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

Tomás Domínguez-Iino†

May 1, 2023

Click here for latest version.

Abstract

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 carbon tax on farmers, I show how alternative market-based policies—downstream agribusiness taxes—reduce upstream emissions but their effectiveness is limited by international leakage and domestic mistargeting, 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, command-and-control tools perform robustly when markets face pre-existing distortions.

*First version: November 2020. This paper is the first chapter of my Ph.D. dissertation at New York University. I thank Alessandro Lizzeri, Guillaume Fréchette, Elena Manresa, and Paul Scott for their unconditional support at the early stages of this project. For helpful feedback, I am grateful to Milena Almagro, Tim Christensen, Christopher Flinn, Martin Rotemberg, Sharon Traiberman, Daniel Waldinger, and seminar audiences at NYU, UPenn, Chicago, Northwestern, Princeton, Yale SOM, LSE, Sciences Po, Warwick, University of Washington, Nottingham, PUC Rio, ITAM, Di Tella, UCA, Ghent, Universitat de Barcelona, Iowa State, Johns Hopkins, SED Annual Meetings, Mid-Atlantic Trade Workshop, Chicago Fed, Fed Board, Bank of Spain, Central Bank of Colombia, World Bank DRG, and CAF Development Bank of Latin America. I have also greatly benefited from conference discussions by Bouwe Dijkstra, Joel Ferguson, Allan Hsiao, Ander Iraizoz-Olaetxea, Pierre Magontier, and Nicolas Morales. Any errors or omissions are my own.

†Federal Reserve Board. E-mail: tdomingueziino@gmail.com. The views expressed are those of the author and not necessarily those of the Federal Reserve Board or the Federal Reserve System.

1 Introduction

Many of the major industries contributing to climate-change produce goods that are tradable and are subject to pre-existing distortions beyond the environmental externality, market power being a case in point. How can we regulate such industries 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 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 main goal of this paper is to evaluate how environmental policy is transmitted along a supply chain, in particular when the stage where it can be feasibly implemented differs from where emissions are generated. To do so, 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 the geographic location 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. After estimating my model, I use its implied counterfactuals to compare the performance of feasible environmental regulations—both market-based and command-and-control—in terms of their efficiency and distributional impact along the supply chain.

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 non-regulated markets and the externality remains uncorrected. Third, agricultural supply chains in developing countries are often fragmented and funnel-shaped, with atomistic upstream farmers selling their output to a concentrated sector of downstream intermediaries. Monopsony power of intermediaries over farmers, especially in remote locations, introduces an additional market distortion on top of the environmental externality. Hence, market-based policies such as carbon taxes may perform poorly if the markets they have to operate through are subject to such pre-existing distortions.

To incorporate the stylized features of developing world agriculture into my analysis, my empirical model includes the following ingredients: (i) an upstream production stage with rich spatial heterogeneity and the key margins driving agricultural emissions, (ii) a concentrated sector of agribusiness intermediaries that hold market power over the farmers they source from, and (iii) international trade to capture leakage effects across consumer markets. Agriculture provides an ideal setting to study how environmental policy is transmitted along a supply chain because it is characterized by having most of its emissions generated at the atomistic stage, where regulation is most challenging to enforce. This naturally raises the question of how poorly targeted an environmental policy is if it cannot be imposed directly on the upstream farmers whose incentives it aims to correct, but rather on the downstream agribusiness firms whose compliance is easier to enforce. I now proceed to describe the main findings from my counterfactual policy analysis.

Main results. The first part of my counterfactual analysis shows how a feasible market-based 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 allocation through the textbook Pigouvian mechanism, in reality this type of policy is largely absent in developing world agriculture due to logistical enforcement challenges as well as its political infeasibility. Motivated by these constraints, I evaluate a policy that is considered feasible among policymakers but which is second-best: a downstream tax levied on agribusiness firms by consumer markets. An example of this would be an environmental tariff imposed by trading partners on their imports from South America, which would be levied at the port on downstream agribusiness firms instead of directly on upstream farmers. In terms of the policy's effectiveness, I show that if only a subset of consumer markets implement the tax then there is substantial emissions leakage. If the tax is implemented unilaterally by the European Union, 80% of the emissions reductions achieved by the EU are undone by the re-routing of shipments to non-regulated consumer markets.

Apart from its ineffectiveness due to leakage across consumer markets, the downstream tax is also poorly targeted across upstream producers. This occurs because farmers in the most remote and emissions-intense regions have the least elastic supply. Therefore, the tax is spatially mistargeted because it leads production to drop least in the upstream locations where the environmental cost is highest. Because of inelastic supply, these remote regions are also where farm-gate prices drop most. Since these locations are among the poorest ones, 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 to farmers to shift their production incentives, agribusiness monopsony power plays a role. Specifically, market power erodes the Pigouvian signal contained in the tax because of incomplete pass-through to upstream farmers, resulting in muted emissions abatement.

The second part of my counterfactual analysis shows how a regulator's optimal choice between a market-based and a command-and-control policy tool depends on two features of agricultural supply chains: the degree of heterogeneity in carbon footprints across upstream farmers and the degree of intermediary market power. Market-based policies that are implemented downstream, such as the taxes from the first part of my analysis, are mistargeted because they do not take into account the spatial heterogeneity in carbon footprints across upstream farmers. Naturally, this mistargeting becomes worse as the degree of upstream heterogeneity increases. Furthermore, the mistargeting is amplified by market power because the less competitive upstream locations tend to be more emissions-intense. Therefore, the upstream pass-through of the policy is especially incomplete in emissions-intense locations.

Overall, the results indicate that market-based policy can perform poorly when the market it has to operate through is distorted in ways beyond the environmental externality. In such cases, I show command-and-control policies such as conservation zones can be better targeted and robust to market structure, precisely because they do not operate through the market mechanism. However, their main drawback is their high enforcement cost. Thus, the regulator faces a trade-off between targeting and enforcement costs when choosing between market-based and command-and-control tools, with the starkness of the trade-off depending on two model primitives: the degree of upstream spatial heterogeneity and market structure. I now proceed to describe the theoretical and empirical aspects of my analysis in further detail.

Theoretical and empirical methods in detail. First, I present a model of land use incorporating two key margins driving the environmental impact of agriculture: an extensive margin of converting natural forested land into *new* agricultural land, and the choice of which specific agricultural commodity is produced on *existing* agricultural land. Disentangling the two is critical to evaluate policies operating through one margin but not the other, for example, commodity-specific taxes versus commodity-blind deforestation fines. However, existing work typically focuses on a single margin at a time. On the one hand, a recent trade literature uses the Ricardian framework of Eaton and Kortum (2002) to study the determinants of the spatial distribution of agricultural activity, but abstracts from the extensive margin since environmental consequences, such as those arising from deforestation, are not their object of study (Costinot, Donaldson and Smith, 2016; Pellegrina, 2019; Sotelo, 2020). On the other hand, a recent applied microeconomics literature on land-use change addresses the extensive margin, but abstracts from which specific agricultural commodities are produced on the cleared land (Scott, 2013; Souza-Rodrigues, 2019).

I simultaneously incorporate both margins by modeling the land use decisions of farmers as a nested discrete choice problem, with a natural land use nest and an agricultural land use nest. Within the agricultural nest, the model collapses to the Ricardian framework and the substitution patterns between commodities map to trade elasticities, while substitution patterns across nests map to land-use change elasticities. Thus, my land use model incorporates two key margins driving agricultural emissions within a single framework—how much land is cleared and what gets produced on it—while delivering estimates that are consistent with existing work estimating each margin separately. To estimate the model's supply parameters, I address endogeneity by constructing demand shifters from quasi-experimental variation in world commodity prices.

Second, I use granular data on domestic trade flows to document concentration among agribusiness firms, a feature which 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 the supply elasticity of farmers and the degree of agribusiness concentration. Therefore, spatial heterogeneity in supply elasticities and local concentration determine how surplus is shared between farmers and agribusiness firms at each upstream market. Because carbon tax proposals often suggest levying the tax at the concentrated stage of the supply chain due to ease of logistical implementation, market structure matters for quantifying pass-through to the upstream farmers who ultimately make the environmentally-relevant decisions.

Concretely, the intermediary part of my model consists of a simple oligopsony specification that nests the perfectly competitive case in order to flexibly alternate between conduct assumptions, allowing me to evaluate how market structure alters the upstream transmission of environmental policy. The final consumers at the end of the supply chain are located in domestic as well as foreign markets, which opens the door for consumption leakage from incomplete regulation. Although the focus of this paper is on the domestic transmission of downstream taxes to upstream producers, the framework can also be used to analyze leakage across international consumer markets. Given much of the policy debate on agricultural carbon taxes is indeed at the international level, with downstream taxes being implemented as tariffs, I also include results on international leakage for the sake of completeness.

Third, after estimating the key elasticities in my model, I exploit its spatial structure to highlight the unique challenges agriculture presents when it comes to emissions regulation, and the implications for choosing between market-based and command-and-control policy tools. Market-based tools such as carbon taxes are appealing because they achieve broad geographic coverage while avoiding the enforcement costs of direct regulation. However, they can be poorly targeted if not levied directly on the upstream farmers whose incentives they aim to correct, but rather on the downstream agribusiness firms whose compliance is easier to enforce. My analysis shows this mistargeted aspect of the market-based policy becomes worse as heterogeneity in carbon footprints across upstream producers increases. Moreover, market power worsens the mistargeting by lowering pass-through to the least competitive upstream markets, which happen to be the ones with the highest emissions intensity. I conclude by showing how command-and-control policies that target a subset of high-carbon locations can dominate market-based policies when upstream heterogeneity is wide enough, and more so when market power is present.

Finally, note that while the market-based tool is a price regulation, the command-and-control tool is a quantity regulation. Therefore, the findings are reminiscent of the classic trade-off between regulating prices versus quantities in settings where regulators face producers with heterogeneous emissions intensities (Weitzman, 1974). The intuition behind the results follows classic insights from public finance, as the degree of heterogeneity dictates which type of regulatory tool is optimal, while adding the insight that pre-existing distortions such as market power can tilt the trade-off towards the quantity regulation.

Main contribution and relation to existing literature. This paper's main contribution is to show how market-based environmental policy is transmitted along a supply chain, especially when the stage where it can be feasibly implemented differs from where the emissions are generated. Hence, this paper quantifies the extent to which the Pigouvian signal of a policy is distorted along the supply chain before reaching the market actors whose incentives it aims to correct. Agriculture provides an ideal setting to study this mechanism because it is uniquely characterized by having its emissions generated at the atomistic stage of its supply chain, rather than at the concentrated stage where policy implementation and enforcement is easier. Apart from agriculture's natural fit with the paper's research question, the South American agricultural sector in particular is of major importance in and of itself due to its status as a global agricultural powerhouse with a major emissions footprint. Over the next paragraphs, I break down the paper's contributions in further detail in terms of how they depart from the existing literature.

First, I link a trade literature that studies how comparative advantage shapes the spatial distribution of agriculture (Costinot et al., 2016; Pellegrina, 2019; Sotelo, 2020) to a recent land-use change literature in agricultural economics and empirical IO (Scott, 2013; Souza-Rodrigues, 2019). The trade literature studies how different commodities are allocated across existing agricultural land, but abstracts from the extensive margin of land conversion. By contrast, the land-use change 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. The trade literature's implied land-use change elasticities are significantly higher than those from the land-use studies, in part because they are estimated from substitution patterns across commodities on already cleared land, where switching costs might be lower than along the extensive margin. I incorporate both margins by modeling farmers' decisions as a nested discrete choice problem, and I show how this can reconcile the relatively high substitution elasticities between commodities estimated by the trade literature with the relatively low land-use change elasticities from the land-use studies.

Second, this paper relates to a growing literature at the intersection of trade and climate change. I use similar modeling tools as a subset of studies quantifying adaptation mechanisms to climate change, such as trade, migration, and sectoral reallocation (Conte, 2020; Conte, Desmet, Nagy and Rossi-Hansberg, 2021; Nath, 2020; Alvarez and Rossi-Hansberg, 2021). Rather than taking the more macroeconomic approach of this literature, which tends to include multiple sectors and factors, I focus on a single sector with a major impact on global emissions. This narrower approach allows me to estimate the model's relevant elasticities with plausibly valid instruments, but comes at the cost of abstracting away from reallocation across factors or sectors in my counterfactuals. Given land is the single most important factor driving emissions in agriculture, and even more so in South America, the abstraction away from other factors such as labor or capital is less consequential when studying the impact of agriculture on climate change relative to the opposite direction—adaptation of agriculture to climate change. For this reason, my analysis is not on how the damages from climate change can be mitigated by adaptation mechanisms such as trade, but instead considers the reverse direction: how trade policies in particular, and market-based tools more generally, can be used to reduce agriculture's contribution to climate change through its primary channel, land use change.

Given the important role of trade for environmental outcomes in this setting, this paper connects naturally to the extensive literature on the use of trade policy as environmental regulation, recently reviewed in Copeland, Shapiro and Taylor (2022). There are two ways in which this paper departs from most of this work, and in particular from recent studies within this literature (Kortum and Weisbach, 2017; Farrokhi and Lashkaripour, 2021; Hsiao, 2021). The first is the emphasis on how spatial heterogeneity in the effectiveness and incidence of environmental policy varies at the sub-national level, which is important because agricultural productivity and carbon density vary widely within a country. Moreover, understanding incidence at the sub-national level is important because it reflects the internal political constraints a government faces when designing environmental policy. Furthermore, the sub-national analysis is necessary to address the key questions in this paper, such as how mistargeted a flat downstream tax might be in terms of its upstream pass-through to domestic producers. The second departure is the imperfectly competitive setting, and in particular having market power on the demand side as opposed to the supply side. I show how market power can affect the performance of market-based environmental policy in a qualitative sense, and not just quantitatively. The sign of the correlation between market power and emissions intensity determines whether pass-through is higher or lower in the most emissions-intense locations, and hence whether market power improves or worsens targeting.

Third, this paper contributes to a literature at the intersection of industrial organization and environmental economics that goes back at least to Buchanan (1969), with modern approaches typically leveraging empirical IO methods (Fowlie, Reguant and Ryan, 2016). In such studies, firms typically exercise market power *downstream on consumers*, and the question is how to design environmental policy along efficiency criteria. I contribute to this literature by considering the case where firms exercise market power *upstream on their suppliers*: the farmers. This distinction matters because over 80% of agricultural emissions are generated at the upstream stage before the commodities leave the farm gate. This contrasts with fossil fuels, where upstream emissions generated by extraction are small relative to downstream emissions generated by consumption (i.e., the burning of fossil fuels for energy use). While agricultural emissions are mostly generated upstream across atomistic farmers, fossil fuel emissions are more likely to be concentrated downstream across a few large utility companies. This makes direct regulation at the emissions source logistically challenging in agriculture, which is why it is crucial to understand how policies which can be easily implemented at the concentrated stage of the supply chain are transmitted to the upstream farmers who make the environmentally-relevant decisions.

My emphasis on upstream transmission also opens a distributional channel on the supply-side, whereas most work on environmental policy incidence typically focuses on downstream impacts across consumers (Bento, Goulder, Jacobsen and Von Haefen, 2009; Fabra and Reguant, 2014). Understanding supply-side distributional effects is first order in this setting because agricultural policy is often designed with redistribution toward farmers as an explicit goal, thus posing a major barrier to advancing environmental regulation. Finally, the imperfectly competitive setting relates to a recent literature on intermediaries and market power in the developing world, most of which takes place in agricultural markets (Bergquist and Dinerstein, 2020; Chatterjee, 2019; Rubens, 2019; Dhingra and Tenreyro, 2020; Zavala, 2021). While most of these studies focus on the welfare impacts of market power per se, this paper only incorporates market power to show how it interferes with the performance of market-based environmental policy. To the best of my knowledge, I am unaware of work in this area studying this specific interaction.

2 Data

I construct a county-level panel of agricultural supply and demand from 1995-2017 by combining various data sources from Argentina and Brazil. Throughout the paper, I use the term "South America" when referring to the two countries jointly because they account for the bulk of the continent's environmentally relevant agricultural production, holding over 90% of its soybean and maize output and over 75% of its cattle herd. The supply side data consists of a county-level panel of land use, agricultural output, agronomic productivity, and farm-gate prices for the following commodities: beef cattle, soybeans, maize, wheat, rice, sunflower, and sugarcane. These commodities account for over 85% of all agricultural land in Argentina and Brazil. For the demand side, I connect each county's production to its nation-level destination markets using trade flow data. A summary of each data source is listed below.

Geographic unit of analysis and temporal frequency. The smallest administrative unit at which the data is available for Argentina is a department ("partido"), while for Brazil it is a municipality ("município"). Given Brazilian municipality borders have changed over time, I use the procedure from Ehrl (2017) to construct time-consistent spatial units, known as "Áreas Mínimas Comparáveis" (AMC). Throughout the paper, I use the term "counties" when referring to Argentine departments and Brazilian AMCs. Given most of the data in this paper is from decadal agricultural censuses, the time period used in the main estimation exercise is a decade. Therefore, changes over time are to be interpreted as fairly long-run changes.

Agronomic productivity. Data on agricultural productivity for major crops is available from the Food and Agriculture Organization's Global Agro-Ecological Zones project (FAO-GAEZ) at 5 arc-minute resolution for over one million grid cells around the globe (IIASA/FAO, 2012). Productivity is measured as potential yields predicted by an agronomic model based on agro-climatic fundamentals. The model's parameters are estimated from field and lab experiments in the agronomic literature, and its specific inputs are: soil characteristics, land gradient, elevation, temperature, rainfall, and sun exposure. To obtain beef "yields per hectare" I construct a measure of cattle productivity by combining the FAO-GAEZ pasture yield index with county-level data on cattle stocking rates. The procedure is described in Appendix section C.1.

Land-use. County-level data on forested, agricultural, and pasture area are from the agricultural censuses of Argentina and Brazil. For Argentina, the source is the National Statistical Institute (INDEC) and for Brazil it is the the Brazilian Institute of Geography and Statistics (IBGE). The Brazilian census also reports acreage allocated to individual crops. For Argentina, I obtain individual crop acreage from the Ministry of Agriculture's "Datos Agroindustriales" database (DA-MAGYP).

Agricultural output. County-level crop and livestock output for Argentina is from the agricultural census, DA-MAGYP, and the livestock registry at the National Food Safety Agency (SENASA). For Brazil I use the agricultural census, which I complement with higher-frequency municipal survey data from Produção Agrícola Municipal (PAM) and Pesquisa da Pecuária Municipal (PPM).

Trade flows. International trade flows of agricultural commodities at the nation-to-nation level are obtained from FAOSTAT. To determine sourcing within Argentina and Brazil I use domestic supply chain data from TRASE. This data is constructed from customs records and maps annual trade flows (in physical quantities and port of export FOB values) from source counties to national-level destination markets, as well as to the agribusiness firms intermediating the transactions.

Prices. Farm-gate prices of crops and cattle are obtained from production value and quantity data. For Brazil, the sources are PAM and the agricultural census. For Argentina, I use cattle transaction microdata from DA-MAGYP that directly reports transaction prices. Destination prices are obtained from the TRASE data on values and quantities. Since the values are reported as port of export FOB, the destination prices reflect the price the agribusiness firms receive for delivering the goods to the port of export, but not to the final destination market. Therefore, the destination prices include domestic transport costs from farm to port, but not the international transport costs from port to overseas destination (which are paid by the final destination consumers).

Emissions. I compute land-use change emissions using biomass data at 300m spatial resolution from the carbon density maps compiled by Spawn and Gibbs (2020), available at NASA Earthdata. To compute emissions footprints across commodities I use data from Poore and Nemecek (2018) and Clark, Domingo, Colgan, Thakrar, Tilman, Lynch, Azevedo and Hill (2020).

Weather shocks. Data on extreme temperatures at the nation-level are from FAOSTAT. I construct county-level weather shocks from the National Centers for Environmental Prediction CFSR.

3 Stylized facts

First, I introduce the environmental science facts indicating which economic decision margins are most relevant in determining agricultural emissions. Second, I introduce the main economic and institutional features of the agricultural sectors of Argentina and Brazil. Finally, I describe how these environmental and economic facts motivate the key ingredients of my model and how they constrain the counterfactuals I run to a feasible subset of policies.

Fact 1: Agricultural emissions are mostly generated at the upstream stage of the supply chain

I focus on three crucial decision margins that drive emissions in the agricultural sector, all of which take place at the upstream stage of the supply chain:

i. *How much land is cleared.* Over 80% of agricultural emissions are generated upstream before the commodities leave the farm-gate, mostly due to land-use change and on-farm sources such as enteric methane (Figure 1). The land-use change share is especially high in South America (above 70%) relative to the rest of the world (below 40%). The upstream nature of agricultural emissions contrasts with fossil fuels, where upstream emissions generated by extraction are small relative to downstream emissions generated by consumption (i.e., the burning of fossil fuels for energy use).

ii. *Which commodity is produced on the cleared land.* Emissions footprints vary widely across agricultural commodities, even after taking into account differing land requirements (Figure 1). These large differences are robust to whether emissions footprints are specified on a per kcal or per protein content basis. For example, beef contains 25 times more CO_2e/kg of protein than plant-based high-protein alternatives. Substantial variation exists even among plant-based commodities, with rice generating twice as much CO_2e/kcal than wheat.

iii. *Where the clearing and producing takes place.* Emissions footprints vary widely across space due to the uneven geographic distribution of carbon stocks that would be potentially released into the atmosphere from land clearing (Figure 2).

Figure 1: Emissions footprints and sources along the agricultural supply chain.

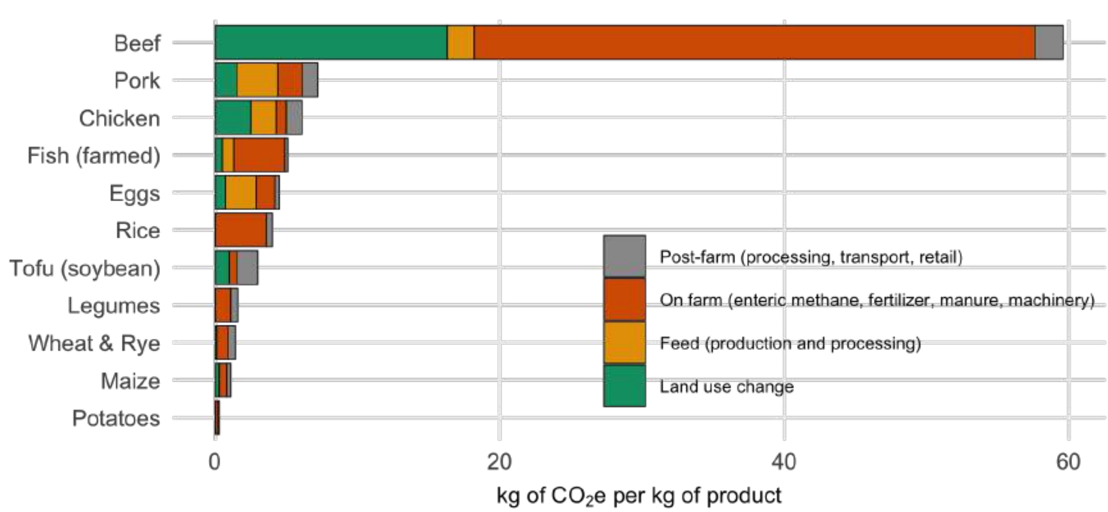

Notes: This figure shows emissions footprints at their global average, using data from Poore and Nemecek (2018). For South America, the land use change share of total emissions is significantly higher, around 70% (FAOSTAT). For robustness of results to emissions footprints specified on a per kcal or per protein content basis, see Figure 12 in Appendix B.

Figure 2: Spatial distribution of carbon density (tC/ha) in South America.

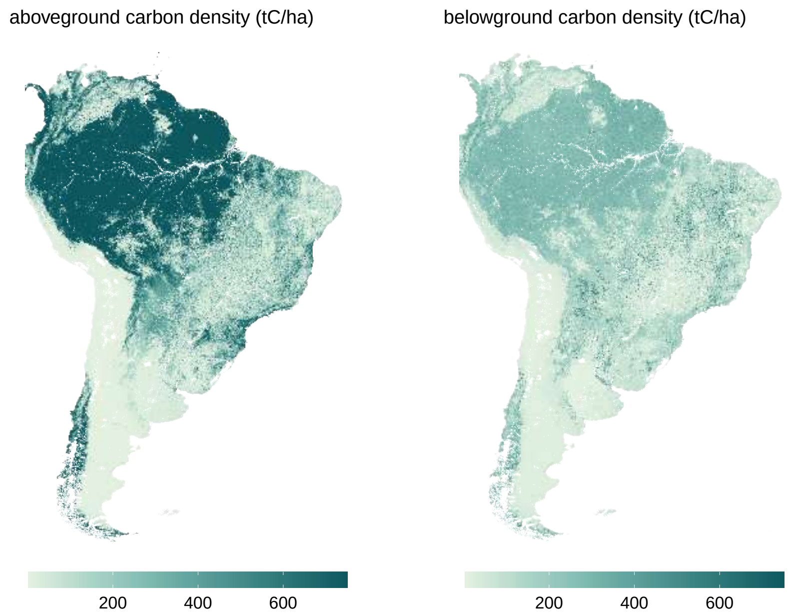

Notes: Figures are constructed using the carbon density dataset from Spawn and Gibbs (2020).

All told, agriculture accounts for 26% of anthropogenic emissions (Poore and Nemecek, 2018). South America's annual agricultural emissions have hovered around 3 Gt CO_2e since 1990, roughly 27% of world agricultural emissions (FAOSTAT). Such magnitudes exceed those of any major sector of the US economy (EPA, 2018): industry (1.5 Gt), electricity (1.8 Gt), or transport (1.9 Gt).

Fact 2: International demand has played a major role in the evolution of South American land use, both across time and space

One of the most salient developments in South American agriculture over recent decades has been the dramatic expansion of soybean production. Before 1980, soybean acreage was the lowest of any major crop, yet by 2005 it exceeded all other major crops combined (Figure 3). Growing international demand, especially from Asia, has been a major driver behind such trends: over 70% of soybean output is exported and over 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 shifted from the most soybean-suitable areas (central-south Brazil, mid-east Argentina) to cheaper land markets in frontier agricultural regions, which is where the forests lie (northern Brazil and parts of the west and north of Argentina). Hence, although soybean expansion may not directly lead to deforestation, it may do so indirectly by displacing land-intensive cattle grazing to the agricultural frontier (Figure 4). Accounting for interactions between agricultural commodities is therefore crucial for understanding deforestation in the South American context.[1]

[1]The extent to which these interactions matter for deforestation is context-specific. For example, Indonesian deforestation is driven almost entirely by palm oil, so abstracting from such interactions seems reasonable (Hsiao, 2021).

Figure 3: The South American soybean boom across time.

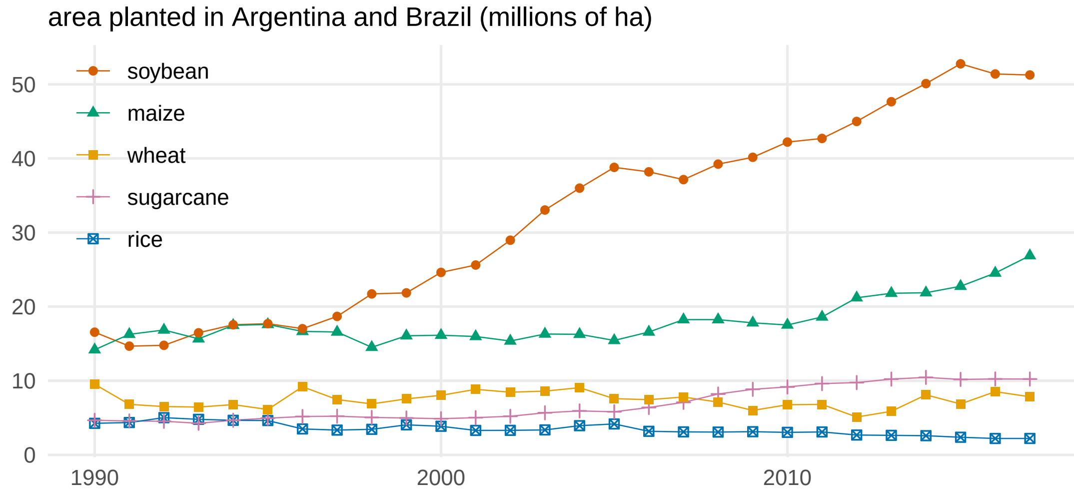

Notes: The figure shows acreage allocated to major crops in Argentina and Brazil (1990-2017). By 2017, 91% of planted land in Argentina was concentrated among 4 crops: soybeans (46%), maize (24%), wheat (16%), sunflower (5%). By 2017, 85% of planted land in Brazil was concentrated among 4 crops: soybeans (44%), maize (21%), sugarcane (11%), beans (4%), wheat (3%), rice (2%). Sources: DA-MAGYP (Argentina), CONAB (Brazil).

Figure 4: The South American soybean boom across space.

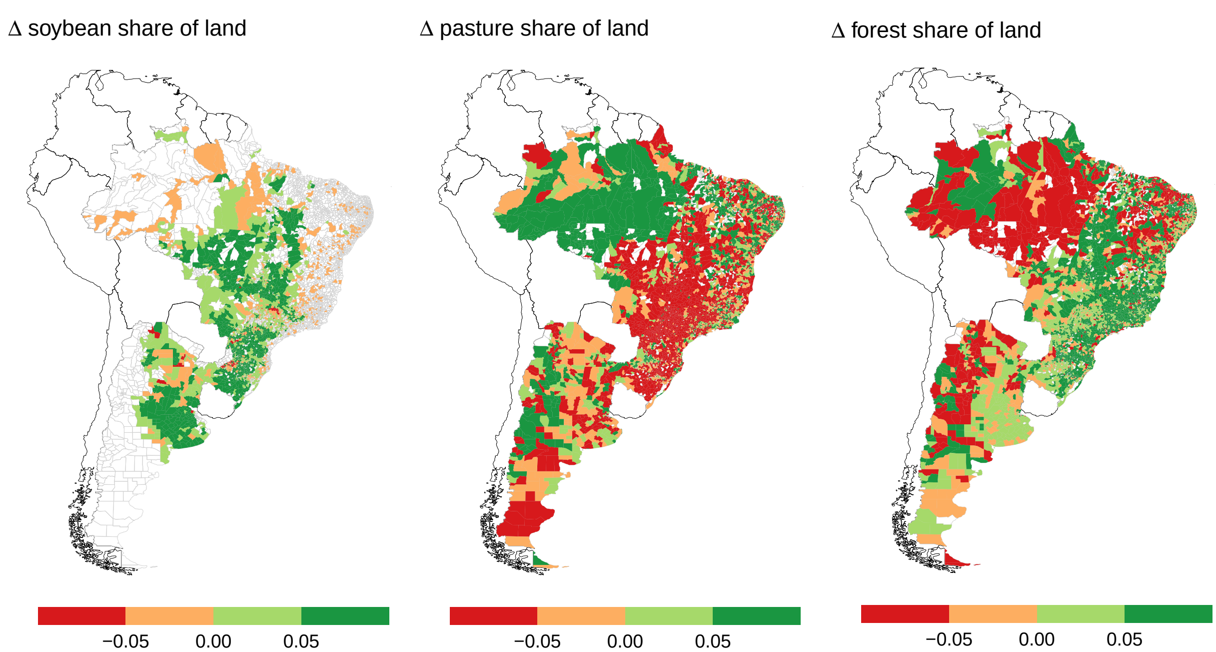

Notes: The figure shows changes in acreage allocated to soybeans, pasture, and forest (1995-2017). Maps are constructed using land-use data from the decadal agricultural censuses of Argentina and Brazil.

Fact 3: 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 less than 50 head of cattle, facing a concentrated sector of downstream agribusiness firms.[2] In the median county, the top three agribusiness firms account for 95% of sourced beef, with the top firm accounting for over 60% (Table 1). JBS, the industry leader, accounted for 36% of purchases nationwide, sourcing from 46% of all counties, and with a median market share of 28% (Table 2).

Table 1: Agribusiness concentration measures (2017).

	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)	118	110	181	34
Number of source counties	2,803	807	1,390	207
Number of destination countries	130	85	69	78
CR-1 (national market)	0.36	0.18	0.16	0.13
CR-3 (national market)	0.69	0.46	0.43	0.36
CR-1 (local market, median)	0.63	1.00	1.00	0.21
CR-3 (local market, median)	0.95	1.00	1.00	0.51
Share of source counties with 1 agribusiness firm	0.13	0.86	0.77	0.07

Notes: Local markets are defined at the county-level. Sources: 2017-2018 agricultural censuses and TRASE.

Table 2: Major agribusiness firms (2017).

			Market share of firm		Counties sourced by firm	
Country	Commodity	Firm	National	Local*	Number	As share of all counties
Brazil	Beef	Jbs	0.36	0.28	1,284	0.46
		Marfrig	0.18	0.30	1,146	0.41
		Minerva	0.15	0.19	1,104	0.39
	Maize	Cargill	0.18	0.47	50	0.06
		Bunge	0.16	0.32	45	0.06
		Amaggi	0.12	0.32	34	0.04
	Soybean	Bunge	0.16	1.00	176	0.13
		Cargill	0.14	1.00	194	0.14
		Adm	0.13	0.93	100	0.07
Argentina	Soybean	Vicentin	0.13	0.13	168	0.81
		Cargill	0.12	0.09	190	0.92
		Bunge	0.12	0.11	182	0.88

*Notes: The reported local market share is the firm's median share across all the counties it sources from.

[2] Among Brazil's ranching establishments, 73% hold less than 50 ha, 79% hold less than 50 head of cattle, and 76% are family farms. See Appendix Figure 14. Sources: IBGE Table 6783, Table 6910, and Table 6783.

Figure 5 shows how agribusiness concentration varies across space and how it correlates with a crude accounting-based measure of a markdown—the ratio of the farm-gate price received by the farmer with respect to the price the agribusiness firm receives at the port. Farm-gate prices are subject to wider markdowns in production locations with a higher concentration of buyers, even after taking into account differences in a location's remoteness and implied transport costs. Needless to say, these stylized facts 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).[3]

Figure 5: Agribusiness concentration (beef markets, Brazil, 2017).

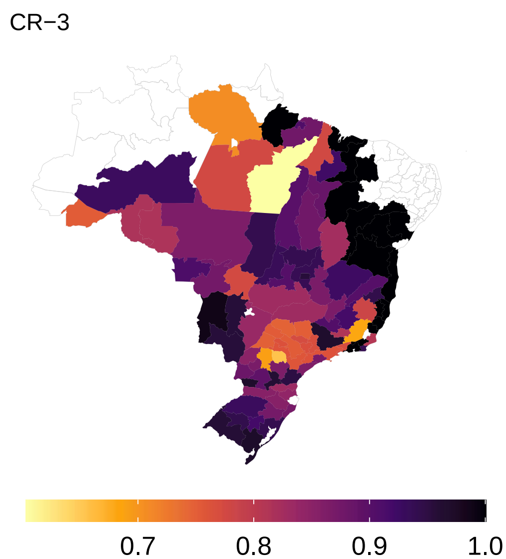

Notes: Maps are displayed at the mesoregion level for ease of visualization. For the scatterplot, each grey bubble is a mesoregion (with size proportional to beef output), while dark bubbles are a binscatter overlay. The reported slope and R^2 are from an OLS regression of the farm-gate/agribusiness price ratio on the concentration ratio and a measure of market access to control for remoteness. Appendix Table 8 shows the detailed results for these regressions.

Implications for model specification

Fact 1 suggests deforestation, commodity choice, and the geographic location of production are the key margins driving agricultural emissions. Therefore, I propose a model where farmers make decisions along two separate margins: first, how much land to clear, and second, which specific commodity to produce on it. Furthermore, the model has high geographic resolution to incorporate rich spatial heterogeneity in agricultural productivity and emissions intensity.

Fact 2 suggests international demand shocks are a major driver of farmers' land use decisions, and hence trade policy can serve as environmental policy. I therefore allow for output to be pur-

[3] The correlation between concentration and markups can be positive or negative depending on how such primitives are chosen (Syverson, 2019). Higher concentration is associated with higher profit margins in models with a fixed number of firms, such as the standard Cournot model. However, allowing entry the relationship can be reversed: markets with low profit margins can be highly concentrated because the gains for potential entrants are too small.

chased by both domestic and foreign consumers, opening the door for consumption leakage from incomplete regulation. Fact 2 also suggests different commodities compete with each other in local land markets. Thus, a positive demand shock for a specific commodity raises land prices in locations where its production is most suitable, displacing other commodities to locations with cheaper land. I incorporate this mechanism by having farmers choose among different commodities, as mentioned previously, and by having commodity markets clear locally at the county-level.

Fact 3 suggests agribusiness firms may plausibly hold buyer market power over farmers, and the extent of such power may vary across space. Therefore, I model agribusiness firms as oligopsonists in local upstream markets. I do so in a way that nests the perfectly competitive case in order to alternate between conduct assumptions, and thus evaluate how market structure alters the transmission of environmental policy to upstream farmers.

Implications for policy feasibility

Beyond the political constraints a government might face when considering a first-best carbon tax on farmers, there are also logistical hurdles to enforcing such a tax that are unique to agriculture. Many of these challenges stem from Fact 1: agricultural emissions are dispersedly generated across millions of upstream farmers with heterogeneous carbon footprints. By contrast, in sectors such as electricity generation the emissions from burning fossil fuels are concentrated downstream among a few large firms, making enforcement at the emissions source logistically easier.

In our setting, a first-best carbon tax could be implemented as an output tax that varies by origin of production due to the wide spatial heterogeneity in carbon intensity. For example, beef from the Amazon would be subject to a higher tax per ton of output than elsewhere because of the high carbon density of the land it was produced on. Levying such a tax directly on upstream farmers would require enough state capacity to successfully enforce tax compliance at high spatial resolution. Alternatively, the tax could be levied downstream on agribusiness firms, but this would require reliable tracing of the commodity's origin to determine its carbon content and set the appropriate output tax rate, i.e., an effective certification scheme. Recent empirical work on policy compliance suggests these requirements are unlikely to be met in our setting, rendering the successful implementation of a first-best carbon tax infeasible.[4] In reality, we observe a range of command-and-control and market-based policies, all of which are second-best.

Command-and-control policies often take the form of conservation zones, which due to their high enforcement costs typically only target a narrow geographic subset of high-carbon density areas.[5] By contrast, market-based tools such as carbon taxes are appealing because they avoid the

[4]For empirical evidence on the difficulties Brazilian environmental authorities have in effectively monitoring upstream farmers, as well as the practical challenges meat-packers face in tracing the ultimate origin (and hence, carbon content) of their cattle purchases, see Barreto, Pereira, Jr. and Baima (2017); Pereira, Rausch, Carrara and Gibbs (2020); Skidmore, Moffette, Rausch, Christie, Munger and Gibbs (2021), among many others. Among different agricultural commodities, the beef supply chain is notorious for how long and fragmented it is, making certification across its multiple stages and actors especially challenging. See Appendix section F.2 for more details.

[5]One of the better known command-and-control policies in our setting is Brazil's "Priority Municipality List". Under this policy, IBAMA (the Brazilian environmental protection agency) increased monitoring and enforcement efforts

implementation costs of direct regulation by working through the market mechanism, thus allowing for broader geographic coverage. However, in practice such taxes are not levied directly on the farmers whose incentives they aim to correct, for the reasons discussed in the previous paragraph. Instead, second-best proposals typically suggest implementing the corrective tax where the supply chain becomes concentrated—downstream on the agribusiness firms at the port—because it can be easily implemented at low administrative cost as an export tax. The shortcoming of such proposals is that due to difficulties in certifying the origin of a given commodity once it arrives at the port, a national average emissions footprint would be used to determine the size of the corrective tax, regardless of whether the commodity was produced on land with high or low carbon density. Hence, such downstream taxes can be poorly targeted.

To conclude, first-best carbon taxes at the upstream source of the externality are not considered feasible solutions in the context of South American agricultural emissions. Second-best policies, both command-and-control and market-based, are more frequently proposed in policy circles and are often traded-off against one another. The former are well-targeted, but come at a high enforcement cost and only cover a narrow geographic area. The latter have broader coverage and face low implementation costs because they operate through the market mechanism, but they can be spatially mistargeted because they are not levied directly at the emissions source. Hence, there is a trade-off between targeting and enforcement costs when choosing between command-and-control and market-based tools. The features of my model allow me to show how this trade-off depends on i) the degree of upstream heterogeneity in carbon footprints, and ii) whether the market faces pre-existing distortions such as market power.

4 Model

On the supply side, atomistic farmers choose between leaving their land in its natural forested state or converting it to agricultural use and producing a specific commodity. Final demand consists of consumers distributed across domestic and foreign markets. However, farmers do not access consumer markets directly, but instead sell their output to intermediating agribusiness firms.

4.1 Supply side: Upstream farmers

Land use decision. Each county i contains a continuum of fields indexed ω. Each field ω is owned by a farmer, who chooses a land use from a discrete choice set consisting of a natural-use option $\mathcal{N}$ and a nest of agricultural commodities $\mathcal{C}$. Field ω's output of commodity $c \in \mathcal{C}$ is,

$$Q_i^c(\omega) = A_i^c(\omega) L_i^c(\omega) \quad \text{with} \quad A_i^c(\omega) = A_i^c \exp(\varepsilon_i^c(\omega)), \tag{1}$$

where $A_i^c(\omega)$ is the field's productivity in commodity c and $L_i^c(\omega)$ is its size. A field's productivity is decomposed into a county-level mean A_i^c and a field-level idiosyncratic shock $\varepsilon_i^c(\omega)$. If p_i^c is

in a selected sample of high deforestation risk municipalities from the Amazon. Assunção, McMillan, Murphy and Souza-Rodrigues (2019) find the policy reduced deforestation by 43 percent.

the commodity's farm-gate price, then the payoff per unit of land allocated to commodity c is $p_i^c A_i^c(\omega)$. Let $A_i^{\mathcal{N}}(\omega)$ denote the payoff per unit of land when left in its natural state, which is also decomposed into a county-level mean $A_i^{\mathcal{N}}$ and a field-level idiosyncratic shock $\varepsilon_i^{\mathcal{N}}(\omega)$.

Payoff from natural use. The payoff from allocating land to a commodity c is market-based and observable from data: it is a dollar-value constructed from market prices and yields. This differs from the payoff to natural use because farmers are generally not paid to keep their land forested, and even in cases in which they are, we typically don't have comprehensive data on such payments. Therefore, the interpretation of $A_i^{\mathcal{N}}(\omega)$ is that it captures the dollar-value of any incentives farmers have to keep part of their land forested, and which are unobserved to the econometrician. Such incentives may be pecuniary or non-pecuniary, and static or dynamic. Examples of static incentives are: the aesthetic value of trees to landowners, unobserved forestation payments, or non-pecuniary benefits (e.g., prevention of soil erosion), as in Souza-Rodrigues (2019). An example of a dynamic incentive is the option value of deforesting in the future, as in Scott (2013).

Nesting assumption. The nesting assumption is that the field-level idiosyncratic shocks are correlated between agricultural commodities, but not between a commodity and the natural-use option. There are two crucial parameters to keep track of: θ governs the dispersion of shocks across fields, while $\lambda \in (0,1)$ governs the correlation of shocks between commodities. Higher values of θ correspond to lower dispersion across fields, and higher values of λ correspond to lower correlation between commodities.[6] Under these distributional assumptions, the probability commodity c is chosen, conditional on the farmer choosing the agricultural nest $\mathcal{C}$, is given by,

$$\pi_i^{c|\mathcal{C}} = \frac{\left(p_i^c A_i^c\right)^{\frac{\theta}{\lambda}}}{\sum_{c' \in \mathcal{C}} \left(p_i^{c'} A_i^{c'}\right)^{\frac{\theta}{\lambda}}}, \tag{2}$$

while the choice probability of the agricultural nest $\mathcal{C}$ is,

$$\pi_i^{\mathcal{C}} = \frac{\left(P_i^{\mathcal{C}}\right)^{\lambda}}{(A_i^{\mathcal{N}})^{\theta} + \left(P_i^{\mathcal{C}}\right)^{\lambda}} \quad \text{with} \quad P_i^{\mathcal{C}} \equiv \sum_{c' \in \mathcal{C}} \left(p_i^{c'} A_i^{c'}\right)^{\frac{\theta}{\lambda}}. \tag{3}$$

$P_i^{\mathcal{C}}$ is defined as the payoff of the agricultural nest as a whole, since it is an index comprising the returns of all the nest's commodities—technically, $\ln P_i^{\mathcal{C}}$ is the nest's inclusive value in the nested

[6]Formally, we have a nested logit model of land-use with the following log returns per hectare of land,

$$r_i^k(\omega) = \begin{cases} \theta \ln\left(p_i^c A_i^c\right) + \varepsilon_i^c(\omega)^* & \text{if } k = c \in \mathcal{C} \\ \theta \ln\left(A_i^{\mathcal{N}}\right) + \varepsilon_i^{\mathcal{N}}(\omega)^* & \text{if } k = \mathcal{N}. \end{cases}$$

$\varepsilon_i^c(\omega)$ is distributed type I EV with location parameter 0 and standard deviation $\sigma \frac{\pi}{\sqrt{6}}$, which is equivalent to having $A_i^c(\omega)$ distributed type II EV (Fréchet) with location parameter 0, scale parameter $\Gamma\left(1-\sigma\right)^{-1} A_i^c$ and shape parameter $\theta \equiv \sigma^{-1}$. Either case implies $E[A_i^c(\omega)] = A_i^c$. The type I EV formulation conveniently casts the nested choice problem as a nested logit model. Notice that we have rescaled payoffs by σ^{-1}, so that $\varepsilon_i^k(\omega)^*$ is a standardized type I EV error: its location parameter is 0 and its standard deviation is $\frac{\pi}{\sqrt{6}}$.

logit model. The share allocated to natural use is $\pi_i^{\mathcal{N}} = 1 - \pi_i^{\mathcal{C}}$. The nested structure implies the unconditional choice probabilities, which map to land shares in the data, can be written as $\pi_i^c = \pi_i^{c|\mathcal{C}} \pi_i^{\mathcal{C}}$. If the county's total surface is $\bar{L}_i$, then the county's total acreage of commodity c is $L_i^c = \pi_i^c \bar{L}_i$. Finally, we have a closed form expression for the county-level supply of commodity c,

$$Q_i^c = \int_\omega Q_i^c(\omega) d\omega = A_i^c \left(\pi_i^{c|\mathcal{C}}\right)^{\frac{\lambda}{\theta}} L_i^c. \tag{4}$$

Notice that as $\lambda \to 1$, correlation between commodities goes to zero and the nested model collapses to a multinomial model, a common specification in Ricardian models of agricultural trade (Costinot et al., 2016; Sotelo, 2020). The nested structure is important for my setting because a multinomial model would restrict substitution between commodities to be just as easy as substitution between natural and agricultural use.[7] First, such restrictions are unrealistic if we expect land clearing to be costlier than switching between commodities on existing agricultural land. Second, disentangling the two margins allows for evaluation of policies operating through one margin but not the other. For example, the substitution margin within the agricultural nest matters for evaluating commodity-specific policies, such as maize-ethanol subsidies. By contrast, the impact of deforestation fines, which are commodity-blind because they disincentivize agriculture as a whole, are determined by the across-nest substitution margin.

Key supply-side elasticities. From 4 we can derive the price-elasticity of output,

$$\frac{\partial \ln Q_i^c}{\partial \ln p_i^c} = \left(\frac{\theta}{\lambda} - 1\right)\left(1 - \pi_i^{c|\mathcal{C}}\right) + \theta \pi_i^{c|\mathcal{C}} \left(1 - \pi_i^{\mathcal{C}}\right). \tag{5}$$

Notice supply becomes more elastic when $\theta \to \infty$ or $\lambda \to 0$. To understand why, recall θ governs the dispersion of productivity *across fields*: as $\theta \to \infty$, marginal fields become identical to inframarginal fields, so county-level supply curves become flat. On the other hand, λ governs the correlation of productivity *between commodities*: as $\lambda \to 0$ correlation becomes perfect, implying all fields order their commodity choices in the same way (although the levels of payoffs may differ across fields). Because all fields make the same commodity choice, heterogeneity across fields disappears, and we obtain a flat supply curve at the county-level. Therefore, a highly elastic supply curve can be explained by low dispersion of productivity across fields (high θ) or high correlation of productivity between commodities (low λ). To separate the role of θ from λ it is useful to

[7] In the nested model, substitution patterns between commodities are stronger than between a commodity and natural use: $\left|\frac{d \ln \pi_i^c}{d \ln p_i^{c'}}\right| > \left|\frac{d \ln \pi_i^{\mathcal{N}}}{d \ln p_i^{c'}}\right|$ for $c \neq c'$. Notice the proportional substitution property holds within-nest $\mathcal{C}$ but not across nests: $\frac{d \ln \pi_i^{c|\mathcal{C}}}{d \ln p_i^{c'}} = -\frac{\theta}{\lambda} \pi_i^{c'|\mathcal{C}}$. A multinomial model is more restrictive since it imposes proportional substitution across all choices, including the natural-use option.

consider the odds ratios within- and across-nests,

$$\ln\left(\frac{\pi_i^{c|\mathcal{C}}}{\pi_i^{c'|\mathcal{C}}}\right) = \frac{\theta}{\lambda}\ln\left(\frac{p_i^c A_i^c}{p_i^{c'} A_i^{c'}}\right) \quad \text{and} \quad \ln\left(\frac{\pi_i^{\mathcal{C}}}{\pi_i^{\mathcal{N}}}\right) = \lambda \ln\left(\frac{\sum_{c\in\mathcal{C}}\left(p_i^c A_i^c\right)^{\frac{\theta}{\lambda}}}{\left(A_i^{\mathcal{N}}\right)^{\frac{\theta}{\lambda}}}\right). \tag{6}$$

The elasticity of substitution within-nest, $\frac{\theta}{\lambda}$, can be high because of low dispersion across fields (high θ) or high correlation between commodities (low λ) for the reasons mentioned in the preceding paragraph. The elasticity of substitution across nests, which we interpret as the deforestation elasticity, is equal to λ. What is the intuition for why λ, the correlation of productivity between commodities, is connected to the deforestation elasticity? As $\lambda \to 0$ the correlation across commodities increases, hence the within-nest heterogeneity falls relative to the across-nest heterogeneity. As fields become relatively more heterogeneous along the across-nest margin, county-level supply curves (of agriculture as a whole) become less elastic, i.e., the deforestation elasticity falls.[8] To conclude, consider the deforestation elasticity in terms of levels rather than shares, i.e., how the amount of land allocated to the agricultural nest, $L_i^{\mathcal{C}}$, responds to the payoff of agriculture as a whole, rather than to the payoff of any individual commodity. To do so, we use $P_i^{\mathcal{C}}$ from equation 3 as the "price" of the agricultural nest, and then derive the "price"-elasticity of agricultural land,

$$\frac{\partial \ln L_i^{\mathcal{C}}}{\partial \ln P_i^{\mathcal{C}}} = \lambda(1-\pi_i^{\mathcal{C}}). \tag{7}$$

Since increases in $L_i^{\mathcal{C}}$ necessarily reduce natural land use, 7 can be interpreted as a deforestation elasticity. As mentioned before, the parameter governing this across-nest adjustment margin is λ.

4.2 Demand side: Downstream agribusiness intermediaries and final consumers

The demand side consists of two stages along the supply chain. First, agribusiness intermediaries buy commodities from upstream farmers in sources indexed $i \in \mathcal{I}$. Second, these intermediaries sell the commodity 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. I abstract from market power of intermediaries in their role as sellers for two reasons. First, the environmentally-relevant decisions are made by farmers that are upstream of the intermediaries, not by downstream consumers. Hence, when considering environmental regulations implemented on agribusiness firms, what matters is how market power affects the upstream transmission of such policies rather than the downstream transmission. Second, my data is simply not rich enough to also incorporate downstream market power.[9]

[8] The right hand side of equation 6 says that for a fixed elasticity of substitution within-nest $\frac{\theta}{\lambda}$, an increase in λ results in a higher deforestation elasticity. It is key to notice we are keeping $\frac{\theta}{\lambda}$ fixed when changing λ: an increase in λ must therefore be matched with an increase in θ, which implies we are reducing heterogeneity across fields within a county. As fields become more homogeneous, county-level supply becomes more elastic, and hence deforestation elasticities increase.

[9] I observe how much of a commodity is purchased by agents in the first stage of the destination market's supply chain, which are mostly food processing companies rather than final retail consumers. Understanding impacts on

Agribusiness intermediaries. There are N_i^c identical intermediary firms, each purchasing q_i^c units of 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_i^c. Apart from transporting the commodities from source to destination, we allow firms to add value by transforming the commodity into a processed version (e.g., soybeans into soybean oil) by using a technology $f_c(q_i^c)$. Firms then sell the processed version at the port closest to source county i, obtaining a free-on-board price $\bar{p}_i^c$. Hence, the transport cost from the port to the final destination is paid by final consumers: a destination j consumer pays $p_{ij}^c = \bar{p}_i^c \tau_{ij}^c$, where τ_{ij}^c is an iceberg trade cost. We can now pose each firm's maximization problem, taking demand of the other firms as given,

$$\max_{q_i^c} \quad \bar{p}_i^c f_c(q_i^c) - p_i^c\left(Q_i^c\right) q_i^c,$$

where q_i^c is an individual firm's demand, Q_i^c is total demand from source i, and $p_i^c\left(Q_i^c\right)$ is source i's inverse supply equation. From the first order conditions we obtain the farm-gate price is a fraction μ_i^c of the marginal revenue it generates for the intermediary,

$$\frac{p_i^c}{\bar{p}_i^c f_c'(q_i^c)} = \underbrace{\left(1 + \frac{1}{\epsilon_i^c N_i^c}\right)^{-1}}_{\equiv \mu_i^c} \quad \text{where} \quad \frac{1}{\epsilon_i^c} \equiv \frac{\partial \ln p_i^c}{\partial \ln Q_i^c}. \tag{8}$$

A farmer from source i obtains μ_i^c cents for every dollar the intermediary makes from the commodity. I define μ_i^c, the ratio of the input's price to its marginal revenue, as the "markdown-wedge". Throughout the rest of the paper, when using the term "markdown" I am referring to this "markdown-wedge". Intuitively, markdowns follow an inverse-elasticity rule: sources with inelastic supply (high $\frac{1}{\epsilon_i^c}$) are subject to large markdowns (low μ_i^c). Markdowns are also larger in sources with few competing firms (low N_i^c).

The setup of the intermediary problem is purposefully simple—firms are identical and there is no entry—the goal being to obtain the smallest departure from the perfectly competitive setting typically assumed by the agricultural trade literature as well as to parsimoniously nest it. Perfect competition is obtained by imposing $\mu_i^c = 1$, and trade of unprocessed commodities by imposing $f_c(q_i^c) = q_i^c$. In these limiting cases, the farm-gate price is equal to the free-on-board price, and the destination market price is simply the farm-gate price adjusted by trade costs, i.e., $p_{ij}^c = p_i^c \tau_{ij}^c$. I discuss how the model can be extended to admit firm heterogeneity and exit/entry, and the implications of doing so, in Appendix D.3. Most importantly, these extensions do not change the qualitative insights of the paper.

Consumers. To interpret what destination market "consumers" are in this model, it is worth clarifying how we measure them in the data. The demand-side data measures how much of a com-

final consumers would require a host of assumptions about how food-processors transform the commodity into differentiated retail food products (production function specification, conduct assumptions, mark-ups), all of which are untestable with the available data and beyond the scope of this paper.

modity arrives at the destination port, and not how much is purchased by final consumers at the final retail stage. Therefore, "consumers" in this model should be interpreted as the agents in the first stage of the destination market's supply chain, which are mostly food processing companies that transform the commodity into retail food products.

The modeling of consumers is the most standard part of the model: I use a three-level CES demand system. In the upper level, they substitute between commodities (e.g., maize vs. wheat). In the middle level, they substitute between source nations of a given commodity (e.g., Brazilian maize vs. US maize). In the lower level, they substitute between counties within a nation (e.g., maize from Northern Brazil vs. maize from Southern Brazil). The lower level is necessary to obtain demand at the county-level.

Given our interpretation of what a "consumer" is, 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 final retail consumers perceive significant quality differences between maize from one county versus another. Instead, the different levels should be interpreted as the degree to which a food processor substitutes inputs across different sources. Hence, the lower-level reflects the degree to which food processors perceive the process of sourcing from one county versus another as differentiated, even if the underlying product being sourced from both counties is identical. Concretely, I assume each destination j has a representative consumer with the following three-level CES utility function,

$$U_j = \left(\sum_c (a_j^c)^{\frac{1}{\eta_u}} (C_j^c)^{\frac{\eta_u - 1}{\eta_u}} \right)^{\frac{\eta_u}{\eta_u - 1}}, \text{ where } C_j^c = \left(\sum_n (a_{nj}^c)^{\frac{1}{\eta_m}} (C_{nj}^c)^{\frac{\eta_m - 1}{\eta_m}} \right)^{\frac{\eta_m}{\eta_m - 1}} \text{ and } C_{nj}^c = \left(\sum_{i \in n} (C_{ij}^c)^{\frac{\eta_l - 1}{\eta_l}} \right)^{\frac{\eta_l}{\eta_l - 1}}.$$

C_j^c is consumption of good c aggregated across source nations indexed n. C_{nj}^c is consumption of good c aggregated across source counties indexed i belonging to nation n. η_u is the upper elasticity of substitution between goods, η_m is the middle elasticity of substitution between source nations, and η_l is the lower elasticity of substitution between counties within a nation. The a's are preference shifters across goods and sources. These preferences deliver the following county-level demand equation,

$$C_{ij}^c = \left(\frac{p_{ij}^c}{P_{nj}^c} \right)^{-\eta_l} a_{nj}^c \left(\frac{P_{nj}^c}{P_j^c} \right)^{-\eta_m} a_j^c \left(\frac{P_j^c}{P_j} \right)^{-\eta_u} \frac{X_j}{P_j} \quad \forall i \in n, \tag{9}$$

where X_j is destination j income, and price indices for each level are defined as follows,

$$P_{nj}^c \equiv \left(\sum_{i \in n} (p_{ij}^c)^{1-\eta_l} \right)^{\frac{1}{1-\eta_l}} \quad P_j^c \equiv \left(\sum_n a_{nj}^c (P_{nj}^c)^{1-\eta_m} \right)^{\frac{1}{1-\eta_m}} \quad P_j \equiv \left(\sum_c a_j^c (P_j^c)^{1-\eta_u} \right)^{\frac{1}{1-\eta_u}}.$$

4.3 Equilibrium

An equilibrium is a set of farm-gate prices $\{p_i^c\}_{i,c}$ such that supply in each county is equal to the total demand from that county, and this holds for every commodity,

$$Q_i^c(p_i^c) = \sum_j C_{ij}^c(p_{ij}^c)\tau_{ij} \quad \forall i, c, \quad \text{where} \quad p_{ij}^c = \frac{p_i^c}{\mu_i^c}\frac{\tau_{ij}}{f_c'(q_i^c)}.$$

It is worth stressing that market clearing occurs county-by-county, i.e., at a sub-national level. This is key to allow for within-nation spatial heterogeneity in the effectiveness and incidence of environmental policy. If the downstream carbon taxes are uniform, i.e., they cannot take into account the sub-national origin of the commodity because of certification challenges, then they will be spatially mistargeted because of the wide geographic heterogeneity in carbon density within a nation. Quantifying the extent to which they are mistargeted requires understanding how quantities differentially respond to policy across sub-national markets.

5 Estimation

In section 4, the model is presented for expositional purposes without time subscripts because it is static. I now explicitly introduce time subscripts given I will combine cross-sectional and temporal variation for estimation. A time period is a decade because that is the frequency of the census data. The observed outcomes across time are therefore interpreted through the lens of the model as a sequence of static equilibria separated by a substantial time lag.[10] The static model is therefore used as a first approximation for studying decisions with long time lags, in part because dynamic considerations such as switching costs become less relevant the longer the temporal horizon is. Examples of recent studies taking such an approach—estimating static discrete choice models using data with decadal frequency—are Diamond (2016) and Donaldson (2018).

5.1 Supply elasticities

To understand the variation in the data that is used to estimate the supply-side parameters, it is useful to consider the odds ratio between two commodities c and c' within the agricultural nest,

$$\ln\left(\frac{\pi_{it}^c}{\pi_{it}^{c'}}\right) = \frac{\theta}{\lambda}\ln\left(\frac{p_{it}^c A_i^c}{p_{it}^{c'} A_i^{c'}}\right) + u_{it}^{cc'}, \tag{10}$$

where π_{it}^c is county i's land share in commodity c at time t, p_{it}^c is the farm-gate price, A_i^c is the county's mean productivity, and $u_{it}^{cc'}$ is an unobservable error term. Given 10 is a supply equation of commodity c relative to c', we interpret $u_{it}^{cc'}$ as an unobservable supply shifter of c relative to

[10] Such an interpretation is common in the spatial economics literature, where having rich spatial heterogeneity is the main priority when choosing a dataset, but which often comes at the cost of lacking the requirements for estimating a fully dynamic model. This is common for census data, which has rich cross-sectional heterogeneity but individual decision makers cannot be linked across time and the frequency of the data is too low to incorporate full dynamics.

c'. The ratio of parameters $\frac{\theta}{\lambda}$ is the elasticity of substitution between commodities, i.e., within the agricultural nest. A useful interpretation of $\frac{\theta}{\lambda}$ is as a supply elasticity of c relative to c'. This elasticity is large when productivity dispersion across fields goes to zero ($\theta \to \infty$) because marginal fields in a county become identical to infra-marginal fields, resulting in a flat supply curve. This elasticity can also be large when productivity is perfectly correlated across commodities ($\lambda \to 0$).[11] Thus, knowing this supply elasticity is not enough to separately identify θ from λ. The additional restriction that is needed for identification exploits variation *across nests*: it is the odds ratio between nest $\mathcal{C}$ and natural use $\mathcal{N}$,

$$\ln\left(\frac{\pi_{it}^{\mathcal{C}}}{\pi_{it}^{\mathcal{N}}}\right) = \lambda \ln\left(P_i^{\mathcal{C}}\right) - \theta \ln\left(A_i^{\mathcal{N}}\right) + u_{it}^{\mathcal{CN}} \quad \text{with } P_i^{\mathcal{C}} \equiv \sum_{c\in\mathcal{C}} \left(p_{it}^c A_i^c\right)^{\frac{\theta}{\lambda}}, \tag{11}$$

where $\ln P_i^{\mathcal{C}}$ is the inclusive value of the agricultural nest in county i at time t, $\theta \ln\left(A_i^{\mathcal{N}}\right)$ is unobservable and time-invariant so it is estimated as a county fixed effect, and $u_{it}^{\mathcal{CN}}$ is an unobservable supply shifter of agricultural land relative to forested land. The parameter λ is the substitution elasticity across nests, i.e., the deforestation elasticity. Given the composite parameter $\frac{\theta}{\lambda}$ estimated from 10, we can construct the inclusive value term in 11, and then λ is identified.

Instruments. OLS estimates will be biased towards zero due to simultaneity bias, specifically because the unobservable supply shocks $u_{it}^{cc'}$ will be correlated with relative land shares and relative returns. For example, if the unobservable local productivity of commodity c increases relative to c', its relative land share would increase and its relative price would drop, biasing the estimate of $\frac{\theta}{\lambda}$ downwards. Since we are estimating a supply equation, an appropriate instrument is a demand shifter varying at the county-year it and commodity-pair cc' level. I construct such an instrument from the export network data as follows,

$$z_{it}^{cc'} = \sum_j s_{ij}^{cc'} d_{jt}^{cc'} \quad \text{with } s_{ij}^{cc'} \equiv \frac{s_{ij}^c}{s_{ij}^{c'}}, \quad d_{jt}^{cc'} \equiv \frac{d_{jt}^c}{d_{jt}^{c'}}, \tag{12}$$

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 for commodity c in destination j. Intuitively, if demand conditions for crop c relative to c' increase in destination j, counties that historically supplied j are more exposed and receive larger demand shocks. I use destination j's imports from every nation except Argentina and Brazil as the demand measure d_{jt}^c, thus purging away supply-side effects in Argentina and Brazil that directly affect the imports of j. The identifying assumption is that the exposure measure $s_{ij}^{cc'}$ is uncorrelated with *changes* in the error term $\Delta u_{it}^{cc'}$, whereas correlation with *levels* $u_{it}^{cc'}$ is allowed (Goldsmith-Pinkham, Sorkin and Swift, 2020).[12]

[11] Zero field dispersion means all fields in a county are identical; that is, any two fields ω and ω' satisfy $A_i^c(\omega') = A_i^c(\omega)\ \forall c$. Perfect correlation across commodities is a weaker restriction because it allows any two fields ω and ω' to be different in the sense that $A_i^c(\omega) \neq A_i^c(\omega')\ \forall c$, but it restricts every field to have the same ordering over commodities. Because ordering is all that matters in discrete choice problems, all fields choose the same commodity, and within-county heterogeneity in choices disappears.

[12] The assumption allows for counties with high unobservable productivity of c relative to c' to selectively export

We also need an instrument for equation 11 for the same reasons as for equation 10: if unobserved agricultural productivity increases overall for all commodities, then the agricultural nest's share would increase and its price index $P_i^{\mathcal{C}}$ would decrease, biasing the estimate of λ towards zero. We now need a shifter of demand for agriculture overall. We construct it similarly to 12, but now summing across all commodities,

$$z_{it}^{\mathcal{C}} = \sum_j \sum_{c \in \mathcal{C}} s_{ij}^c d_{jt}^c \tag{13}$$

Results. OLS and IV estimates of $\frac{\theta}{\lambda}$, the substitution elasticity between commodities, are shown in columns 1 and 2 of Table 3. The values are comparable to trade elasticity estimates from Ricardian models of agriculture, which are estimated from variation across commodities within the agricultural nest and abstract from deforestation, typically ranging between 1.5-4 (Costinot et al., 2016; Pellegrina, 2019; Sotelo, 2020). Columns 3 and 4 add interactions with a frontier region indicator to allow for spatial heterogeneity. The negative sign of the interaction coefficients imply substitution across commodities is costlier for farmers in frontier regions.

Estimates of λ, the deforestation elasticity, are shown in Table 4 and are broadly consistent with the agricultural and empirical IO literature. Given the long time lags in my decadal data, I interpret my estimates as long-run elasticities. Scott (2013) finds long-run elasticities of 0.3 by estimating a dynamic model on annual data from the United States. Berry and Schlenker (2011) and Roberts and Schlenker (2013) both estimate land-use change elasticities for Brazil between 0.2-0.4. Within Brazil and focusing on the Amazon biome, Souza-Rodrigues (2019) finds land-use change elasticities near zero. My estimates for frontier regions such as the Amazon are in line with these findings, given they are substantially lower than those from core agricultural regions.[13]

Role of the nesting structure. How important is the nesting structure? Estimating a non-nested multinomial model amounts to imposing $\lambda = 1$ and estimating a single parameter θ from variation between commodities and natural use simultaneously. Table 5 shows the results for such a model. The multinomial model mixes variation within and across nests to deliver a single elasticity: notice the multinomial OLS estimates are sandwiched between the nested model's within- and across-nest elasticities. In this case, the land-use change elasticity is restricted to equal the substitution elasticity between commodities,

$$\frac{d \ln \pi_i^c}{d \ln p_i^c} = \theta(1 - \pi_i^c).$$

Because estimates of θ are well above 1 in the trade literature, the estimates from the multinomial model are hard to reconcile with land-use change elasticities that are found to be well below 1 in the agricultural and empirical IO literature. The trade elasticities are estimated from substitution

to specific destinations. It also allows counties that historically exported to specific destinations to experience faster growth in their overall unobservable productivity, but not in their *relative* productivity of commodity c relative to c'.

[13] Souza-Rodrigues (2019) uses cross-sectional data, so his static framework is appropriate to estimate long-run elasticities.

Table 3: Nested model - substitution elasticity between commodities (within-nest).

	OLS (1)	IV (2)	OLS (3)	IV (4)
$\frac{\theta}{\lambda}$	0.768***	2.243***	0.641***	2.316***
	(0.080)	(0.135)	(0.093)	(0.146)
$\frac{\theta}{\lambda}$ × frontier region			0.681***	−0.844***
			(0.168)	(0.198)
Time FE	X	X	X	X
Location FE	X	X	X	X
Observations	5,618	5,618	5,618	5,618
Adjusted R^2	0.197	0.104	0.200	0.100
Notes:	First stage F-statistic = 105.6. SE clustered at county level.			

Table 4: Nested model - deforestation elasticity (substitution elasticity across nests).

	OLS (1)	IV (2)	OLS (3)	IV (4)
λ	−0.040**	0.281*	−0.042**	0.341*
	(0.018)	(0.155)	(0.020)	(0.185)
λ × frontier region			0.009	−0.298**
			(0.034)	(0.152)
Time FE	X	X	X	X
Location FE	X	X	X	X
Observations	7,630	7,630	7,630	7,630
Adjusted R^2	0.661	0.638	0.661	0.636
Notes:	First stage F-statistic = 53.47. SE clustered at county level.			

Table 5: Multinomial model - single substitution elasticity between all land uses.

	OLS (1)	IV (2)	OLS (3)	IV (4)
θ	0.472***	0.517***	0.472***	0.535***
	(0.010)	(0.010)	(0.011)	(0.011)
θ × frontier region			−0.003	−0.065**
			(0.029)	(0.030)
Time FE	X	X	X	X
Location FE	X	X	X	X
Observations	10,846	10,846	10,846	10,846
Adjusted R^2	0.480	0.476	0.480	0.474
Notes:	First stage F-statistic = 130.27. SE clustered at county level.			

between commodities on existing agricultural land, ignoring the extensive margin of switching from non-agricultural to agricultural land. By contrast, the agricultural and empirical IO literature often ignores substitution between commodities within agriculture in order to focus on the binary extensive margin (Scott, 2013; Souza-Rodrigues, 2019). One would expect the extensive-margin elasticities to be smaller if switching from forest to cropland is costlier than switching between crops on already cleared land. The nested model's objective is to allow changes in a commodity's acreage to be decomposed into both margins:

$$\frac{d\ln \pi_i^c}{d\ln p_i^c} = \underbrace{\frac{d\ln \pi_i^{c|\mathcal{C}}}{d\ln p_i^c}}_{\text{commodity substitution: } \frac{\theta}{\lambda}(1-\pi_i^{c|\mathcal{C}})} + \underbrace{\frac{d\ln \pi_i^{\mathcal{C}}}{d\ln p_i^c}}_{\text{land conversion: } \frac{\theta}{\lambda}\pi_i^{c|\mathcal{C}}\times\lambda(1-\pi_i^{\mathcal{C}})}.$$

The first term is identical to the land-use change elasticity implied by Ricardian models and indicates how crop c acreage increases by stealing land shares from other crops. The second term tells us how crop c land use increases by stealing land shares from natural land. Total agricultural land responds to overall agricultural returns, measured as the price index $P_i^{\mathcal{C}}$, as follows:

$$\frac{d\ln \pi_i^{\mathcal{C}}}{d\ln P_i^{\mathcal{C}}} = \lambda(1-\pi_i^{\mathcal{C}}),$$

By separating the two margins, the model can reconcile relatively high substitution elasticities identified from variation within the agricultural nest (high $\frac{\theta}{\lambda}$) with relatively low land-use change elasticities identified from variation across nests (low λ).

5.2 Trade costs and the geography of market power

To quantify trade costs in our imperfectly competitive setting we first need to obtain the geographic distribution of markdowns. We combine the estimated elasticities of substitution within and across nests from the previous section with the land share data to obtain the price-elasticities of supply, using equation 5. Next, we combine these supply elasticities with data on the number of intermediaries to obtain the markdown μ_i^c for each location, using equation 8. The first and second panel of Figure 6 show the results for beef markets. Supply is less elastic in in-land regions on the agricultural frontier, resulting in wider markdowns on farm-gate prices.

Given our estimates for the spatial distribution of markdowns, we now proceed to quantify trade costs. My approach is to back out trade costs from price gaps between an origin location i and a destination location j, in the same spirit as Donaldson (2018). Implementing such a strategy requires data on origin producer prices and destination consumer prices *for the goods shipped from that specific origin*. I use county-level farm-gate prices from agricultural censuses as origin producer prices p_i^c, while origin-destination prices p_{ij}^c are from the agribusiness sourcing data.

Figure 6: Supply elasticities, markdowns, and model-implied trade costs (beef markets).

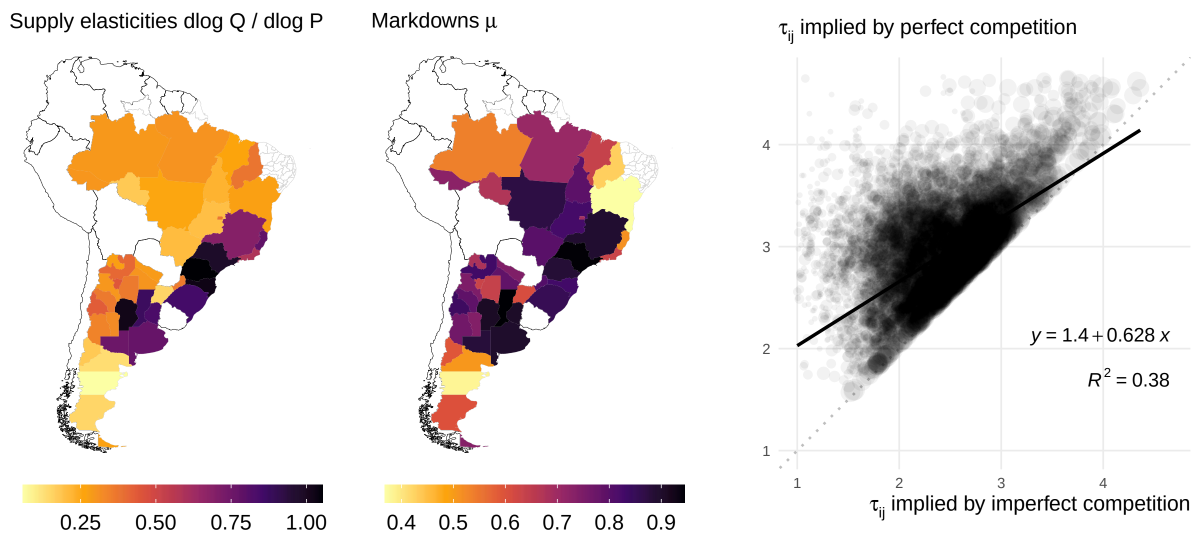

Notes: For the purposes of the maps in this figure, the supply elasticities and markdowns are computed at the county-level and then averaged at the state/province level for visualization purposes. The maps therefore show the average supply elasticities and markdowns for each Brazilian state and Argentine province. For the scatter plot, each bubble is an origin county-destination market pair. Bubble sizes are proportional to the number of agribusiness buyers at the origin (i.e., smaller bubbles reflect less competitive upstream markets).

In perfectly competitive settings such as Donaldson (2018), the price gap between origin i and destination j is sufficient to pin down bilateral trade costs,

$$\frac{p^c_{ij}}{p^c_i} = \tau^c_{ij}. \tag{14}$$

However, if intermediaries hold market power, the price gap reflects a combination of trade costs and markdowns (Atkin and Donaldson, 2015). Furthermore, if intermediaries are processing the commodity rather than trading it raw, we also need to account for value added (i.e., the marginal product of the commodity input to the processed good),

$$\frac{p^c_{ij}}{p^c_i} = \frac{\tau^c_{ij}}{\mu^c_i f'_c(q^c_i)}. \tag{15}$$

Without further assumptions, trade costs τ^c_{ij}, markdowns μ^c_i, and marginal products $f'_c(q^c_i)$ cannot be separately identified from price gaps. The conduct assumption is therefore an identifying restriction: it pins down μ^c_i as a function of supply elasticities and the number of firms. By evaluating 15 only on trade flows of raw commodities (e.g., unprocessed soybeans) we can impose $f'_c(q^c_i) = 1$ and trade costs are identified.[14] Notice different conduct assumptions will imply different trade costs through μ^c_i. Hence, conduct assumptions can be tested by how reasonable their implied trade costs are—for example, by how they compare to benchmark values from the gravity liter-

[14] Alternatively, we would need to estimate or calibrate $f_c(q^c_i)$ if dealing with trade flows of processed commodities, such as soybean oil.

ature (Anderson and Van Wincoop, 2003). This follows the same spirit as the "menu approach" from the empirical IO literature, where a menu of alternative conduct assumptions are tested by how reasonable their implied marginal costs are (Nevo, 2001).

Results. As equations 14-15 make clear, the markdown μ_i^c drives a wedge between the trade costs implied by assuming market power or not. Hence, if the imperfectly competitive model is the true one, incorrectly assuming perfect competition will inflate trade costs because the portion of the price gap generated by markdowns is attributed to trade costs. In our case, this upward bias is largest in inland regions because they have wider markdowns (i.e., lower values of μ_i^c, as shown in Figure 6). Thus, imposing perfect competition will inflate trade costs more in the uncompetitive (inland) regions than in the competitive (coastal) ones. The right panel of Figure 6 compares the trade costs implied by each conduct assumption and quantifies the size of this upward bias.

5.3 Demand elasticities

We start at the lowest level of the CES demand system and move up. First, the lower-level elasticity η_l is identified from expenditure variation across source origins i within a source nation n:

$$\ln\left(\frac{X_{ijt}^c}{X_{njt}^c}\right) = (1-\eta_l)\ln\left(p_{ijt}^c\right) + \lambda_{njt}^c + \varepsilon_{ijt}^c \quad \forall i \in n, \tag{16}$$

where X_{ijt}^c is destination j's expenditure on commodity c from county i and $X_{njt}^c = \sum_{i\in n} X_{ijt}^c$.[15] Because 16 is a demand equation, the issue of classic simultaneity bias arises when estimating via OLS. Therefore, I instrument for price with a supply shifter which I construct as $z_{ijt}^c \equiv s_{ij}^c \times w_{it}$, where s_{ij}^c is the share of origin i production that goes to destination j in a baseline year, and w_{it} is a local weather shock (measured as deviations from a historical average). The relevance condition of the instrument is straightforward: given a negative supply shock caused by adverse weather in origin i, the size of the supply shock effectively faced by destination j depends on how exposed it is to i as reflected by its historic trading relationship. The exclusion restriction is that the origin's destination shares s_{ij}^c are not predictive of changes in unobservable demand shocks—the shares are allowed to be correlated with the contemporaneous demand shock ε_{ijt}^c, but not with its change over time, $\Delta\varepsilon_{ijt}^c$. Therefore, the assumption allows for an origin i that historically exported most of its output to j to consistently experience large unobservable demand shocks from j, but not to experience systematic *changes* in such demand shocks over time.

The middle-level elasticity η_m is identified from expenditure variation across nations:

$$\ln\left(\frac{X_{njt}^c}{X_{jt}^c}\right) = (1-\eta_m)\ln\left(P_{njt}^c\right) + \lambda_{jt}^c + \varepsilon_{njt}^c, \tag{17}$$

[15] We treat $\lambda_{njt}^c \equiv -\ln\left(\sum_{i'\in n}(p_{i'jt}^c)^{1-\eta_l}\right)$ as a commodity-origin nation-destination-time fixed effect.

where X^c_{jt} is destination j's total expenditure on commodity c across all source nations.[16] Since 17 is a demand equation just like 16, 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, however at the origin nation level rather than the origin county level. Finally, the upper-level elasticity η_u is identified from expenditure variation across commodities,

$$\ln\left(\frac{X^c_{jt}}{X_{jt}}\right) = (1-\eta_u)\ln\left(P^c_{jt}\right) + \lambda_{jt} + \varepsilon^c_{jt}, \tag{18}$$

where X_{jt} is destination j's total expenditure on agricultural imports.[17] The required instrument is now a supply shifter varying across commodities, which I construct as supply shocks at the destination-commodity level as $z^c_{jt} = \sum_n m^c_{nj} w_{nt}$, where m^c_{nj} is the share of destination j imports coming from an origin nation n in a baseline year and w_{nt} is the origin nation-level weather shock.

Table 6: Demand substitution elasticities.

	Dependent variable: ln expenditure share					
	Lower level (across counties)		Middle level (across nations)		Upper level (across commodities)	
	OLS	IV	OLS	IV	OLS	IV
ln price	0.983***	−12.624***	−0.950***	−4.118***	−0.407***	−1.486
	(0.024)	(0.920)	(0.053)	(0.595)	(0.136)	(2.601)
Observations	168,112	168,112	83,512	83,512	1,899	1,899

Notes: Lower, middle, and upper specifications include origin nation-destination-year-commodity fixed effects, destination-year-commodity fixed effects, and destination-year fixed effects, respectively.

Results. Table 6 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. Intuitively, the results confirm lower substitutability across commodities than across commodity sources, and lower substitutability across nations than across counties. At the nation and commodity levels, the implied CES parameter values are $\eta_m = 5.12$ and $\eta_u = 2.49$. These results are of similar magnitude to those in the related literature: Costinot et al. (2016) estimate demand substitution elasticities across nations and commodities, finding $\eta_m = 5.40$ and $\eta_u = 2.82$.

[16] The term $\lambda^c_{jt} \equiv -\ln\left(\sum_{n'} a^c_{n'jt}(P^c_{n'jt})^{1-\eta_m}\right)$ and $\varepsilon^c_{njt} \equiv \ln\left(a^c_{njt}\right)$.

[17] To construct the price indices required for the upper-level estimation I use the residuals from the middle-level equation 17. That is, $P^c_{jt} \equiv \left(\sum_n \hat{a}^c_{njt}\left(P^c_{njt}\right)^{1-\eta_m}\right)^{\frac{1}{1-\eta_m}}$, where $\hat{a}^c_{njt} = \exp\left(\hat{\varepsilon}^c_{njt}\right)$. Furthermore, $\lambda_{jt} \equiv -\ln\left(\sum_{c'} a^{c'}_{jt}\left(P^{c'}_{jt}\right)^{1-\eta_u}\right)$ is treated as a destination-time fixed effect and $\hat{a}^c_{jt} = \exp\left(\varepsilon^c_{jt}\right)$.

6 Policy counterfactuals

6.1 Efficiency and distributional impacts of a downstream carbon tax

Motivated by the obstacles to implementing a first-best carbon tax directly on upstream farmers, I use my model to simulate a second-best tax levied at the downstream stage of the supply chain, i.e., on the agribusiness firms. The tax is second-best because of the limitations regulators have in tracing the emissions content of beef shipments once they arrive at the downstream stage—a well documented challenge in the South American beef industry due to its unique multi-stage and vertically unintegrated supply chain. Hence, I assume regulators use a national average footprint of CO_2e per tonne of finished beef, which is then multiplied by the social cost of carbon (assumed to be 30 USD/t CO_2e) to obtain an output tax per tonne of beef. This output tax is levied on all beef arriving at the downstream stage, independently of its upstream location of production. Therefore, the tax is spatially mistargeted because beef produced in high-emission locations is taxed at the same rate as beef produced in low-emission locations.

I implement the downstream tax in two ways. First, I consider the case where only a subset of trading partners implement the tax on their imports from South America. This is motivated by recent EU proposals to veto potential free trade agreements with the South American trade bloc, with the stated goal of reducing deforestation. Second, I compare this "incomplete regulation" case to a "complete regulation" scenario where all downstream consumer markets impose the tax, including the domestic market. Implementation details are in Appendix E.

Efficiency I: Emissions leakage across downstream consumer markets. The top panel of Figure 7 shows the equilibrium effects of an EU-only tax on South American beef. Since regulation is incomplete across consumer markets, the drop in shipments to the EU is offset in equilibrium by increased consumption in non-EU markets, including the domestic market. The tax's corrective potential is therefore substantially limited by this consumption "leakage" effect: over 80% of the emissions reductions attributed to the drop in EU consumption are offset by increased consumption elsewhere. The bottom panel of Figure 7 displays the case with complete regulation. In this case there is no re-routing of shipments, so emissions drop across all destination markets. It is important to note that the emissions changes in Figure 7 take into account substitution away from beef and into crops, i.e., domestic production leakage from beef (which is taxed) to crops (which are not). Specifically, the emission changes reported in the figure include the growth in crop-related emissions generated by substitution from beef into crops. Therefore, the fact that all bars in the bottom right panel of Figure 7 are negative indicate the drop in emissions from beef production more than offsets any increase in emissions from crop production.

Figure 7: Emissions abatement potential of downstream carbon taxes.

Incomplete regulation: unilateral EU tax on its South American imports.

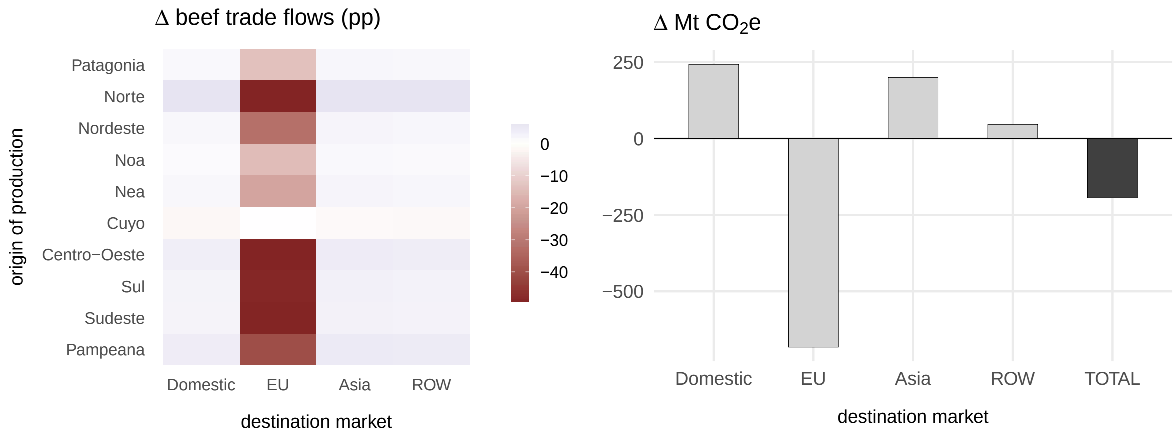

Complete regulation: downstream tax by all consumer markets.

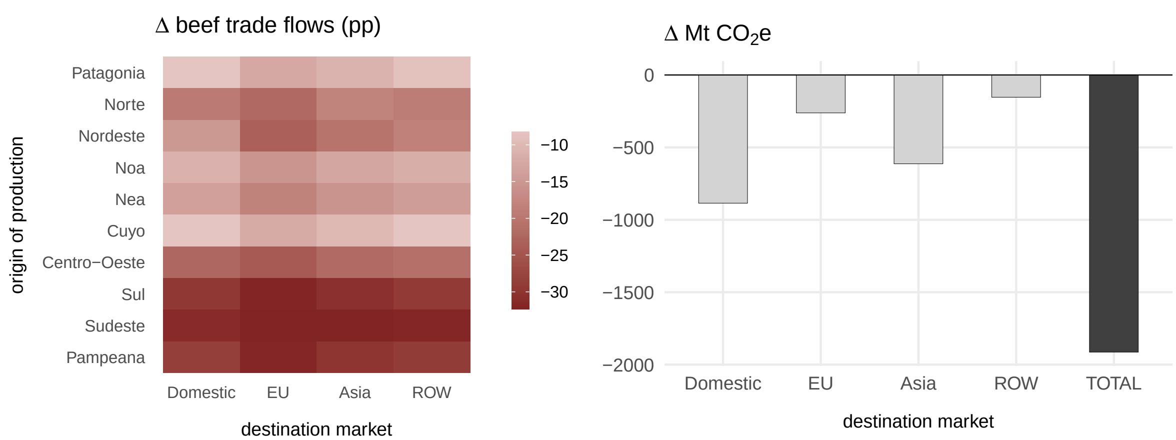

Notes: on the left, each matrix cell shows the change in output (reported as percentage point changes) shipped from an origin South American region to a destination consumer market (including the domestic market) as a result of the tax. For reference, regions are mapped in Appendix Figure 13. To go from output changes to emissions changes, the matrix is overlayed with the emission footprints of each origin. This is done for both beef and crops, in order to take into account emissions from the substitution of beef to crops. The resulting emissions from changes on all commodity flows (both beef and crops) are then aggregated up to the destination level, thus delivering the changes in emissions attributed to the changes in consumption of each destination (shown on the right in grey, with the total emissions impact across all destinations in black).

Efficiency II: Spatial mistargeting across upstream farmers. We now move beyond the international leakage effects to analyze how efficiently the downstream tax is domestically transmitted up the supply chain to farmers. From now on, we focus on the complete regulation case to isolate inefficiencies due to international leakage from inefficiencies due to the tax's lack of domestic targeting. The first thing to note is that the bottom trade matrix in Figure 7 shows larger changes in production for some upstream locations than others. The middle panel of Figure 8 maps these locations, showing that production drops least in regions on the agricultural frontier:

Central/Northern Brazil and the Argentine periphery. Because these regions have the highest emissions intensities, the downstream tax is spatially mistargeted upstream. This mistargeting occurs for two reasons. First, our empirical estimates indicate supply is less elastic in frontier regions, hence quantities drop less in response to the portion of the downstream tax that ends up getting passed through. Second, pass-through is lowest in the upstream locations most subject to agribusiness market power, which happen to be the frontier regions (left panel of Figure 8). Thus, while the spatial mistargeting result is independent of market structure assumptions (since it relies solely on the spatial pattern of supply elasticities), market power amplifies the mistargeting by reducing pass-through most to the highest-emissions locations.

Figure 8: Upstream heterogeneity in policy effectiveness and distributional incidence.

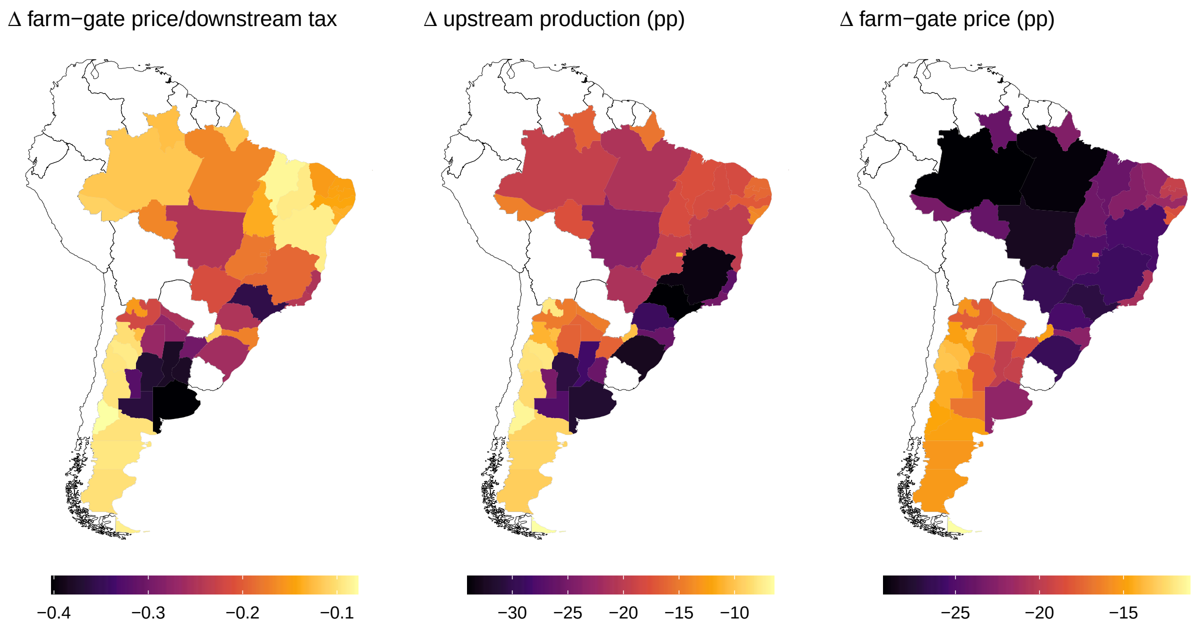

Notes: the maps show the difference between the baseline equilibrium and the downstream tax equilibrium. All values correspond to the beef cattle sector and are displayed at the state-level.

Redistribution: Supply-side regressivity across upstream farmers. Beyond the limited effectiveness of the downstream tax in terms of abatement potential, what are its distributional effects? The right panel of Figure 8 shows the impact on farm-gate prices expressed as percentage point declines, i.e., the policy's implied income tax on farmers. It is in the frontier agricultural regions of Northern Brazil—which are also the poorest regions—where farmer income is taxed at the highest rate. Although pass-through rates are lower in these regions, the implied income tax is higher because of lower farm-gate prices at baseline.[18] Since the implied income tax is highest for farmers in the poorest regions, 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.

[18]The pass-through rate is defined as the change in the farm-gate price relative to the downstream tax. The implied income tax is the defined as the change in the farm-gate price relative to the initial farm-gate price.

Role of market structure: Market power erodes the Pigouvian signal. We already discussed how pass-through rates are lower in frontier agricultural regions, resulting in smaller quantity responses and muted abatement relative to other regions. I now quantify by how much the downstream tax's corrective signal is eroded away by market power. To do so, I run my counterfactual simulation under two market structure assumptions: perfectly competitive and imperfectly competitive agribusiness intermediaries. Results are shown in Figure 9.

Market power results in pass-through rates being cut by more than than half in most locations, with substantial heterogeneity driven by regional variation in supply elasticities and intermediary concentration. The Brazilian Norte region—home to most of the Amazon biome and holder of the highest emissions intensity—has pass-through rates that are only a third of what they would be if markets were competitive. Thus, the Pigouvian signal of the tax is most eroded by market power precisely in the upstream locations with the highest social cost.

Figure 9: Role of market structure for pass-through from downstream firms to upstream farmers.

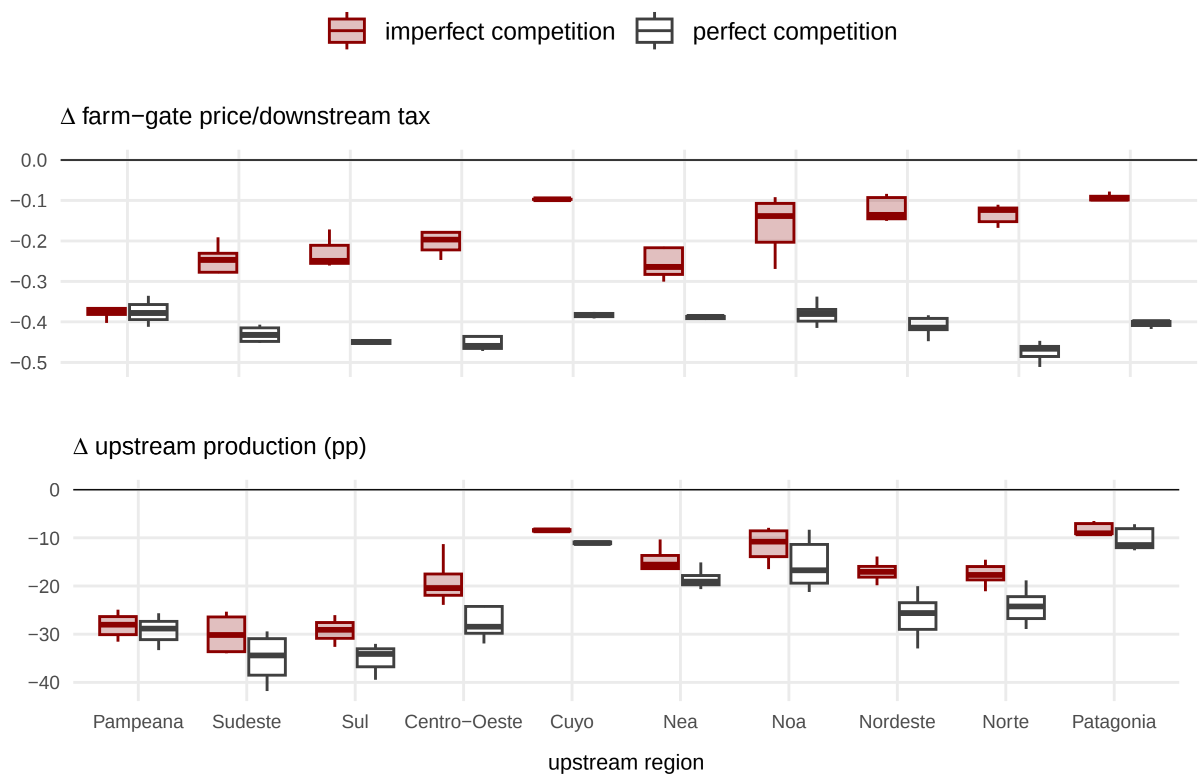

Notes: the figure shows the upstream effects of the downstream tax under each market structure assumption. All values correspond to the beef cattle sector. Box-and-whisker plots show the distribution of effects across locations within each region. For reference, regions are mapped in Appendix Figure 13.

It is important to stress the reason the mistargeting of the downstream tax is worsened by market power is because of the positive correlation between market power and emissions intensity, i.e., the most remote upstream locations are both uncompetitive and emissions-intense. The sign of this correlation is what determines whether pass-through is higher or lower in the most emissions-intense locations. If the correlation is positive, pass-through is lowered in the most emissions-intense locations and the mistargeting is amplified, while if the correlation is negative, the mistargeting is partially corrected. Thus, market power can affect the performance of market-based environmental policy in a qualitative sense, and not just quantitatively.

6.2 Market-based vs. command-and-control policy tools in agricultural supply chains

The funnel-like structure of agricultural supply chains, with millions of atomistic farmers upstream and a few large agribusiness firms downstream, has implications for which kind of policy tool is most effective at reducing emissions. The downstream tax from the previous section is an example of a market-based policy, as it aims to shift farmer incentives by changing the market prices they receive. Motivated by the implementation constraints regulators face, I simulated the policy as a flat tax per output on downstream agribusiness firms. This implies the tax fails to target the heterogeneity in carbon footprints across upstream farmers, making it blunt and second-best. In contrast to the market-based policy, a command-and-control policy in my model would take the form of a conservation zone in a high-carbon density area, where the quantity of deforestation is directly enforced at a specified monitoring cost. The appeal of the command-and-control policy is it is perfectly targeted, but only to a narrow geographic area due to its high enforcement cost. Hence, choosing between a market-based tool (downstream tax) and a command-and-control tool (conservation zone) involves a trade-off between targeting and enforcement costs. Through counterfactuals, this section shows this trade-off depends on the degree of heterogeneity in carbon footprints across upstream producers as well as market structure. If heterogeneity is high enough and market power is significant enough, the market-based tool becomes very mistargeted, allowing the command-and-control tool to achieve larger emissions reductions.

Implementation details. The counterfactual exercise involves the following four steps:

1. Simulate the market-based counterfactual: I use my downstream tax counterfactual from the previous section (with complete regulation) as my market-based counterfactual.

2. Simulate the command-and-control counterfactual: I simulate an equilibrium where direct enforcement is tightened for the counties in the top 10% of the carbon density distribution. In the model, this is implemented by changing the non-pecuniary returns, $A_i^{\mathcal{N}}$.

3. I compute the emissions abated by 1. and 2. relative to the baseline laissez-faire, denoting the emissions abated by each type of policy tool as A_{MB} and A_{CC}, respectively.

4. I repeat steps 1-3 under an alternative spatial distribution of carbon density which is mean-preserving. The baseline distribution is the one observed in the data and displayed in Figure 2, with a mean denoted μ and a standard deviation denoted σ. Hence, the alternative distribution also has mean μ, but an alternative standard deviation denoted $\tilde{\sigma}$. At the end of this step I obtain $A_{MB}(\tilde{\sigma})$ and $A_{CC}(\tilde{\sigma})$. I repeat this for various values of $\tilde{\sigma}$.

Targeting and robustness to market structure. Figure 10 shows how well-targeted each policy is at the baseline level of spatial heterogeneity σ. It plots the relationship between emissions intensity and changes in upstream production—a stronger negative correlation implies better targeting. For the market-based policy under perfect competition there is a positive relationship between emissions intensity and upstream production changes: regions with higher emissions-intensity experience smaller declines in production because their supply elasticities are lower. Under imperfect competition, the intercept of the relationship jumps because pass-through is weaker across all locations once market power is introduced. The slope also rises, because the most emissions-intense locations are the least competitive, hence their pass-through rates are lowered most.

For the command-and-control policy, there is a very clear negative relationship: production drops most in the most emissions-intense regions precisely because they are directly regulated. In equilibrium, production in non-regulated locations responds by increasing, although this is hardly visible from the figure because the total growth is distributed among many locations, while the command-and-control regulation is only imposed on a few locations. The most important takeaway is that the command-and-control results are nearly identical with or without market power. The command-and-control policy is robust to market structure precisely because it does not operate through the market mechanism.

Figure 10: Robustness to market structure (market-based vs. command-and-control policy).

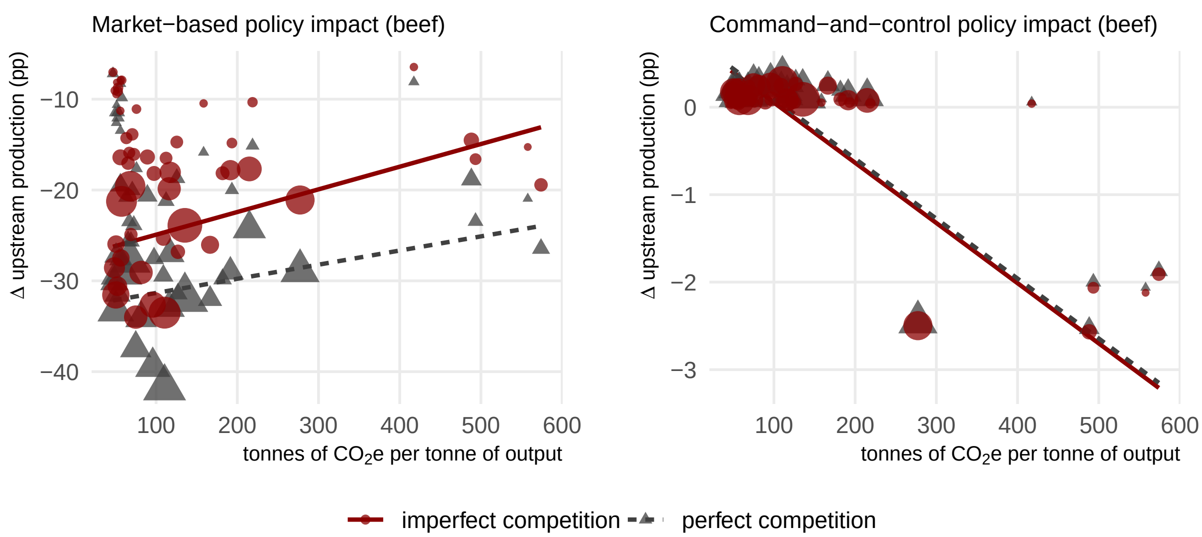

Notes: each marker represents a state, with sizes proportional to baseline production levels. Lines indicate the fit of a linear model to the markers, with observations weighted by baseline production levels. All results are from simulations that use the baseline level of spatial heterogeneity σ.

Choosing between market-based and command-and-control policy tools. For our concluding exercise we consider a regulator choosing between the two policy tools with the goal of maximizing abatement net of enforcement costs. Given a baseline level of heterogeneity σ, denote $A_{MB}(\sigma)$ and $A_{CC}(\sigma)$ as the emissions abated under each policy, C as the additional enforcement cost of the command-and-control tool relative to market-based tool, and SCC as the social cost of carbon. The regulator will choose the command-and-control tool if the social value of abated emissions net of enforcement costs exceeds the social value of abated emissions by the market-based tool,

$$SCC \times A_{CC}(\sigma) - C > SCC \times A_{MB}(\sigma) \iff C < SCC \times \overbrace{(A_{CC}(\sigma) - A_{MB}(\sigma))}^{\overline{A}(\sigma)}. \tag{19}$$

We define $\overline{A}(\sigma)$ as the additional emissions abated under the command-and-control tool relative to the market-based tool, and $\overline{C}(\sigma) = SCC \times \overline{A}(\sigma)$ as its dollar value. One interpretation of the right hand side of 19 is that enforcement costs cannot exceed a threshold $\overline{C}(\sigma)$ for the command-and-control tool to be chosen. As enforcement costs C decline, for example due to improvements in satellite image technology, command-and-control becomes more desirable. We can think of $\overline{A}(\sigma)$ as the abatement premium of the command-and-control tool over the market-based tool—the larger this premium is, the larger the enforcement cost the regulator is willing to tolerate to implement the command-and-control tool. Figure 11 plots $\overline{A}(\sigma)$ under the assumption of imperfect competition on the left panel, and compares it to perfect competition on the right panel.

Figure 11: Abatement performance of command-and-control policy over market-based policy.

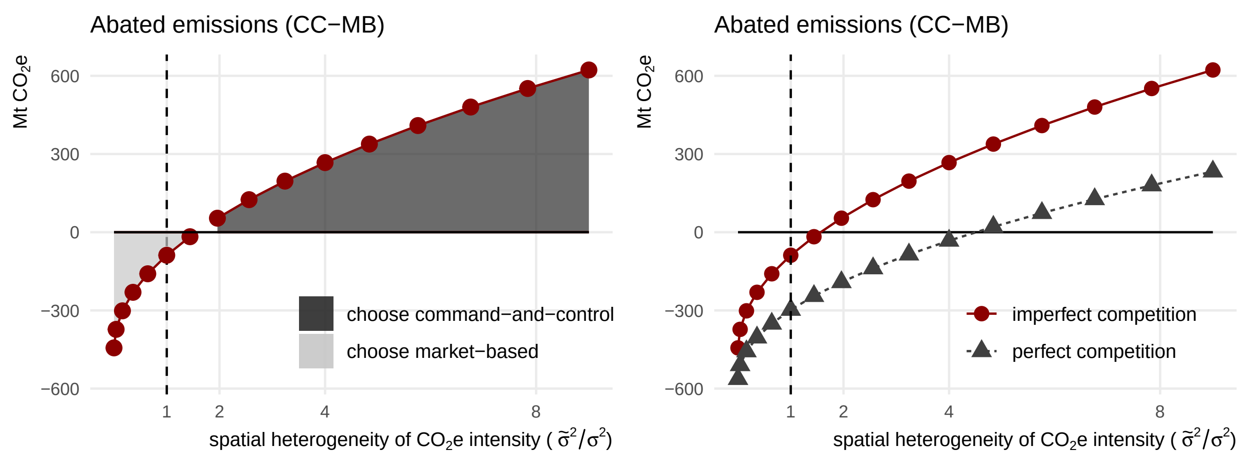

Notes: The vertical axis shows the command-and-control abatement premium—the difference between the emissions abated by the command-and-control tool relative to the market-based tool, for a given degree of heterogeneity $\tilde{\sigma}$. The horizontal axis shows the deviation of the counterfactual degree of heterogeneity $\tilde{\sigma}$ from the observed degree of heterogeneity σ, reported as the ratio $\tilde{\sigma}^2/\sigma^2$. Hence, $\tilde{\sigma}^2/\sigma^2 = 1$ corresponds to the degree of heterogeneity observed in the data. The figure on the left reports the results for the specification with imperfect competition, while the one on the right compares it to perfect competition.

The left panel of Figure 11 shows that at the baseline level of heterogeneity (a value of 1 on the figure's horizontal axis) we have $\overline{A}(\sigma) < 0$: the command-and-control tool delivers less abatement than the market-based tool. More importantly though, the crucial observation is that the command-and-control premium is increasing in the degree of spatial heterogeneity σ. The reason is that the more heterogeneity there is among upstream farmers, the more mistargeted the market-based tool is and the more valuable command-and-control becomes.

The right panel of Figure 11 shows that the command-and-control premium is larger under imperfect competition, and more so when spatial heterogeneity is at its highest. The reason is that the remote locations where markets are least competitive are also the ones with the highest emissions intensities. Hence, market-based policy, being subject to incomplete pass-through under imperfect competition, is especially incomplete in these high-emissions locations, while command-and-control policy remains robust, as we had seen in Figure 10. Thus, market-based policy can perform poorly when the market is has to operate through is distorted in ways beyond the environmental externality. In such cases, targeted command-and-control policies can be especially valuable precisely because they avoid the market mechanism.

To conclude, note that while the market-based tool is a price regulation, the command-and-control tool is a quantity regulation. Therefore, the findings are reminiscent of the classic trade-off between regulating prices versus quantities in settings where regulators face uncertainty about the emissions intensities of producers (Weitzman, 1974). In our setting, the trade-off arises because the regulator is uncertain about the carbon content of the commodity upon arriving at the downstream stage, due to the absence of a certification mechanism. Moreover, in our setting the heterogeneity across producers is manifested across geography through the carbon density of land. The intuition behind the results follows classic insights from public finance, as the degree of heterogeneity dictates which type of regulatory tool is optimal, while adding the insight that pre-existing distortions such as market power can tilt the trade-off towards the quantity regulation.

7 Discussion

The empirical model is parsimonious, with the main goal of showing how market-based environmental policy is transmitted along a supply chain, especially when the stage where it can be feasibly implemented differs from the one where the emissions are generated. Given I have abstracted away from modeling features that would distract from this goal, this section outlines the limitations of my analysis and the implications of re-introducing such features.

No reallocation across sectors, nor of factors other than land. The model has a single factor and a single sector: land and agriculture. I abstract from reallocation of land across sectors because it is not a quantitatively relevant margin for environmental purposes: agriculture uses 50% of the world's habitable land, while the urban and built-up areas where the manufacturing and service sectors reside represent less than 1%.[19] Modern-day deforestation is driven by agriculture, not by manufacturing nor services. I also abstract from reallocation of other factors such as capital and labor because land use is the first order determinant of agricultural emissions—how much land is deforested and which commodity is produced on it. In my South American setting, over 70% of agricultural emissions are attributed solely to land use change. As mentioned in the literature review in section 1, allowing for sectoral and factor reallocation would have been more important

[19] Source: OWID.

if this paper's research question would have been about the agricultural sector's adaptation to climate change, instead of agriculture's contribution to climate change. In this sense, I view this paper as complementary to the adaptation literature.

Robustness of results to market structure. It is important to stress that most of the qualitative results of the paper also hold with perfect competition, so the default assumption of market power is not crucial to deliver the main insights. First, the mistargeted aspect of the market-based policy also holds in the perfectly competitive case because it only relies on supply being less elastic in the more emissions-intense areas. Second, the same is true for the regressivity results, which only rely on principle of incidence: farmers with less elastic supply face a larger implied tax on their income. Third, the trade-off between market-based and command-and-control policy from section 6.2 is qualitatively the same with perfect competition: the command-and-control tool is preferred at high levels of spatial heterogeneity. The additional qualitative insight we obtain from introducing market power is that it amplifies the mistargeting of the market-based tool by reducing pass-through most to the emissions-intense regions, thus tilting the trade-off towards the command-and-control option. Throughout the paper, market power is therefore used as a counterfactual market structure within each counterfactual policy exercise, rather than as a necessary assumption to deliver specific outcomes in the baseline equilibrium.

Assumptions on firm conduct. Conditional on having an imperfectly competitive structure, there are many types of conduct one could choose from (e.g., Cournot, collusion, to name a few). The purpose of this paper is not to provide definitive evidence on the specific conduct of agribusiness firms, nor the welfare implications of their market power per se. Instead, the goal of introducing market power is to understand how it interferes in the transmission of market-based environmental policy relative to a perfectly competitive setting. For this reason, the intermediary part of the model is as simple as possible, following a standard Cournot specification that nests the perfectly competitive case in order to flexibly alternate between market structure assumptions when evaluating a given policy counterfactual. The baseline model does not have entry/exit nor firm heterogeneity, and while I discuss how to include these extensions in Appendix D.3, I ultimately do not include them in my baseline analysis because this would require a host of additional assumptions to estimate the new parameters that would need to be introduced: entry costs and elasticities of substitution across firms. Finally, and as already mentioned in section 4.2, I abstract away from market power of agribusiness firms over consumers because (i) data limitations on the consumer side prevent me from incorporating two-sided market power, and (ii) the environmentally-relevant decisions are made by farmers, so the key object of interest is the firm's upstream, not downstream, transmission of policy. Therefore, this paper is careful to avoid any statements on the welfare impact of market power per se, given such a statement would indeed depend on entry/exit, firm heterogeneity, and the extent of market power over consumers. Instead, this paper simply reports how market power changes the impact of environmental policy on key observables such as land use, farm-gate prices, and emissions. Importantly, the qualitative insights from introducing market power only rely on incomplete pass-through to upstream farmers, which in turn depends on the curvature of supply. Appendix D.2 provides a discussion on curvature and pass-through, adapting the standard analyses of Bulow and Pfleiderer (1983) and Weyl and Fabinger (2013) to the case of buyer market power. In short, to the extent that adding entry/exit, firm heterogeneity, and two-sided market power continues to deliver incomplete pass-through, the qualitative results of the paper would remain the same.

Static framework. I have chosen to use decadal census data because the priority for my research question is to have rich cross-sectional variation in variables such as land use and farm-gate prices, and linking these to agribusiness concentration. However, this comes at the cost of lacking the temporal requirements for estimating a fully dynamic model, since individual decision makers cannot be linked across time and the temporal frequency of the data is too low. Therefore, I opt for a static framework because of the long time lags in my decadal data, with the observed outcomes being interpreted through the lens of the model as a sequence of static equilibria separated by a substantial time lag. The static model is used as a first approximation for modeling long-run decisions, given dynamic considerations such as switching costs become less relevant the longer the temporal horizon is. The static framework's main implication for measurement is that my supply elasticities should be interpreted as long-run elasticities. Typically, the main reason for explicitly incorporating dynamics is for measurement purposes because the frequency of the data being used is annual (Scott, 2013; Araujo, Costa and Sant'Anna, 2020). At such a high frequency, switching costs need to be accounted for to correctly estimate land use change elasticities. While it is hard to predict how much my estimates would change if I had the data to estimate a fully dynamic model, my current estimates are at least broadly in line with long-run elasticities from the literature, as discussed in section 5.1. In short, abstracting from dynamics is more likely to be of quantitative than qualitative consequence for this paper, since the mechanisms driving the main insights are not of an inherently dynamic nature. This brings me to the static framework's main implication for policy analysis, which is that my counterfactual results should be viewed as long-run responses to policies that are permanently implemented, similar to the static analysis of Souza-Rodrigues (2019). Hence, I do not tackle policy questions that would require simulating a transition path, nor dynamic mechanisms such as commitment (Hsiao, 2021). I therefore view this paper as complementary to the work that does choose to tackle such issues.

8 Final remarks

This paper's main goal has been to show how a Pigouvian policy is transmitted along a supply chain, in particular when the stage where it can be feasibly implemented differs from where the externality is generated. Agricultural supply chains provide an ideal setting to study this mechanism because they are uniquely characterized by having the externality generated at the atomistic stage, where policy is more challenging to enforce than at the concentrated stage. Therefore, understanding how a corrective tax is transmitted from downstream agribusiness firms, where enforcement is feasible, to the upstream farmers who ultimately make the environmentally-relevant decisions is a first order issue for this setting.

The general interest takeaway of this paper is that market-based policies such as corrective taxes can be poorly targeted if they cannot be feasibly levied at the source of the externality. This lack of targeting is especially severe under two conditions. First, when there is wide heterogeneity in the intensity of the externality across its sources. Second, when pre-existing market distortions weaken the correlation between policy pass-through and externality intensity. Thus, market-based tools can perform poorly when the markets they operate through face distortions beyond the main externality they aim to correct. In such cases, targeted command-and-control tools can be robust to such distortions, precisely because they avoid the market mechanism.

The empirical application of this paper has been to a specific industry and externality: South American agriculture and its greenhouse gas emissions. While the industry is ideal for the research question and is important in and of itself due to its global environmental footprint, the overall message of this paper may resonate with other industries and other externalities as well. The garment industry has long struggled with the issue of labor exploitation in developing countries. The oil and gas industry is plagued with economic sanctions that aim to punish and deter the non-democratic actions of rogue states. While there is a common thread running through these examples, there are also key differences in terms of (i) the length and shape of the supply chain, (ii) the type of incentives that intend to be corrected (often related to an externality, but not always), and (iii) the pre-existing market distortions that Pigouvian tools need to work through to achieve their desired objective. Understanding how individual industries are characterized by specific iterations along these three dimensions can help yield general insights about the effectiveness and incidence of market-based corrective policies.

References

Alvarez, J.L.C. and Rossi-Hansberg, E. (2021). The economic geography of global warming. Tech. rep., National Bureau of Economic Research.

Anderson, J.E. and Van Wincoop, E. (2003). Gravity with gravitas: A solution to the border puzzle. *American Economic review* 93(1):170–192.

Araujo, R., Costa, F. and Sant'Anna, M. (2020). Efficient forestation in the Brazilian Amazon: Evidence from a dynamic model. *SocArXiv* .

Assunção, J., McMillan, R., Murphy, J. and Souza-Rodrigues, E. (2019). Optimal environmental targeting in the Amazon rainforest. Tech. rep., National Bureau of Economic Research.

Atkin, D. and Donaldson, D. (2015). Who's getting globalized? The size and implications of intra-national trade costs. Working paper 21439, National Bureau of Economic Research.

Barreto, P., Pereira, R., Jr., A.B. and Baima, S. (2017). Will meat-packing plants help halt deforestation in the Amazon? Tech. rep., Instituto Centro Da Vida.

Bento, A.M., Goulder, L.H., Jacobsen, M.R. and Von Haefen, R.H. (2009). Distributional and efficiency impacts of increased US gasoline taxes. *American Economic Review* 99(3):667–99.

Bergquist, L.F. and Dinerstein, M. (2020). Competition and entry in agricultural markets: Experimental evidence from Kenya. *American Economic Review* 110(12):3705–47.

Berry, S. and Schlenker, W. (2011). Technical Report for the ICCT: Empirical Evidence on Crop Yield Elasticities. *Weather* pp. 1–18.

Bisang, R., Santagelo, F., Anlló, G. and Campi, M. (2007). Mecanismos de formación de precios en los principales subcircuitos de la cadena de ganados y carnes vacunas en la Argentina. Tech. rep., Convenio CEPAL-IPCVA.

Bresnahan, T.F. (1989). Empirical studies of industries with market power. *Handbook of Industrial Organization* 2:1011–1057.

Buchanan, J.M. (1969). External diseconomies, corrective taxes, and market structure. *American Economic Review* 59(1):174–177.

Bulow, J.I. and Pfleiderer, P. (1983). A note on the effect of cost changes on prices. *The Journal of Political Economy* 91(1):182–185.

Burgos, M., Mattos, E. and Medina, A. (2014). *La soja en Argentina (1990-2013): Cambios en la cadena de valor y nueva articulación de los actores sociales*. Centro de Economía y Finanzas para el Desarrollo de la Argentina, CEFID AR.

Chatterjee, S. (2019). Market power and spatial competition in rural India. Working paper.

Clark, M.A., Domingo, N.G., Colgan, K., Thakrar, S.K., Tilman, D., Lynch, J., Azevedo, I.L. and Hill, J.D. (2020). Global food system emissions could preclude achieving the 1.5 and 2 degree C climate change targets. *Science* 370(6517):705–708.

Conte, B. (2020). Climate change and migration: the case of Africa. *Unpublished manuscript* .

Conte, B., Desmet, K., Nagy, D.K. and Rossi-Hansberg, E. (2021). Local sectoral specialization in a warming world. *Journal of Economic Geography* 21(4):493–530.

Copeland, B.R., Shapiro, J.S. and Taylor, M.S. (2022). Globalization and the Environment. In: *Handbook of International Economics*, vol. 5, pp. 61–146. Elsevier.

Costinot, A., Donaldson, D. and Smith, C. (2016). Evolving comparative advantage and the impact of climate change: Evidence from 1.7 million fields around the world. *The Journal of Political Economy* 124(1):205–248.

Crespi, J.M. and Saitone, T.L. (2018). Are Cattle Markets the Last Frontier? Vertical Coordination in Animal-Based Procurement Markets. *Annual Review of Resource Economics* 10:207–227.

Dhingra, S. and Tenreyro, S. (2020). The rise of agribusiness and the distributional consequences of policies on intermediated trade .

Diamond, R. (2016). The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000. *American Economic Review* 106(3):479–524.

Donaldson, D. (2018). Railroads of the Raj: Estimating the impact of transportation infrastructure. *American Economic Review* 108(4-5):899–934.

Eaton, J. and Kortum, S. (2002). Technology, geography, and trade. *Econometrica* 70(5):1741–1779.

Ehrl, P. (2017). Minimum comparable areas for the period 1872-2010: an aggregation of Brazilian municipalities. *Estudos Econômicos (São Paulo)* 47(1):215–229.

Fabra, N. and Reguant, M. (2014). Pass-through of emissions costs in electricity markets. *American Economic Review* 104(9):2872–99.

Farrokhi, F. and Lashkaripour, A. (2021). Can Trade Policy Mitigate Climate Change?

Fowlie, M., Reguant, M. and Ryan, S.P. (2016). Market-based emissions regulation and industry dynamics. *The Journal of Political Economy* 124(1).

Gerber, P., Steinfeld, H., Henderson, B., Mottet, A., Opio, C., Dijkman, J., Falcucci, A. and Tempio, G. (2013). *Tackling climate change through livestock – A global assessment of emissions and mitigation opportunities*. Rome: Food and Agriculture Organization of the United Nations (FAO).

Goldsmith-Pinkham, P., Sorkin, I. and Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review* 110(8):2586–2624.

Gonda, H., Feldkamp, C., Jaurena, G., Ferrer, J.M., Arroquy, J.I. and Garcia, F. (2017). *Low-emissions development of the beef cattle sector in Argentina - Reducing enteric methane for food security and livelihoods*. Rome: FAO and New Zealand Agricultural Greenhouse Gas Research Centre.

Herrero, M., Havlík, P., Valin, H., Notenbaert, A., Rufino, M.C., Thornton, P.K., Blümmel, M., Weiss, F., Grace, D. and Obersteiner, M. (2013). Biomass use, production, feed efficiencies, and greenhouse gas emissions from global livestock systems. *Proceedings of the National Academy of Sciences* 110(52):20888–20893.

Herrero, M., Thornton, P.K., Notenbaert, A.M., Wood, S., Msangi, S., Freeman, H., Bossio, D., Dixon, J., Peters, M., van de Steeg, J. et al. (2010). Smart investments in sustainable food production: revisiting mixed crop-livestock systems. *Science* 327(5967):822–825.

Hsiang, S. and Kopp, R.E. (2018). An economist's guide to climate change science. *Journal of Economic Perspectives* 32(4):3–32.

Hsiao, A. (2021). Coordination and Commitment in International Climate Action: Evidence from Palm Oil. *Working paper* .

IIASA/FAO (2012). *Global Agro-ecological Zones (GAEZ v. 3.0)*. IIASA, Laxenburg, Austria and FAO, Rome, Italy.

Inman, M. (2008). Carbon is forever. *Nature Climate Change* pp. 156–158.

IPCC (2014). Climate change 2014: synthesis report. *Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, RK Pachauri and LA Meyer (eds.)]. IPCC, Geneva, Switzerland* .

Klenert, D., Mattauch, L., Combet, E., Edenhofer, O., Hepburn, C., Rafaty, R. and Stern, N. (2018). Making carbon pricing work for citizens. *Nature Climate Change* 8(8):669–677.

Kortum, S. and Weisbach, D. (2017). The design of border adjustments for carbon prices. *National Tax Journal* 70(2):421.

Melitz, M.J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica* 71(6):1695–1725.

Menon, S., Denman, K.L., Brasseur, G., Chidthaisong, A., Ciais, P., Cox, P.M., Dickinson, R.E., Hauglustaine, D., Heinze, C., Holland, E. et al. (2007). Couplings between changes in the climate system and biogeochemistry. Tech. rep., Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States).

Montzka, S.A., Dlugokencky, E.J. and Butler, J.H. (2011). Non-CO2 greenhouse gases and climate change. *Nature* 476.

Nath, I.B. (2020). The food problem and the aggregate productivity consequences of climate change. Tech. rep., National Bureau of Economic Research.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica* 69(2):307–342.

Nordhaus, W.D. (2007). A review of the Stern review on the economics of climate change. *Journal of Economic Literature* 45(3):686–702.

OECD (2019). *Agricultural policy monitoring and evaluation 2019*. Paris: OECD Publishing.

Otsuka, K., Nakano, Y. and Takahashi, K. (2016). Contract farming in developed and developing countries. *Annual Review of Resource Economics* 8:353–376.

Pellegrina, H.S. (2019). Trade and the spatial organization of agriculture: Evidence from Brazil. *Working paper* .

Pereira, R., Rausch, L.L., Carrara, A. and Gibbs, H.K. (2020). Extensive production practices and incomplete implementation hinder Brazil's zero-deforestation cattle agreements in Para. *Tropical Conservation Science* 13:1940082920942014.

Pierri, J., Wesz Junior, V. et al. (2017). "Soyafication" in Argentina and Brazil (1980-2014): Influence of public policies, transnational companies, and the dependent economic structure. *Extensão Rural* 24(1):20–41.

Poore, J. and Nemecek, T. (2018). Reducing food's environmental impacts through producers and consumers. *Science* 360(6392):987–992.

Ripple, W.J., Smith, P., Haberl, H., Montzka, S.A., McAlpine, C. and Boucher, D.H. (2013). Ruminants, climate change and climate policy. *Nature Climate Change* 4(1):2.

Roberts, M.J. and Schlenker, W. (2013). Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate. *American Economic Review* 103(6):2265–95.

Rubens, M. (2019). Ownership consolidation, monopsony power and efficiency: Evidence from the Chinese tobacco industry. Working paper.

Scott, P.T. (2013). Dynamic discrete choice estimation of agricultural land use. *Working Paper* .

Searchinger, T., Waite, R., Hanson, C., Ranganathan, J., Dumas, P. and Matthews, E. (2019). Creating a sustainable food future: A menu of solutions to feed nearly 10 billion people by 2050. *World Resources Institute* .

Skidmore, M.E., Moffette, F., Rausch, L., Christie, M., Munger, J. and Gibbs, H.K. (2021). Cattle ranchers and deforestation in the Brazilian Amazon: Production, location, and policies. *Global Environmental Change* 68:102280.

Sotelo, S. (2020). Domestic trade frictions and agriculture. *The Journal of Political Economy* 128(7):2690–2738.

Souza-Rodrigues, E. (2019). Deforestation in the Amazon: A unified framework for estimation and policy analysis. *The Review of Economic Studies* 86(6):2713–2744.

Spawn, S. and Gibbs, H. (2020). Global Aboveground and Belowground Biomass Carbon Density Maps for the Year 2010. *ORNL DAAC* .

Springmann, M., Clark, M., Mason-D'Croz, D., Wiebe, K., Bodirsky, B.L., Lassaletta, L., de Vries, W., Vermeulen, S.J., Herrero, M., Carlson, K.M. et al. (2018). Options for keeping the food system within environmental limits. *Nature* 562(7728):519.

Stern, N. (2007). *The economics of climate change: the Stern review*. Cambridge University press.

Syverson, C. (2019). Macroeconomics and market power: Context, implications, and open questions. *Journal of Economic Perspectives* 33(3):23–43.

Tilman, D. and Clark, M. (2014). Global diets link environmental sustainability and human health. *Nature* 515(7528):518–522.

Train, K.E. (2009). *Discrete choice methods with simulation*. Cambridge University Press.

Vale, P., Gibbs, H., Vale, R., Christie, M., Florence, E., Munger, J. and Sabaini, D. (2019). The Expansion of Intensive Beef Farming to the Brazilian Amazon. *Global Environmental Change* 57:101922.

Weitzman, M.L. (1974). Prices vs. quantities. *The Review of Economic Studies* 41(4):477–491.

Weyl, E.G. and Fabinger, M. (2013). Pass-through as an economic tool: Principles of incidence under imperfect competition. *The Journal of Political Economy* 121(3):528–583.

Zavala, L. (2021). Unfair Trade? Market Power in Agricultural Value Chains. *Working paper* .