which I use a baseline value of US$50 per tonne of CO_2e.

V. Estimation

A. *Supply Elasticities*

Within-nest substitution elasticities.—To clarify the variation I use to estimate the supply-side parameters, consider the odds ratio between two land use choices c and $c' \in \mathcal{C}$:

$$\ln\left(\frac{\pi_{it}^c}{\pi_{it}^{c'}}\right) = \frac{\theta}{\lambda}\ln\left(\frac{r_{it}^c}{r_{it}^{c'}}\right) = \frac{\theta}{\lambda}\ln\left(\frac{(p_{it}^c)^{1/\gamma_L^c} z_{it}^c}{(p_{it}^{c'})^{1/\gamma_L^{c'}} z_{it}^{c'}}\right) + u_{it}^{cc'}, \tag{7}$$

where the second equality is obtained by substituting the expression for r_{it}^c from section IV.A.2. The first term on the right of equation (7) is a combination of objects that are observable to the econometrician: farm-gate prices p_{it}^c, land intensities γ_L^c, and the observable component of land productivity z_{it}^c (yields). The second term is a combination of unobservables defined as

$$u_{it}^{cc'} \equiv \frac{\theta}{\lambda}\left[\ln\left(\frac{\zeta_{it}^c}{\zeta_{it}^{c'}}\right) + \ln\left(w_{Hit}^{\gamma_{HL}^{cc'}} w_{Mit}^{\gamma_{ML}^{cc'}}\right)\right],$$

where

$$\gamma_{HL}^{cc'} \equiv -\frac{\gamma_H^c}{\gamma_L^c} + \frac{\gamma_H^{c'}}{\gamma_L^{c'}}$$

and

$$\gamma_{ML}^{cc'} \equiv -\frac{\gamma_M^c}{\gamma_L^c} + \frac{\gamma_M^{c'}}{\gamma_L^{c'}}.$$

This term contains unobservable shocks to land productivity ζ_{it}^c and unobserved factor prices w_{Hit}, w_{Mit}. Hence, $u_{it}^{cc'}$ is a combination of unobservable supply shocks to c relative to c'. Ordinary least squares (OLS) estimates of equation (7) are therefore subject to simultaneity bias, for which we need a demand shifter as an instrument.

The parameter of interest in equation (7) is θ/λ: the elasticity of substitution between commodities within the agricultural nest. This elasticity is large when productivity dispersion across fields goes to zero ($\theta \to \infty$) or when productivity is perfectly correlated across commodities ($\lambda \to 0$). When $\theta \to \infty$, there is no field heterogeneity remaining: all fields are on the margin when prices change, so the county-level within-nest elasticity is high. When $\lambda \to 0$, the idiosyncratic shocks become perfectly correlated within nest $\mathcal{C}$, which lowers the variance of their differences, $\varepsilon_i^c(\omega) - \varepsilon_i^{c'}(\omega)$. Because these differences drive the idiosyncratic component of choices, their lower variance implies less heterogeneity for the within-nest choice margin and thus a high within-nest elasticity.

Across-nest substitution elasticities.—Consider the odds ratio between nest $\mathcal{C}$ and $\mathcal{N}$,

$$\ln\left(\frac{\pi_{it}^{\mathcal{C}}}{\pi_{it}^{\mathcal{N}}}\right) = \lambda \ln\left(P_{it}^{\mathcal{C}}\right) + u_{it}^{\mathcal{N}}, \quad \text{with } P_{it}^{\mathcal{C}} \equiv \sum_{c \in \mathcal{C}} \left(r_{it}^{c}\right)^{\theta/\lambda} \quad \text{and} \quad u_{it}^{\mathcal{N}} \equiv -\theta \ln\left(A_{it}^{\mathcal{N}}\right). \tag{8}$$

Note that $\ln P_{it}^{\mathcal{C}}$ is the inclusive value of the agricultural nest and $u_{it}^{\mathcal{N}}$ is an unobservable supply shifter of agricultural land relative to forested land, which maps to the (unobserved) payoff to natural use $A_{it}^{\mathcal{N}}$. The parameter of interest is λ, the substitution elasticity across nests (i.e., the deforestation elasticity). As $\lambda \to 0$, the within-nest heterogeneity falls, as explained in the previous paragraph. Therefore, fields become relatively more heterogeneous in the across-nest margin, leading to a low across-nest elasticity. Given θ/λ from equation (7), we can construct the inclusive value in equation (8), and λ is identified (see app. C, sec. C.2 for details).

Estimation sample.—I use county-level data from agricultural censuses of Brazil (1995, 2006, 2017) and Argentina (2002, 2008, 2018). The commodities considered are beef (pasture), soybean, maize, wheat, rice, sunflower, and sugarcane. Beef, maize, and soybean represent 85% of all agricultural land, while the other commodities are jointly less than 10%.

Factor shares.—I take factor shares from the literature, since they have been estimated for South American countries (Sotelo 2020; Pellegrina 2022). Land intensities for major crops (maize, soybean, wheat) range from 0.4 to 0.6, while values for beef are between 0.5 and 0.7.

Instruments.—OLS estimates of equation (7) will be biased toward zero owing to simultaneity bias, because the unobservable supply shocks will be correlated with relative land shares and relative payoffs. The appropriate instrument is a demand shifter varying at the county-year it and commodity-pair cc' level. To construct such an instrumental variable (IV), I first define a demand shifter for commodity c as $\text{IV}_{it}^{c} = \sum_j s_{ij}^{c} d_{jt}^{c}$, where s_{ij}^{c} is the share of commodity c output from county i that historically goes to destination j

and d_{jt}^{c} is a time-varying measure of demand conditions in destination j.[10] Intuitively, if demand conditions for crop c increase in destination j, the counties that historically supplied j are more exposed and receive a larger demand shock. The relevance of the instrument relies on persistence in the sourcing relationships between downstream consumer markets and upstream producers, as mediated by the agribusiness firms, which is ultimately confirmed via the first-stage F-statistic in the results. Finally, I define a relative demand shifter between two commodities c and c' as $\text{IV}_{it}^{cc'} \equiv \text{IV}_{it}^{c}/\text{IV}_{it}^{c'}$. The identifying assumption is that the county's relative exposure is uncorrelated with changes in the error term, $\Delta u_{it}^{cc'}$, whereas correlation with levels $u_{it}^{cc'}$ is allowed. We also need an instrument for equation (8) for the same reasons as for equation (7): if the unobserved payoff to natural use increases, then the commodity nest's share would fall, and its price index $P_{it}^{\mathcal{C}}$ would increase, biasing the estimate of λ toward zero. We now need a shifter of demand for agriculture overall, which we construct by aggregating the demand shifters across all commodities, $\text{IV}_{it}^{\mathcal{C}} \equiv \Sigma_{c \in \mathcal{C}} \text{IV}_{it}^{c}$.

Frontier region definition.—I allow heterogeneity in substitution parameters across "frontier" and "core" agricultural regions. I do this to allow my land use change estimates to be comparable to the subset of studies focusing exclusively on the Amazon region. Hence, I define the frontier region for Brazil as the government-designated area of Amazonia Legal. The definition of core region for Argentina is based on the *zona núcleo* designation from agricultural censuses, which consists of the Pampeana region.

Results.—Panel A of table 3 shows OLS and IV estimates for θ/λ, the substitution elasticity between commodities. Panel B shows estimates for λ, the deforestation elasticity. Note that the within-nest estimates are in line with values from the agricultural trade literature, which typically finds substitution elasticities across commodities ranging between 1.5 and 4 (Costinot, Donaldson, and Smith 2016; Sotelo 2020; Pellegrina 2022). Columns 3 and 4 add interactions with a frontier region indicator to allow for spatial heterogeneity: the negative sign suggests that farmers in frontier regions are less responsive to price changes along both within- and across-nest margins.

While I cannot fully disentangle the microfoundations driving the regional heterogeneity, most mechanisms are likely related to frontier regions being less developed overall. Farmers in these regions may find it harder to switch between commodities owing to poorly functioning input

[10] For the construction of the instruments, the destination markets are all those in the Trase data, more than 190 countries from all continents that are aggregated into eq. (7) trading blocs: Asia, Europe, Africa, Middle East, Central America, South America, and North America. I use destination j's imports from every nation except Argentina and Brazil as the demand measure d_{jt}^{c}, thus purging supply-side effects in Argentina and Brazil that directly affect the imports of j.

TABLE 3
SUPPLY-SIDE SUBSTITUTION PARAMETERS

	OLS (1)	IV (2)	OLS (3)	IV (4)
	A. Within-Nest Substitution Elasticity: Estimate for Equation (7)			
θ/λ	.201***	2.220***	.164***	2.219***
	(.034)	(.144)	(.035)	(.143)
$\theta/\lambda \times$ frontier region			.725***	−1.106***
			(.123)	(.180)
Time fixed effect	✓	✓	✓	✓
Location fixed effect	✓	✓	✓	✓
Observations	7,008	7,008	7,008	7,008
Kleibergen-Paap Wald first-stage *F*-statistic		171.2		176.6
	B. Across-Nest Substitution Elasticity: Estimate for Equation (8)			
λ	.052***	.279***	.052***	.281***
	(.009)	(.068)	(.010)	(.069)
$\lambda \times$ frontier region			.002	−.229**
			(.031)	(.076)
Time fixed effect	✓	✓	✓	✓
Location fixed effect	✓	✓	✓	✓
Observations	9,255	9,255	9,255	9,255
Kleibergen-Paap Wald first-stage *F*-statistic		143.4		151.4

NOTE.—Standard errors are clustered at the county level.
** $p < .05$.
*** $p < .01$.

markets (machinery, credit), thus rationalizing lower within-nest substitution. The lower across-nest response to commodity prices is consistent with deforestation in frontier areas being driven relatively more by nonmarket forces, most notably land-grabbing incentives.

B. *Markdowns*

As indicated by equation (4), markdowns μ_{it}^{c} are a function of the price elasticity of supply and the number of firms sourcing from each location. Hence, we can compute markdowns before imposing the demand side of the model.[11] Figure 3*A* shows the geography of markdowns for the beef cattle sector, the main driver of land use. Average markdowns are 0.74 for beef and 0.86 and 0.87 for maize and soybeans. Markdowns are wider

[11] The price elasticity of supply is derived from eq. (3) and is computed with the following objects: substitution parameters (θ, λ), land intensity parameters (γ_L^c), and observed land use shares ($\pi_{it}^{c|\mathcal{C}}$, $\pi_{it}^{\mathcal{C}}$) as

$$\frac{\partial \ln Q_{it}^{c}}{\partial \ln p_{it}^{c}} = \left[\left(\frac{\theta}{\lambda} - 1\right)\left(1 - \pi_{it}^{c|\mathcal{C}}\right) + (\theta - 1)\pi_{it}^{c|\mathcal{C}}\left(1 - \pi_{it}^{\mathcal{C}}\right)\right]\frac{1}{\gamma_L^c} + \left(\frac{1}{\gamma_L^c} - 1\right). \quad (9)$$

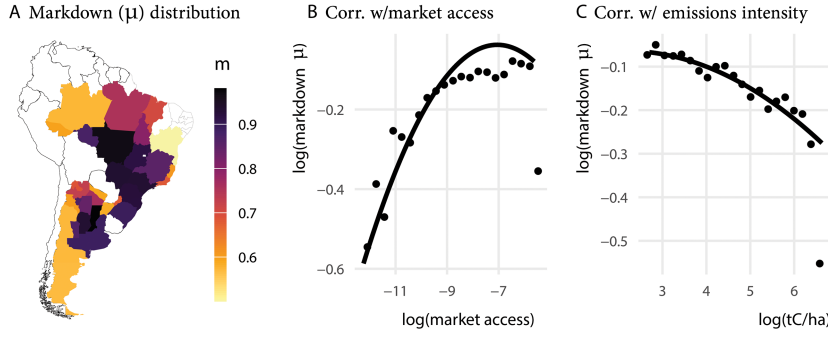

FIG. 3.—Spatial distribution of markdowns and correlation with local observables. All plots are constructed from county-level markdown estimates, market access terms, and carbon intensities. Markdowns are computed using equations (4) and (9) and averaged over time. Market access is constructed as $\Sigma_j s_j d_{ij}^{-1}$, where d_{ij} is distance between county i and hub j and s_j is hub trade volume. Carbon intensity is measured as tonnes carbon per hectare. For the scatterplots, curves are fitted polynomials of county-level observations, and markers are binned scatterplots. All results are for upstream beef markets. Additional results for soybean and maize are displayed in appendix C, section C.3.2.

in frontier regions where supply is least elastic (Brazil's north, Argentina's west), likely because of the reasons explained in the preceding section.

Panels B and C of figure 3 show that frontier regions are indeed more remote (as measured by market access) and more emissions intense (as measured by the carbon density of land). Most importantly, the positive correlation between the two market distortions—market power and the environmental externality—will play a significant role in the policy analysis of section VI. Finally, in appendix C, section C.3.3, I validate the model-implied markdowns against external accounting-based markdowns. The key takeaway is that model-implied markdowns are positively correlated across space with external markdowns. If external markdowns are to some degree indicative of market power, then the fact that the model-implied markdowns move in the same direction across space is reassuring.

C. *Demand Elasticities*

First, the lower-level elasticity η_1 is identified from expenditure variation across source origins i within a source nation n:

$$\ln\left(\frac{X_{ijt}^c}{X_{njt}^c}\right) = (1 - \eta_1)\ln\left(p_{ijt}^c\right) + \lambda_{njt}^c + \varepsilon_{ijt}^c \quad \forall\ i \in n, \tag{10}$$

where X_{ijt}^c is destination j's expenditure on commodity c from county i, $X_{njt}^c = \Sigma_{i' \in n} X_{i'jt}^c$, $\lambda_{njt}^c \equiv -\ln(\Sigma_{i' \in n} a_{i'jt}^c (p_{i'jt}^c)^{1-\eta_1})$ is a commodity-source nation-destination-time fixed effect, and $\varepsilon_{ijt}^c \equiv \ln(a_{ijt}^c)$. Because equation (10) is a

demand equation, simultaneity bias arises when estimating via OLS. Therefore, I instrument for price with a supply shifter that I construct as $IV^c_{ijt} \equiv s^c_{ij} \times b_{nt}$, where b_{nt} is a weather shock at the origin nation (measured as deviations from historical average) and s^c_{ij} is the share of origin i production going to destination j in the first year of the sample. The instrument's relevance condition is straightforward: given a negative weather shock at the origin, the size of the supply shock faced by destination j depends on how exposed it is to i through its historic trading relationship. The exclusion restriction is that the origin-destination shares s^c_{ij} are not predictive of changes in unobservable demand shocks. This allows an origin i that historically exports most of its output to j to consistently experience large demand shocks from j but not to experience systematic changes in such shocks over time. The middle-level elasticity η_{m} is identified from expenditure variation across nations:

$$\ln\left(\frac{X^c_{njt}}{X^c_{jt}}\right) = (1 - \eta_{\mathrm{m}}) \ln\left(P^c_{njt}\right) + \lambda^c_{jt} + \varepsilon^c_{njt}, \tag{11}$$

where X^c_{jt} is destination j's total expenditure on commodity c across all source nations. The term $\lambda^c_{jt} \equiv -\ln(\Sigma_{n'} a^c_{n'jt}(P^c_{n'jt})^{1-\eta_{\mathrm{m}}})$ is a commodity-destination-time fixed effect, and $\varepsilon^c_{njt} \equiv \ln(a^c_{njt})$. To construct the price indexes in equation (11), I use the residuals from equation (10), that is, $P^c_{njt} \equiv (\Sigma_{i \in n} \hat{a}^c_{ijt}(p^c_{ijt})^{1-\eta_{\mathrm{i}}})^{1/(1-\eta_{\mathrm{i}})}$, where $\hat{a}^c_{ijt} = \exp(\hat{\varepsilon}^c_{ijt})$. Since equation (11) is a demand equation just like equation (10), the simultaneity problem and its solution are the same, with the only difference being the lower geographic resolution (origin locations are now nations instead of counties). Hence, I again instrument for price by constructing a supply shifter but at the origin nation level. Finally, the upper-level elasticity η_{u} is identified from expenditure variation across commodities:

$$\ln\left(\frac{X^c_{jt}}{X_{jt}}\right) = (1 - \eta_{\mathrm{u}}) \ln\left(P^c_{jt}\right) + \lambda_{jt} + \varepsilon^c_{jt}, \tag{12}$$

where X_{jt} is destination j's total expenditure on agricultural imports. To construct the price indexes in equation (12), I use the residuals from equation (11), that is, $P^c_{jt} \equiv (\Sigma_n \hat{a}^c_{njt}(P^c_{njt})^{1-\eta_{\mathrm{m}}})^{1/(1-\eta_{\mathrm{m}})}$, where $\hat{a}^c_{njt} = \exp(\hat{\varepsilon}^c_{njt})$. Furthermore, $\lambda_{jt} \equiv -\ln(\Sigma_{c'} a^{c'}_{jt}(P^{c'}_{jt})^{1-\eta_{\mathrm{u}}})$ is a destination-time fixed effect, and $\hat{a}^c_{jt} = \exp(\varepsilon^c_{jt})$. The required instrument is now a supply shifter varying across commodities, which I construct as supply shocks at the destination-commodity level as $IV^c_{jt} = \Sigma_n s^c_{nj} b_{nt}$, where s^c_{nj} is the share of destination j imports coming from an origin nation n in a baseline year and b_{nt} is the origin nation-level weather shock.

Results.—Table 4 shows the substitution elasticity estimates for each level of the demand system. At all levels, IV estimates are larger in absolute value than OLS estimates, consistent with simultaneity bias. Substitutability

TABLE 4
DEMAND SUBSTITUTION ELASTICITIES, DEPENDENT VARIABLE: LOG (expenditure share)

	Lower Level (across counties)		Middle Level (across nations)		Upper Level (across commodities)	
	OLS	IV	OLS	IV	OLS	IV
log (price)	1.266***	−8.641***	−1.478***	−4.992***	−1.030***	−4.107**
	(.020)	(.691)	(.079)	(.772)	(.163)	(1.173)
Kleibergen-Paap Wald first-stage *F*-statistic		257		32		9
Observations	409,558	409,558	38,272	38,272	4,059	4,059

NOTE.—Lower, middle, and upper specifications include origin nation–destination–year–commodity fixed effects, destination-year-commodity fixed effects, and destination-year fixed effects, respectively. Standard errors are clustered by origin-destination (for lower and middle levels) or destination (for upper level).

** $p < .05$.

*** $p < .01$.

is higher across commodity sources than across commodities and across counties than across nations. At the nation and commodity levels, the implied CES parameter values are $\eta_{\mathrm{m}} = 5.99$ and $\eta_{\mathrm{u}} = 5.11$. These magnitudes are similar to those from related literature (Costinot, Donaldson, and Smith 2016), with deviations likely driven by differences in the set of commodities and countries.

Additional parameters and simulation details.—The remaining parameters are observed in the data (land endowment $\bar{L}_i$, population $\bar{H}_i$, nonagricultural wages $\underline{w}_{Hit}$, intermediate input prices w_{Mt}), taken from the literature (factor shares γ_L, γ_H, γ_M, across-sector labor supply elasticity ψ), or calibrated (unobservable land productivity ζ_{it}^{c}, trade costs τ_{ij}). Implementation details for these parameters, as well as for equilibrium simulations, can be found in appendix D, section D.1.

VI. Policy Counterfactuals

The counterfactual analysis proceeds in three steps. Section VI.A evaluates three stylized policies often proposed in the public debate in order to highlight the regulatory trade-off between targeting and scale. Section VI.B reruns the counterfactuals under different conduct assumptions to show how the trade-off is shaped by competition. Section VI.C compares all policies in a full welfare analysis where the first best is used as a benchmark.

A. *Policy Trade-Off between Targeting and Scale*

1. Downstream Taxes: Large-Scale but Poor Upstream Targeting

This first policy counterfactual is a uniform downstream tax for each commodity (based on each commodity's average emissions footprint across

upstream origins).[12] The motivation for this exercise is based on the notion that a commodity's upstream origin cannot be perfectly traced. Hence, this stylized counterfactual captures how well a regulator can do when constrained to commodity-specific but spatially uniform downstream taxes.

Figure 4*A* shows the impact of the downstream tax on upstream production (for the beef sector, the main driver of total emissions).[13] The figure plots the relationship between an upstream location's emissions intensity and its change in production: a negative relationship implies better targeting. First, the average impact on upstream production is large even when using a modest SCC of US$50 per tonne of CO_2e for the downstream tax. This is consistent with existing work, which finds that carbon taxes at even lower SCC values would nearly eliminate all agricultural land in the Amazon (Souza-Rodrigues 2019). Second, while the policy has a large impact on average, it is bluntly targeted: output drops least in frontier regions where emissions intensities are highest. Mistargeting occurs because (i) supply is less elastic in frontier regions and (ii) pass-through is lower in these regions because of the stronger degree of agribusiness market power. While this spatial mistargeting result holds regardless of conduct assumptions because it relies solely on the spatial pattern of supply elasticities, market power exacerbates it by eroding pass-through most in locations where the environmental cost is highest. Beyond the limited effectiveness of the downstream tax, I also analyze its distributional effects. In appendix D, section D.5.1, I show that farmer income is implicitly taxed at a higher rate in poorer regions (i.e., the policy is regressive across space). Hence, apart from increasing food prices for consumers, the downstream tax has an extra layer of regressivity on the supply side.

2. Upstream Regulation: Effective Upstream Targeting but Low Scale

Recall that the A_i^N term in the landowner choice problem is a residual: it captures any incentives farmers have to deforest that are not explained by

[12] I implement the policy as a downstream retail tax (for goods consumed domestically) combined with a downstream export tax (on exported goods). Technically, it is implemented by setting an output tax $t_{ij}^c = t^c = \text{SCC} \times E^c \ \forall i, j, c$, where E^c is the average tonnes of emissions per tonne of output for commodity c. The reason for having both a domestic tax and an export tax is to shut down standard inefficiencies from consumption-side leakage: without the domestic tax, there would be consumption leakage to unregulated domestic consumers. Hence, fully regulating all consumer markets isolates inefficiencies from domestic mistargeting, which are the main focus of the paper. Appendix D, sec. D.5.2, shows a version of this counterfactual where only a foreign trade partner implements this tax, which delivers similar issues of upstream targeting while also introducing consumption-side leakage (to domestic and foreign unregulated consumer markets).

[13] I focus on beef because it accounts for 70% of all agricultural land, and its emissions footprint is orders of magnitude above that of crops. As a result, it is the main driver of aggregate emissions. I account for emissions from crops when discussing aggregate emissions in the welfare analysis of sec. VI.C.

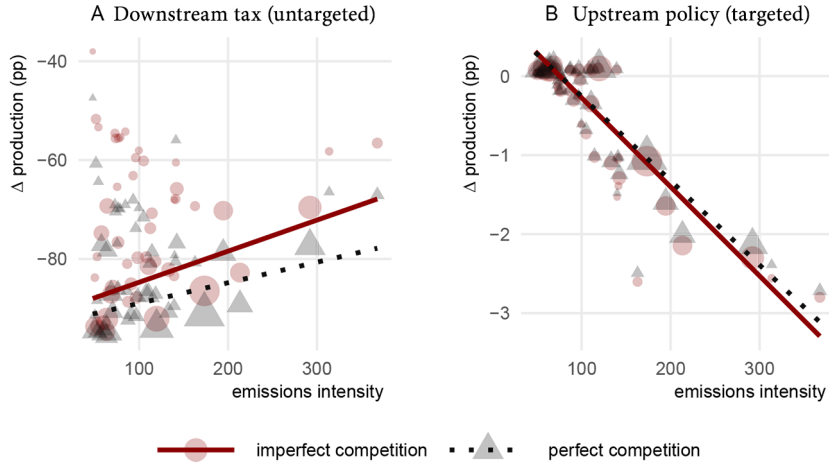

FIG. 4.—Targeting performance of downstream and upstream regulation. Scatterplots are constructed from upstream outcomes averaged to the state level and for the beef sector. Marker sizes are proportional to baseline production. Lines indicate the fit of a linear model to the markers, with observations size weighted by baseline production. The horizontal axis reports the emissions intensity of each upstream location (tonnes of CO_2e per tonne of output). The vertical axes report the percentage point change in upstream output as a result of each policy. Outcomes are reported for the default conduct assumption of imperfect competition and perfect competition.

changes in commodity prices. Hence, a change in $A_i^{\mathcal{N}}$ is consistent with any policy that is enforced upstream and changes incentives to deforest directly, rather than indirectly via commodity prices. Technically, I implement this section's policy as a forest subsidy that increases the $A_i^{\mathcal{N}}$ term for upstream locations in the top decile of emissions intensity. This differs from the previous downstream taxes in that it is directly implemented upstream on farmers. To ensure that farmers do not just claim the subsidy without reforesting, I assume that the targeted benefits of this policy come at an enforcement cost that limits its scalability.[14] Figure 4*B* plots the policy's impact on upstream production and its spatial correlation with emissions intensity. The contrasting slopes between panels A and B of figure 4 indicate that the upstream policy is well targeted, while the downstream tax is not. More subtly, the outcomes delivered by the upstream

[14] The enforcement cost depends on the size of the conservation zone. For the purpose of this stylized counterfactual, the enforcement cost is thus contingent on the conservation zone being defined as counties in the top 10% of emissions intensity (which is based on typically observed coverage rates for priority-zone policies in this setting). The subsidy rate is set based on the average carbon intensity of land: a landowner in location i is paid SCC $\times$ e_i^{LUC} dollars per hectare of forest, where e_i^{LUC} is the location's CO_2e per hectare of land and SCC is the social cost of carbon.

policy are similar with or without market power, which is not true for the downstream tax. Because of its direct implementation, the upstream policy avoids the pass-through distortions that market power introduces for the downstream tax.

3. Downstream Taxes with Certification: Improving Targeting While Preserving Scale

The previous two policies were at two extremes in terms of targeting assumptions: the upstream regulation was perfectly targeted but at small scale, while the downstream tax operated at large scale but without any targeting. We now consider an intermediate policy: a downstream tax that preserves scale while allowing coarse targeting at the macroregion level. Given that there are only 10 of these large macroregions, such a policy would need to certify commodities only by their coarse macroregion origin rather than their precise county origin.[15] How well this coarse policy performs relative to finer county-level targeting depends on how much of the variation in market power and emissions intensities is across regions relative to within.

Figure 5 plots the correlation between an upstream location's emissions intensity and its counterfactual change in production for (i) the uniform downstream tax, (ii) a macroregion targeted tax, and (iii) a county-level targeted tax. Macroregion targeting flips (and corrects) the correlation between emissions intensity and abatement. While county-level targeting naturally performs even better, the key takeaway is that even coarse macroregion certification makes a qualitative difference for targeting.

B. Role of Competition for Policy Trade-Off

The previous section suggests that when choosing among two policy alternatives, it is natural to compare the abatement benefits each tool delivers relative to its implementation cost. Let ΔE_{U} and ΔE_{D} denote the emissions reductions (in tonnes of CO_2e) achieved by the upstream regulation and the downstream tax, respectively. Moreover, define $\mathcal{T} \equiv \Delta E_{\mathrm{U}} - \Delta E_{\mathrm{D}}$ as the "targeting premium": the additional emissions reductions achieved by the upstream tool (targeted but small scale) relative to the downstream tool (large scale but nontargeted). Finally, let C denote the additional enforcement cost of the upstream tool relative to the downstream one. A regulator with the goal of maximizing abatement at minimal enforcement cost

[15] I implement this with an output tax $t^c_{ij} = t^c_{r(i)} = \mathrm{SCC} \times E^c_{r(i)} \; \forall \; i, j, c$, where $E^c_{r(i)}$ is average tonnes of emissions per tonne of output in the macroregion r that county i belongs to. Hence, the tax varies by commodity and macroregion but not within a macroregion.

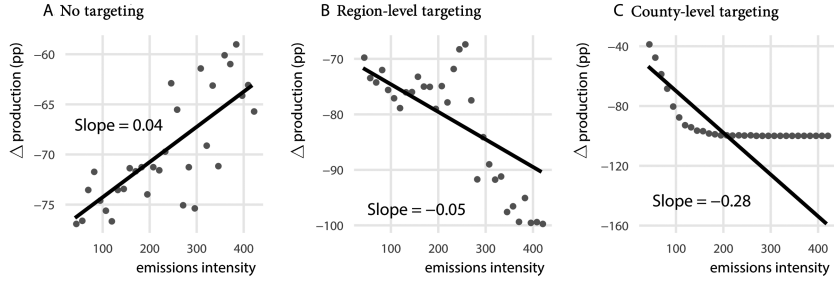

FIG. 5.—Coarse targeting flips the correlation between CO_2e intensity and abatement. All plots are constructed from county-level outcomes for the beef sector. Markers are binned scatterplots, and lines indicate fit to county-level outcomes. Horizontal axes report the emissions intensity of each upstream location (tonnes of CO_2e per tonne of output). Vertical axes report percentage point change in upstream output as a result of each policy. From left to right, results are shown for downstream taxes that are untargeted (*A*), coarsely targeted by upstream region (*B*), and finely targeted by upstream county (*C*). Outcomes are reported for the default conduct assumption of imperfect competition.

chooses the upstream tool if its dollar value of emissions reductions, net of additional enforcement costs, exceeds that of the downstream tool:

$$\text{SCC} \times \Delta E_{\text{U}} - C > \text{SCC} \times \Delta E_{\text{D}} \Rightarrow C < \text{SCC} \times \mathcal{T}. \tag{13}$$

The right-hand side of equation (13) says that the regulator chooses the targeted tool if its additional enforcement cost (C) is below a certain threshold: the dollar value of the targeting premium, $\text{SCC} \times \mathcal{T}$. SCC is observable given that it is a parameter that we incorporate exogenously into our analysis, and the targeting premium $\mathcal{T}$ is also observable because it is an outcome of the model simulations. Hence, while we do not observe enforcement costs C, we can compute an upper bound under which small-scale upstream regulation is preferred to a large-scale downstream tax.

Note that $\mathcal{T}$ depends on model primitives such as the correlation between emissions intensity and market power: $\mathcal{T}$ is higher when this correlation is positive because that is when pass-through is lower to emissions-intense locations, worsening the downstream tax's effectiveness relative to upstream regulation. To illustrate this point, I use the model to compute $\mathcal{T}$ for varying degrees of spatial correlation, which involves the following steps:

1. Let the baseline spatial distribution of carbon density be the one observed in the data, denoted by a mean of μ and a standard deviation of σ.
2. Simulate the uniform downstream tax from section VI.A.1. Denote the emissions reductions it attains as $\Delta E_{\text{D}}(\sigma)$, given that it depends on the carbon-density dispersion σ. Simulate the upstream policy

from section VI.A.2 and denote its emissions reductions as $\Delta E_{\mathrm{U}}(\sigma)$. Finally, compute the targeting premium, $\mathcal{T}(\sigma) \equiv \Delta E_{\mathrm{U}}(\sigma) - \Delta E_{\mathrm{D}}(\sigma)$.

3. Construct a mean-preserving spread (MPS) of the baseline distribution of carbon density: it also has mean μ but an alternative standard deviation denoted $\tilde{\sigma}$. Repeat step 2, computing the targeting premium under the MPS, $\mathcal{T}(\tilde{\sigma}) \equiv \Delta E_{\mathrm{U}}(\tilde{\sigma}) - \Delta E_{\mathrm{D}}(\tilde{\sigma})$. Repeat for various MPSs (i.e., for various values of $\tilde{\sigma}$).

Figure 6*A* shows that $\mathcal{T}(\tilde{\sigma})$ is increasing in $\tilde{\sigma}$: the more heterogeneity in emissions intensity, the higher the value of targeting and therefore the higher the enforcement cost society is willing to pay for targeting. Moreover, this value is higher when there is market power—upstream targeting is valuable not just because of the standard efficiency argument that holds under any kind of market structure but also because it avoids the pass-through distortions that market power introduces for the downstream policy. Importantly, if the correlation between emissions intensity and market power had been of the opposite sign, then market power would have improved targeting by raising pass-through to emissions-intense locations (i.e., the solid line in fig. 6 would be below the dashed line). Thus, market power may worsen or improve an environmental policy's targeting depending on the direction in which it distorts pass-though, which in turns depends on how industry primitives shape the aforementioned correlation.

These qualitative results add nuance to the classic intuition that by depressing quantities, "the monopolist is the conservationist's friend" (Solow 1974). While we can confirm this proenvironment intuition of

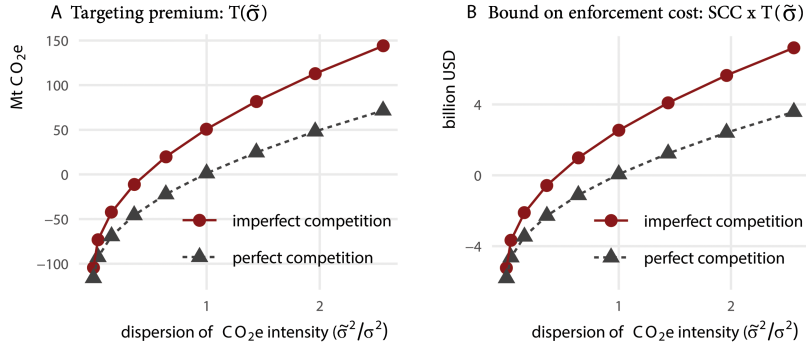

FIG. 6.—Interaction between emissions heterogeneity, market structure, and targeting. The vertical axis shows the targeting premium $\mathcal{T}(\tilde{\sigma})$ for a given degree of emissions dispersion $\tilde{\sigma}$. *A* reports the targeting premium in units of CO_2e, while *B* reports its dollar value under a SCC of US$50 per tonne. The horizontal axis in shows the deviation of the counterfactual dispersion $\tilde{\sigma}$ from the observed degree of dispersion σ, reported as the ratio $\tilde{\sigma}/\sigma$.

market power by comparing the laissez-faire equilibrium with and without market power—indeed, forest cover is 25% higher when market power is present—the role of market power for the transmission of environmental policy is more ambiguous. In this setting, the empirical object that resolves this ambiguity is the sign of the spatial correlation between market power and the environmental externality.

C. *Performance of Feasible Second-Best Policies Relative to First Best*

The previous section analyzed how emissions heterogeneity and market power interact to shape the trade-off between targeting and scale. To isolate core mechanisms, we focused on emissions reductions (net of enforcement costs) as the main policy goal. While this approach resonates with the notion of regulators being entrusted with narrow mandates rather than total welfare maximization, this section complements the preceding analysis by considering a broader welfare metric that accounts for producer and consumers surplus, as well as the government's fiscal balance. Moreover, I benchmark results against the first-best allocation, which I introduce next.

1. First-Best Benchmark

The equilibrium that results from a perfectly competitive setting with perfectly targeted Pigouvian taxes is the first-best efficient benchmark—there is no market power distortion, and environmental damages are fully internalized. To characterize this allocation, I simulate the equilibrium under perfectly competitive conduct and with the following Pigouvian tax per unit of output and subsidy per unit of forested land:

$$t_i^{c*} = \text{SCC} \times e_i^{c,\text{NLUC}} \quad \text{and} \quad s_i^{\mathcal{N}*} = \text{SCC} \times e_i^{\text{LUC}}, \tag{14}$$

where $e_i^{c,\text{NLUC}}$ is CO_2e on-farm emissions per tonne of output (e.g., emissions from cattle methane) and e_i^{LUC} is CO_2e emissions per hectare of forested land. The output tax targets on-farm emissions (which vary by commodity), while the forest subsidy targets land use change emissions (independent of commodity). We label the allocation resulting from this equilibrium with a superscript asterisk. The Pigouvian output taxes derived in equation (14) will not decentralize the first best if we impose imperfectly competitive conduct instead. Hence, they need to be adjusted to account for market power. Let t_i^c be the output tax that decentralizes the first best in the imperfectly competitive case. Appendix A, section A.5, shows that the expression for such a tax is $t_i^c = t_i^{c*} - p_i^{c*}[1 - \mu_i^c(L_i^{c*})]$, where $\mu_i^c(L_i^{c*}) < 1$ is the markdown evaluated at the first-best land allocation L_i^{c*}. The takeaway

is $t_i^c < t_i^{c*}$: the output tax is lower than under perfect competition because market power is already pushing quantities part of the way toward the first best. This is the classic insight from Buchanan (1969) on optimally taxing a polluting monopolist.

2. Welfare Analysis of Second-Best Policies

I follow standard approaches from single-industry empirical industrial organization studies by constructing welfare as the sum of its constituent parts (Fowlie, Reguant, and Ryan 2016; Hsiao 2021). I define social welfare W as the sum of consumer surplus (CS), producer surplus (PS), government surplus (G) and emissions (E) evaluated at the SCC (i.e., $W = \text{CS} + \text{PS} + G - \text{SCC} \times E$). Details on welfare accounting for each component are given in appendix D, section D.4. Table 5 reports the welfare impact of each policy, as percentage point changes from the laissez-faire as well as share of the first-best efficiency gains.

First, within the category of downstream domestic taxes, emissions declines are larger the more finely targeted the policy is. Producer and consumer surplus declines are also smaller the more targeted the tax is. Both forces contribute to the larger welfare gains attained by finer targeting. In the case of unilateral tariffs by foreign trade partners, impacts on all

TABLE 5
DECOMPOSITION OF WELFARE IMPACTS

	Change Relative to Laissez-Faire Equilibrium (percentage points)					First-Best Efficiency Gains (%)	
	ΔE (1)	ΔPS (2)	ΔCS (3)	ΔCS_{SA} (4)	ΔCS_{ROW} (5)	ΔE (6)	ΔW (7)
	A. Downstream Domestic Taxes						
Untargeted (sec. VI.A.1)	−26.88	−21.49	−28.07	−28.55	−27.57	29.27	26.40
Region targeted (sec. VI.A.3)	−30.48	−20.74	−26.36	−27.12	−25.60	33.20	30.85
County targeted (sec. VI.A.3)	−34.47	−18.96	−23.43	−24.26	−22.59	37.54	35.99
	B. Downstream Foreign Taxes						
Unilateral EU tariff (app. D, sec. D.5.2)	−2.92	−2.44	−5.67	2.36	−13.80	3.18	2.44
	C. Upstream Policies						
Forest subsidy (sec. VI.A.2)	−28.09	354.60	−.10	−.10	−.10	30.60	32.55
First best (sec. VI.C.1)	−91.81	853.99	−37.62	−38.09	−37.14	100.00	100.00

NOTE.—Details on how each welfare component is computed are in app. D, sec. D.4. CS_{SA} refers to consumer surplus of South America, and CS_{ROW} refers to that of foreign consumers. Columns 1–5 display percentage point changes between the policy counterfactual and the laissez-faire. Columns 6 and 7 display the changes in emissions and overall welfare, as a percent of the changes delivered by the first best. Percentage point changes for government surplus are not displayed because the baseline level for G in the laissez-faire is zero (however, government surplus does indeed enter ΔW).

welfare components are small because of consumption-side leakage. Moreover, surplus for South American consumers increases, while that of foreign consumers declines. Second, moving on to upstream policies, the limited-scale but targeted forest subsidy delivers more emissions abatement than the untargeted downstream tax. However, once we allow for some targeting of the downstream tax, even as coarsely as by region, the upstream policy is outperformed in abatement terms. The forest subsidy generates smaller declines in consumer surplus than the downstream tax though, because it targets production for farmers only in the most emissions-intense areas. Producers also gain significantly because they are being paid to keep their land forested. Despite the transfer to producers, aggregate welfare gains are modest because the forest subsidy implies a large fiscal cost. Finally, the first-best welfare gains come from large reductions in emissions damages as well as increases in producer surplus because they are subsidized to keep land forested. Both forces outweigh the decline in consumer surplus and the net fiscal cost of the forest subsidies. Overall, the range of second-best policies I consider attain 29%–38% of first-best emissions reductions and 26%–36% of welfare gains.

VII. Final Remarks

Environmental externalities often coexist with other distortions in real-world markets. In this paper, I study how the transmission of environmental policy occurs along a supply chain, in particular when preexisting distortions lie between the stage where emissions are generated and the stage where regulation is feasible. I do so in the empirical context of an agricultural supply chain, where the link between upstream farmers and downstream consumers is intermediated by a concentrated layer of agribusiness intermediaries with plausible monopsony power. While policy interventions are more feasibly implemented at the supply chain's bottleneck, I show that the impact on the upstream producers making the environmentally relevant decisions can be poorly targeted and especially so when market power is present. Concretely, if the upstream producers most subject to market power are also the most emissions intense, then the Pigouvian signal of a downstream tax is eroded most where social cost is highest.

More generally, market power can be theoretically ambiguous for the transmission of environmental policy, and this paper proposes a specific mechanism—differential pass-through—that microfounds this ambiguity. Moreover, the empirical object that resolves this theoretical ambiguity is the correlation between market power and the environmental externality. In this setting, market power worsens targeting because the correlation is positive. In other settings the correlation could be negative, in which case market power would improve targeting by raising Pigouvian pass-through

where social cost is highest. Understanding how specific industries vary in the primitives shaping such correlations can help reveal general insights about the efficiency of corrective policies.

Data and Code Availability

Code replicating the tables and figures in this article can be found in Domínguez-Iino (2025) in the Harvard Dataverse: https://doi.org/10.7910/DVN/RY2CVR.

References

Assunção, J., R. McMillan, J. Murphy, and E. Souza-Rodrigues. 2023. "Optimal Environmental Targeting in the Amazon Rainforest." *Rev. Econ. Studies* 90 (4): 1608–41.

Barreto, P., R. P. A. Brandão Jr., and S. Baima. 2017. "Will Meat-Packing Plants Help Halt Deforestation in the Amazon?" Technical report, Inst. Centro Da Vida.

Bento, A. M., L. H. Goulder, M. R. Jacobsen, and R. H. Von Haefen. 2009. "Distributional and Efficiency Impacts of Increased US Gasoline Taxes." *A.E.R.* 99 (3): 667–99.

Bergquist, L. F., and M. Dinerstein. 2020. "Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya." *A.E.R.* 110 (12): 3705–47.

Bresnahan, T. F. 1989. "Empirical Studies of Industries with Market Power." *Handbook Indus. Org.* 2:1011–57.

Buchanan, J. M. 1969. "External Diseconomies, Corrective Taxes, and Market Structure." *A.E.R.* 59 (1): 174–77.

Bulow, J. I., and P. Pfleiderer. 1983. "A Note on the Effect of Cost Changes on Prices." *J.P.E.* 91 (1): 182–85.

Chatterjee, S. 2023. "Market Power and Spatial Competition in Rural India." *Q.J.E.* 138 (3): 1649–711.

Conte, B. 2020. "Climate Change and Migration: The Case of Africa." BSE Working Paper no. 1411, Barcelona School Econ.

Conte, B., K. Desmet, D. K. Nagy, and E. Rossi-Hansberg. 2021. "Local Sectoral Specialization in a Warming World." *J. Econ. Geography* 21 (4): 493–530.

Copeland, B. R., J. D. Shapiro, and M. S. Taylor. 2022. "Globalization and the Environment." In *Handbook of International Economics*, vol. 5, edited by Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, 61–146. Amsterdam: Elsevier.

Costinot, A., D. Donaldson, and C. Smith. 2016. "Evolving Comparative Advantage and the Impact of Climate Change: Evidence from 1.7 Million Fields around the World." *J.P.E.* 124 (1): 205–48.

Cruz, J. L., and E. Rossi-Hansberg. 2024. "The Economic Geography of Global Warming." *Rev. Econ. Studies* 91 (2): 899–939.

Dhingra, S., and S. Tenreyro. 2020. "The Rise of Agribusiness and the Distributional Consequences of Policies on Intermediated Trade." CEPR Discussion Paper no. DP14384, Centre Econ. Policy Res., London.

Domínguez-Iino, T. 2025. Replication Data for: "Efficiency and Redistribution in Environmental Policy: An Equilibrium Analysis of Agricultural Supply Chains". Harvard Dataverse, https://doi.org/10.7910/DVN/RY2CVR.

Eaton, J., and S. Kortum. 2002. "Technology, Geography, and Trade." *Econometrica* 70 (5): 1741–79.

Ehrl, P. 2017. "Minimum Comparable Areas for the Period 1872–2010: An Aggregation of Brazilian Municipalities." *Estudos Econ.* 47 (1): 215–29.
EPA (Environmental Protection Agency). 2018. "Greenhouse Gas Inventory Data Explorer." https://cfpub.epa.gov/ghgdata/inventoryexplorer/.
Fabra, N., and M. Reguant. 2014. "Pass-Through of Emissions Costs in Electricity Markets." *A.E.R.* 104 (9): 2872–99.
Farrokhi, F., and A. Lashkaripour. 2021. "Can Trade Policy Mitigate Climate Change?" *Econometrica* 93 (5): 1561–99.
Farrokhi, F., and H. S. Pellegrina. 2023. "Trade, Technology, and Agricultural Productivity." *J.P.E.* 131 (9): 2509–55.
Fowlie, M., M. Reguant, and S. P. Ryan. 2016. "Market-Based Emissions Regulation and Industry Dynamics." *J.P.E.* 124 (1): 249–302.
Goldberg, P. K. 1995. "Product Differentiation and Oligopoly in International Markets: The Case of the US Automobile Industry." *Econometrica* 63 (4): 891–951.
Hsiao, A. 2021. "Coordination and Commitment in International Climate Action: Evidence from Palm Oil." Working paper, Univ. Chicago.
IIASA/FAO (International Institute for Applied Systems Analysis/UN Food and Agriculture Organization). 2012. "Global Agro-Ecological Zones (GAEZ v. 3.0)." https://iiasa.ac.at/models-tools-data/gaez.
Kortum, S., and D. Weisbach. 2017. "The Design of Border Adjustments for Carbon Prices." *Nat. Tax J.* 70 (2): 421.
McFadden, D. 1977. "Modelling the Choice of Residential Location." Discussion Paper no. 477, Cowles Foundation, New Haven, CT.
Mitra, S., D. Mookherjee, M. Torero, and S. Visaria. 2018. "Asymmetric Information and Middleman Margins: An Experiment with Indian Potato Farmers." *Rev. Econ. Statis.* 100 (1): 1–13.
Nath, I. 2025. "Climate Change, the Food Problem, and the Challenge of Adaptation through Sectoral Reallocation." *J.P.E.* 133 (6): 1705–56.
Pellegrina, H. S. 2022. "Trade, Productivity, and the Spatial Organization of Agriculture: Evidence from Brazil." *J. Development Econ.* 156:102816.
Pereira, R., L. L. Rausch, A. Carrara, and H. K. Gibbs. 2020. "Extensive Production Practices and Incomplete Implementation Hinder Brazil's Zero-Deforestation Cattle Agreements in Para." *Tropical Conservation Sci.* 13:1940082920942014.
Poore, J., and T. Nemecek. 2018. "Reducing Food's Environmental Impacts through Producers and Consumers." *Science* 360 (6392): 987–92.
Roberts, M. J., and W. Schlenker. 2013. "Identifying Supply and Demand Elasticities of Agricultural Commodities: Implications for the US Ethanol Mandate." *A.E.R.* 103 (6): 2265–95.
Rubens, M. 2023. "Market Structure, Oligopsony Power, and Productivity." *A.E.R.* 113 (9): 2382–410.
Ryan, S. P. 2012. "The Costs of Environmental Regulation in a Concentrated Industry." *Econometrica* 80 (3): 1019–61.
Scott, P. T. 2014. "Dynamic Discrete Choice Estimation of Agricultural Land Use." TSE Working Paper no. 14-526, Toulouse School Econ.
Skidmore, M. E., F. Moffette, L. Rausch, M. Christie, J. Munger, and H. K. Gibbs. 2021. "Cattle Ranchers and Deforestation in the Brazilian Amazon: Production, Location, and Policies." *Global Environmental Change* 68:102280.
Solow, R. M. 1974. "The Economics of Resources or the Resources of Economics." *A.E.R.* 64 (2): 1–14.
Sotelo, S. 2020. "Domestic Trade Frictions and Agriculture." *J.P.E.* 128 (7): 2690–738.
Souza-Rodrigues, E. 2019. "Deforestation in the Amazon: A Unified Framework for Estimation and Policy Analysis." *Rev. Econ. Studies* 86 (6): 2713–44.

Spawn, S., and H. Gibbs. 2020. "Global Aboveground and Belowground Biomass Carbon Density Maps for the Year 2010." ORNL DAAC, Oak Ridge, TN. https://doi.org/10.3334/ORNLDAAC/1763.

Train, K. E. 2009. *Discrete Choice Methods with Simulation*. Cambridge: Cambridge Univ. Press.

Weyl, E. G., and M. Fabinger. 2013. "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition." *J.P.E.* 121 (3): 528–83.

Zavala, L. 2021. "Unfair Trade? Market Power in Agricultural Value Chains." Working paper, Princeton Univ., Princeton, NJ.