



R I D G E

RESEARCH INSTITUTE FOR
DEVELOPMENT, GROWTH
AND ECONOMICS

SINCE 2014

Farm to Firm: Clustering and Returns to Scale in Agricultural Supply Chains

James E. Sayre

RIDGE Working Paper 062

**This paper was presented at the
2025 RIDGE December Forum**

RIDGE Working Papers are not formal publications. They are preliminary versions of research presented at the RIDGE Forum and posted on the RIDGE website with the author's permission. They remain the intellectual property of the authors and may be submitted for publication in academic journals.

FARM TO FIRM: CLUSTERING AND RETURNS TO SCALE IN AGRICULTURAL SUPPLY CHAINS

James E. Sayre*
UC Davis

September 21, 2025

Abstract

Agricultural production in many countries is strikingly clustered. I ask why such concentration persists, focusing on Mexico, where high value agricultural production is geographically concentrated even as lower-yield subsistence farming persists throughout the country. Combining novel supply chain data with precise agronomic suitability and agricultural microdata, I show that this clustering is driven not by growing conditions, but instead by agricultural supply chains. I document that these supply chains require large, coordinated investments by both private firms and public actors. Firms in supply chains largely source agricultural products only from nearby regions, leaving regions further away unable to produce for higher-value domestic and international markets. I then develop a spatial general equilibrium model that explicitly models the coordinated investments necessary to establish agricultural supply chains. My approach not only rationalizes observed concentration patterns, but also has implications for understanding optimal agroindustrial policies and responses to mitigate climate shocks.

Keywords: agricultural supply chains, fixed costs, concentration

JEL Classification: Q17, F12, O19, R32

*This paper was the first chapter of my Ph.D. dissertation at UC Berkeley. I thank Matilde Bombardini, Thibault Fally, and Marco Gonzalez-Navarro for their guidance and mentorship from the inception of this project. I am grateful for the feedback of Ellen Bruno, Lorenzo Caliendo, Davin Chor, Abdoulaye Cisse, Benjamin Faber, Joel Ferguson, Simon Fuchs, Juan Carlos Hallak, Alain de Janvry, Andrés Rodríguez-Clare, Elisabeth Sadoulet, Leo Simon, Sofia Villas-Boas, and David Zilberman, as well as seminar participants at AAEA, APRG, Boston U, IATRC, Montana State, PacDEV, Peking, Stanford, TIGN, UC Berkeley, UC Davis, and the UEA Summer School. I also thank Carmen Estrades, Zoë Plakias, and Clayton Webb for their detailed discussions and comments on my work, as well as Melissa Dell, Sunghun Lim, Pablo Querubin, and Sebastian Sardon, who kindly provided me with data. This project was funded by the Giannini Foundation of Agricultural Economics. I appreciate the invaluable guidance and hospitality of Natalia Volkow and the INEGI microdata laboratory staff which allowed me to contribute to this paper under project number LM1325. The conclusions and opinions expressed therein are my own and do not constitute official statistics of INEGI nor that of the National System of Statistical and Geographical Information (SNIEG). All errors are my own. Contact: jsayre@ucdavis.edu.

1 Introduction

A key feature of the process of industrialization is the movement away from subsistence farming and into more specialized agricultural production. As agriculture has become more specialized in developed regions, the locations of agricultural consumption and production have become increasingly decoupled. As a result, modern agricultural production features clusters of intensive production located far away from downstream consumers. To provide two current examples, 88% of the strawberries produced in the United States are grown on the central California coast and 85% of Mexican avocados are produced in a small region that spans the states of Michoacán and Jalisco. In both cases, these patterns of concentration are relatively recent in historical perspective – in 1950, the same two states in Mexico accounted for only 15% of total domestic avocado production, and California accounted for 13% of total strawberry production¹.

Understanding the drivers of this process of agricultural specialization and concentration are crucial to think about the development of agrarian societies and regional welfare, where rural wages are strongly linked to the productivity of agriculture. For instance, in Mexico, farm laborers in more specialized areas earn wages up to 4 times larger than their counterparts elsewhere (see figure A.4). The mechanisms behind this process, however, remain unclear. On the one hand, examples of such productive agricultural clusters appear consistent with the stylized view of agriculture proposed by Krugman (1991). In that model, given intensive use of immobile land and constant returns to scale in agriculture, “the geographical distribution of [agricultural] production will be determined largely by the exogenous distribution of suitable land.” In other words, the unique climatic conditions of the California coast provide for the ideal location in terms of soil conditions and climate for growing strawberries. This source of Ricardian comparative advantage, coupled with decreases in trade costs, serves as the predominant explanation for these clusters in the literature (Costinot and Donaldson, 2012). This view is implicitly embedded in modern spatial models of agricultural production featuring Fréchet distributed land productivities, wherein land productivities and trade costs are the primary factors which can explain differences in the concentration of production.

However, these spatial models ignore the role that granular firms play in facilitating trade in agricultural products between producers and consumers in increasingly distant locations. This trade is enabled by firms in agricultural value chains (AVCs or AGVCs), which include participating farms, processors, packers, exporters, and retailers. These chains allow for innovations that enhance storability and phytosanitary quality, and thus enable agricultural goods to be shipped inter-regionally as well as internationally. These firms in agricultural value chains are highly concentrated relative to the number of farmers producing a given crop. They also facilitate the vast majority of international

¹See figures A.9 and A.11 for reference to how Mexican avocado production has shifted from being dispersed to highly clustered from 1950 to 2020, and Olver and Zilberman (2022) for data and reference about the development of strawberry production in the United States.

trade in agriculture. Further, they are often located in only a few regions of a country in those same productive clusters.

Therefore, to understand patterns of agrarian development and key drivers of rural welfare, it is important to disentangle the factors of agricultural suitability, trade costs, as well as granular agricultural value chains in influencing patterns of concentration in agricultural production. To make progress here, I assemble a novel dataset on granular agricultural supply chains in Mexico and their firm-level trade patterns. I combine this information with farm level microdata that provides information on agricultural production decisions, which I link to geographically precise measures of crop-specific suitability and weather shocks. Using this data, I provide a number of stylized facts that illuminate the dynamics at play between farm level decisions and the role of agricultural supply chains.

To start, I find that agricultural value chain presence explains concentration patterns at the subnational level in Mexico far more than agronomic suitability measures. Next, I document large differences between the number of farms and agroexporters in my setting, where there are greater than 7,000 times more farms than agroexporters. I show that the majority of exporting farms are in close proximity to agroexporters (within 25-75 kilometers), indicating that the distance to agricultural value chains plays a large role in determining who are able to export. Further, I show that in aggregate, Mexico's areas of intensive production of a given crop are not driven by agricultural suitability as much as in other contexts such as the United States. I hypothesize that the existence of barriers to complete Ricardian specialization of land operate through agglomeration forces inherent in agricultural value chains. If barriers to enter into more modern production and exporting are large, then areas unable to pay these costs (that are nonetheless suitable in terms of agroclimatic factors and trade costs) will be left out of these production opportunities. To support this hypothesis, I provide both qualitative surveys and firm microdata as quantitative evidence to suggest that the fixed costs of entry in agricultural supply chains are large and in part driven by the non-tariff barriers in agricultural exporting, often phytosanitary measures such as pest inspections.

Next, to examine the causal effects of agroexporters on their nearby regions, I consider an event study analysis where I examine the timing of the first entry of an agroexporter producing a given crop across a number of regions in Mexico. I find that in the years after firm entry, the surrounding municipality has a land share 2 percentage points higher dedicated to the crop the firm produces. Therefore, I conclude that these firms play significant roles in shaping the specialization of the areas they operate in.

Taking these features of agricultural production and supply chains into account, I develop a model of agricultural production and trade which incorporates granular agricultural supply chains, as well as captures the aggregate decisions made at the farm level. My model incorporates atomistic agrointermediaries with market power and large fixed costs of entry. Additionally, farmers decide

which crops to plant based upon input and output prices as well as the fundamental suitability of their land for that crop. The model features two separate production regimes that are consistent with Mexico's agricultural history, where agricultural production has shifted from domestic, dispersed production, to export-oriented and concentrated. One is autarkic, where farmers can only sell their output in local markets. In the second, farmers can sell to nearby agricultural supply chains, who can then sell the processed output domestically or internationally (subject to trade costs). I conclude that my model is better suited than others in the literature that do not explicitly include agointermediaries to match observed patterns of agricultural clustering.

Building on the framework I develop, I am able to use the structure of my model to tractably estimate the fixed costs of entry for agointermediaries producing a variety of crops in Mexico. I note that my estimated fixed costs are considerably larger for specialty crops such as avocados and sugar than cereals like maize. I compare these fixed costs with regional revenue in each crop, and find that most smaller production regions are not able to bear these costs, especially for specialty crops. These fixed costs are related to trade policy, as they are correlated with the number of non-tariff barriers to imports into the United States.

Ultimately, these fixed costs to entry of agricultural supply chains affect which regions are winners and losers from agricultural production and trade, relative to a world with constant returns to scale in the entire agricultural sector. Finally, developing a quantitative model allows me to examine a number of counterfactuals. I examine whether policies aimed at increasing supply chain entry will help expand access to export opportunities to farmers throughout Mexico. Such policies may include lowering the fixed costs of entry for intermediaries, either through government subsidies or trade agreements targeted at lowering the trade barriers imposed by phytosanitary regulation, or more general infrastructure investments to reduce trade costs. I also show that policies such as these may help mitigate the losses from climate change.

2 Related Literature

My paper touches upon several strands of literature in economics. In examining the features driving the agglomeration of agricultural production, my work relates back to the work of David Ricardo in understanding the patterns of comparative advantage as well as Marshall (1890) in understanding some of the features that generate external economies of scale. More recently, my paper relates to the literature in international trade studying economies of scale dating back to the seminal work of Krugman (1979). In follow up work, Krugman (1991) presents a core-periphery model where manufacturing firms cluster based on consumer demand. However, the agricultural sector is assumed to be constant returns to scale (CRS) with no transport costs, and in this way any potential clustering of agricultural production is ruled out. While Costinot and Donaldson (2012)

documents that patterns of national specialization in agriculture closely follow a pattern expected by the Ricardian theory of comparative advantage, I argue that these patterns are more muted at the subnational level.

I argue that these patterns can be explained by the role of global value chains (GVCs) in agricultural production. Recent work models optimal value chain structure more broadly (e.g. Antràs and Chor, 2013 & Antràs and De Gortari, 2020), as well as examined the relationship between AGVC production and structural transformation (Lim, 2023). Several papers demonstrate the ability for knowledge spillovers in agricultural supply chains to generate innovations “in the air” in the words of Marshall (1890) such as Reardon et al. (2019) and Zilberman et al. (2019). In particular, I argue that these patterns can be rationalized by the large fixed costs required to establish agricultural supply chains, which are driven by both internal and external economies of scale. Although I am not the first to observe that these intermediaries may have large fixed costs to establish², I am the first to estimate these internal as well as external economies of scale for agricultural supply chains. In this sense, my work touches on papers studying industries with potential national external economies (Grossman and Rossi-Hansberg, 2010 and Kucheryavyy et al., 2016), and work estimating external economies of scale in other (primarily manufacturing) sectors (Bartelme et al., 2019). If external economies of scale exist for certain agricultural supply chains, this literature establishes a theoretical basis for government policy to encourage additional clusters. Such policies may include subsidizing the costs of entry for agroexporters, improving infrastructure – particularly the infrastructure that facilitates international trade such as large ports (Ganapati et al., 2021), or negotiating for the removal of non tariff barriers to entry in agricultural exporting.

In studying the persistence of fixed investments, I touch upon the literature examining long run persistence in equilibrium outcomes stemming from differing initial conditions such as Bleakley and Lin (2012) and Bleakley and Lin (2015). Such work has argued for the importance of historical factors in path dependence of production, rather than innate natural factors (Lin and Rauch, 2020). In recent work, Allen and Donaldson (2020) develop a spatial model that features path dependence and characterize parameters required for the uniqueness of steady-state equilibria and provide bounds on steady-state welfare. Similarly, a number of recent papers provide existence and uniqueness

²For instance, Bergquist and Dinerstein (2020) examine agricultural markets in Kenya and observe that traders earn medium markups of 39% percent. Using experimental variation in a subsidy paid to induce entry of new traders (ranging from 49 to 148 US dollars), Bergquist and Dinerstein (2020) develop a model that rationalizes their low takeup of entry subsidies, which they conclude comes from relatively high fixed costs to entry into these markets. In examining the distribution of traders in their sample, they conclude that the largest ones have the biggest impacts on consumer welfare, and that effective competition policies must target the larger, much more profitable agents in agricultural value chains. In this paper, I focus my attention on the largest actors in agricultural value chains (such as exporters and processors). In doing so, I estimate the fixed costs for large agricultural supply chains, which would be more difficult to study experimentally than smaller traders.

for equilibria in similar models with increasing returns to scale (such as de Gortari, 2020 and Kucheryavyy et al., 2021, although the latter makes similar assumptions for the agricultural sector as Krugman, 1991).

On the methodological side, my paper stems from a number of recent papers studying trade and agriculture, particularly how trade affects local crop allocations, such as Bergquist et al. (2025), Donaldson (2018), Fajgelbaum and Redding (2022), and Sotelo (2020). In particular, my two-tier production function over crop varieties and input levels is similar to that featured in Farrokhi and Pellegrina (2023). This work also relates to a number of recent papers studying the market power of agriculture intermediaries in space. Chatterjee (2019) and Jung et al. (2021) study the market power of intermediaries in relation to the spatial distance to the farms they source from, respectively in Indiana and India, both concluding distance plays a major role in markdowns. Dhingra and Tenreyro (2020) examine the shift to agricultural-value-chain-oriented production in India, and find that farmers selling to such intermediaries saw larger declines in income than those not selling to such firms. Examining intermediaries in non-agricultural settings, Grant and Startz (2025) conclude that chains of intermediaries arise due to economies of scale in trade costs, and concludes that intermediary entry may benefit consumers through increased competition and gains from variety. Domínguez-Iino (2025) studies the consequences of potential emissions regulations on agriculture in a setting featuring agricultural value chains with market power in South America. Despite using a similar framework, the primary question is quite different and examines the interplay between agointermediaries, trade, and environmental outcomes such as deforestation. Although Domínguez-Iino (2025) considers an extension with entry costs of intermediaries, the main results ignore the possibility of entry and does not attempt to estimate them. In contrast, my model aims more explicitly at understanding the role of intermediaries in affecting the spatial distribution of agriculture, and attempts to estimate entry costs directly.

The most related work is by Zavala (2025), which documents the low share of revenue farmers in developing countries receive relative to agointermediaries. The paper uses a similar production structure to microfound market power in firms that source crops from smallholders, and assumes that farms can shift between intermediaries anywhere with a certain elasticity of substitution, subject to receiving a lower (trade cost adjusted) farm gate price. In my context and based on motivating statistics I present in 3.2, I note that farms by and large only supply nearby intermediaries, despite fairly low estimated costs of trade in my context. In contrast, I assume that while firms have market power in sourcing within their local region, they necessarily must have such market power in order to cover their large fixed costs of entry. These fixed costs simultaneously generate the need to pay markdowns on input costs as well as generate increasing returns to scale in agricultural value chain production. These fixed costs prohibit the entry of AGVC firms elsewhere, particularly in smaller regions. Therefore, relative to the work of Zavala (2025), my paper concludes that subsidizing the

entry of agricultural supply chains at the regional level is key to increasing farmer income, rather than the entrance of fair-trade constrained firms, since very few regions feature AGVCs for a given crop, as I note in section 3.1.

My paper also speaks to the literature examining the effect of distance on agricultural production decisions (such as Pellegrina, 2020), and seeks to disentangle the effects of agricultural supply chain presence from trade costs on production decisions in a quantitative model. For instance, Gáfaró and Pellegrina (2022) conclude that larger farms that are closer to urban areas are more productive, and more likely to sell in non-local markets (possibly through AGVCs, although this is not modeled explicitly), and develop a quantitative spatial model which can accommodate these two stylized facts. The authors find that removing geographic barriers to farm participation may raise output by 14%. However, although the authors examine similar stylized facts to my work, Gáfaró and Pellegrina (2022) model that these barriers arise from farm level fixed costs, and not from fixed costs in intermediary entry, as I argue. My estimates of the farm level agricultural production function in Section 6 suggest that fixed costs operate at the intermediary level, where I find that farms producing most crops display constant or decreasing returns to scale. In the same context of Mexico, Rivera-Padilla (2020) argues that lowering trade costs would lead farmers to shift their production from staple crops such as maize to cash crops. I argue instead that connections to AGVCs are crucial to sell cash crops in larger domestic and international markets, but that AGVCs choose to site in locations correlated with good market access, which feature low trade costs. In my model, I am able to distinguish the importance of agricultural supply chains and trade costs, which is difficult otherwise given the endogenous location decisions of supply chains.

Finally, I address the literature studying the low productivity in agriculture in developing countries such as Rao (1993), Tombe (2015), Gollin et al. (2014). In particular, the recent work of Adamopoulos and Restuccia (2021) concludes that cross country differences in agricultural productivities are not driven by differences in land suitabilities across countries, as measured by the FAO GAEZ project. Instead, they conclude that aggregate yield gaps between rich and poor countries would be reversed if areas could easily change crops to ones that have the highest value yields, and that spatial reallocation of crop production would raise agricultural productivity particularly in low income countries. In this work, I highlight some of the barriers to such a theoretically optimal spatial reallocation of crop production. A large body of work also examines the effects of participation in AGVCs on smallholders in developing countries, including Bellemare and Bloem (2018), de Janvry and Sadoulet (2019), Lence (2016), Meemken and Bellemare (2020), and Nuhu et al. (2021). In studying the effects that firms (often multinational firms) have on their suppliers, my work relates to work such as Alfaro-Ureña et al. (2021).

3 Data and Descriptive Statistics

3.1 Data Sources

A necessary input to understanding the dynamics between farms and agointermediary firms (henceforth, firms) is knowledge of the structure of agricultural supply chains, which can often vary by crop (i.e., many agroexporting facilities process only one commodity, and certain inputs are suitable only for one crop). This can be difficult with standard firm and agricultural censuses. For instance, most agricultural censuses detail input decisions, crop choice, and sales behavior, but do not contain any information on the intermediaries who buy farmers' output. Likewise, agricultural processing and packing firms can be found in firm or economic censuses, but information on the sourcing of inputs, input prices, and even which crop a firm is exporting can be lacking, depending on the detail of the firm censuses³. I integrate several sources of information to document how these farms and firms interact.

Setting

In this paper, I focus on the context of North American agriculture, with a focus on Mexico. Mexico is an ideal setting for studying the distributional effects of agricultural supply chains. Highly productive, export-oriented agriculture (much of it bound north towards the United States and Canada, and distributed by North American multinationals⁴) is prevalent in some regions, while subsistence agriculture is the main feature of others. Markets between Mexico and the United States are well integrated – the lack of tariff barriers due to the North American Free Trade Agreement (NAFTA)/United States-Mexico-Canada Agreement (USMCA), close proximity, and good infrastructure mean that fresh produce can be readily supplied across the border. Many of the large agroexporters in Mexico are headquartered in the United States (most commonly in the border states of Arizona, California, and Texas), and many rely on Mexico's differences in seasonality to supply the American market when it cannot produce during the winter. As a share of total employment within the entire agricultural food system, which includes agriculture, non-food agriculture, and food processing and distribution, agriculture comprises roughly 50% of employment. For comparison, this places Mexico as an "emerging" country, where agriculture comprises a smaller share of employment than most developing countries where agriculture directly involves roughly 90% of employment, but a much higher share than the United States, where roughly 15% of total employment is in the agriculture-food sector (Ambikapathi et al., 2022).

³An added difficulty is that at times industry classifications such as NAICS do not provide sufficient detail to determine which firms act as packing and exporting firms. For instance, among the lists that SENASICA reports for avocado, coffee, mango, and wheat packing firms with certification to export, packing firms can be found in five separate NAICS 3 (also 2) digit codes: 115, 311, 431, 461, and 493.

⁴See appendix figure A.13 for the total exports reported by firms with and without foreign direct investment from the United States; many Mexican crop agroexporters are wholly owned by US stakeholders.

Information on agroexporters and other firms in the agricultural value chain

I obtain information on agroexporters for a wide range of crops from digitized trade records from SENASICA phytosanitary certificates. These trade records provide firm-to-firm trade, in this case export, data for Mexico from 2011 to the present. The trade records report this information at the crop level.

To provide information on agroexporters prior to 2011, I combine various editions of the Mexican Agricultural Exporter Directory put together by the Secretary of Agriculture and Rural Development (*Secretaría de Agricultura y Desarrollo Rural*, SADER). The exporter directory details the crops a firm produces, the location by state, the destinations a firm exports to, as well as the certifications it has obtained. The directory is provided biennially, allowing for inference of when a given firm started exporting a given crop. To validate this source of data, I exploit phytosanitary rules for certain crops which require farms and downstream plants wishing to sell their output to register with a national agency in charge of phytosanitary regulation (also known as a national plant protection organization). The Mexican National Service of Food Sanitation, Safety, and Quality (*Servicio Nacional de Sanidad, Inocuidad y Calidad Agroalimentaria*, SENASICA) regularly publishes the currently approved Mexican farms and processing facilities for several crops for which phytosanitary regulations require certification. I find that for matching years, both of these data sources have almost a complete match, confirming the accuracy of the Mexican Agricultural Exporter Directory.

I can then link the plants of agroexporters to the Mexican firm censuses (*Censos Económicos*) based on name and address information⁵. This allows me to obtain standard firm measures such as headcount, capital stock, and revenue.⁶

Agricultural microdata

To understand detailed patterns of crop production across space, I rely on microdata from 1991, 2007, and 2022 rounds of the Mexican Agricultural Census (*Censo Agrícola*) provided by the National Institute of Statistics and Geography (INEGI, or *Instituto Nacional de Estadística y Geografía*). The Agricultural Census reports standard measures of inputs and outputs, as well as sales characteristics of farms, such as whether a farm sells to supermarkets, packing houses, or exports *directly*. Crucially, the Agricultural Census reports the universe of farm units within Mexico. The agricultural census does not provide information on the prices farmers receive for their output nor their input costs. To link farms with information about their input costs, I rely on the national agricultural survey (*Encuesta Nacional Agropecuaria*, ENA), which provides

⁵I find the plants for each agroexporter by searching for the name of the firm in the National Statistical Directory of Economic Units (*Directorio Estadística Nacional de Unidades Económicas*, or DENUE). This results in a match rate of 44% for firms in the Mexican Agricultural Exporter Directory and almost 20% for firms in the trade records.

⁶The firm censuses are conducted in all localities with a population of more than 2500 every 5 years (i.e. 1989, 1994, 1999, 2004, 2009, 2014, 2019, 2024)

information regarding crop choice, input decisions, prices, production, and sales at the farm level⁷.

Land use

I supplement this fine-grained production survey with aggregate yearly agricultural production data from the Service of Agrofood and Fisheries Information (*Servicio de Información Agroalimentaria y Pesquera*, or SIAP) which supplies data on crop production at the municipality level from 2003 to 2024⁸.

Agricultural suitability

I combine distinct sources of information regarding agronomic and ecological crop suitabilities. The more commonly used source is the FAO Global Agro-Ecological Zones (GAEZ) database (Fischer et al., 2021), which provides productivity information on 29 separate crops at a 5 arc-minute worldwide grid (30 arc seconds for selected crops). To supplement this source of data for crops that the GAEZ database does not cover (most speciality crops such as avocados and mangoes), I use the FAO EcoCrop database (Hijmans et al., 2001) to provide information on crop growing requirements as well as gridded products for temperature, rainfall, soil conditions, and other variables. I merge these suitability variables with the most precise level of geographic aggregation available for the 1991, 2007, and 2022 Agricultural Censuses, which is called the *área de control* or *manzana rural*.⁹ I describe the creation of these suitability indices in an [online appendix \(section B\)](#).

US-Mexico data

Given the prominence of US-Mexico trade in agriculture, I model the rest of world as based upon the United States. To do so, I use information from the US Department of Agriculture (USDA) National Agricultural Statistics Service on county level crop acreage data. I produce a correspondence table to link these data, SIAP data, and agricultural production found in the Mexican Agricultural Census. I also obtain information on the main border and sea ports for agricultural trade between the US and Mexico from the US Census Bureau and the Mexican Secretariat of Communications and Transport. I obtain information on urban and rural counties and their populations from the USDA Economic Research Service (ERS).

⁷The ENA is available in 2012, 2014, 2017 and 2019, and is representative at the crop by state level.

⁸In addition, SIAP provides the same agricultural production information at the state level going back to 1980.

⁹Áreas de control are statistical units that aim to group together geographically clustered farms or farm units with similar characteristics in terms of agricultural production, land ownership status (e.g. private or communitarian, known as *ejidal*) and socioeconomic characteristics. The median *área de control* comprises 10 farm units, with a minimum of one farm unit, and a maximum of 700 farm units (in the state of Oaxaca, which is known for having small plots of land). In terms of land size, the median *área de control* is less than 129 hectares, far smaller than the median municipality with a land size of 23,178 hectares.

3.2 Descriptive statistics

Using these sources of data, I produce a number of statistics regarding the characteristics of AGVCs and farms, effects of AGVCs on upstream farms nearby to them, the placement of agricultural value chains in space, and the difficulties farmers encounter in trying to export their crops – either the regulations required to export or the role of distance in determining who can sell to AGVCs.

Observation 1. Farms wishing to export need to sell to concentrated agroexport sector

The first observation I make is the large imbalance between the number of farms and the number of agroexporters in Mexico, so that any farm wishing to export their output to international markets will almost certainly need to do so through an agroexporting firm. I report the summary statistics for farms and firms in my sample in Table 1 below.

In the 2007 Agricultural Census, there are 3,656,611 farm units, and in 2022, 4,373,956 farm units. Of these, in 2007 only 73,334 farm units contract directly with commercial buyers, but having a explicit contract is not necessary to sell to agroexporters, so there are many more farms selling than is captured in the data¹⁰. Farms that fulfill the necessary prerequisites are allowed to export directly, but this is relatively uncommon – in 2007, I observe that only 3,112 farms export directly to foreign markets. In 2022, there are even fewer farms exporting directly – only 1,841 farms.

In 2007, I obtain a list of 649 agroexporting firms operating in Mexico, with 754 separate plants in total based upon the agroexporter directory. Examining the share of municipality-crop observations with operating exporting firms (there are 2439 municipalities and 157 crops in my sample, so 382,923 municipality-crop pairs in total), there are less than 0.12% of observations with an operating packer, and only 7.7% of metropolitan area-crop pairs¹¹ contain an operating packer. This implies that even aggregating up to larger geographic unit, farmers producing most crops will not have relatively easy access to an agroexporter.

Observation 2. Increasing returns to scale in agricultural exporting may come from both internal and external economies of scale

The reasons that drive agricultural supply chains to cluster in certain locations may be explained by the role of fixed costs in establishing supply chains, wherein only the largest, most productive areas will be able to bear those costs. These fixed costs come from both internal and external economies of scale. The observation that firms may cluster around specific locations to take advantage of external economies of scale dates back to Marshall (1890), who argued that these may arise from 1) access to a common labor market and infrastructure, 2) lower transportation and transaction costs along the supply chain, and 3) economies from knowledge spillovers. Although

¹⁰In some crops such as avocados, contracting is relatively uncommon, and fruit is purchased directly in spot markets, whereby afterwards it is immediately harvested and shipped.

¹¹I define my 75 metropolitan areas in section A.6.

Table 1: Summary statistics for farms and agroexporters

Farms	Number (#) of farms	# of farms contracting directly w/ buyers	Number (#) of farms exporting directly
	3,656,611	73,334	3,112
Agroexporting Firms	# of agroexporting plants	# of agroexporting firms	# of municipalities w/ an agroexporter
	754	649	122
Share of farms with access to agroexporter	In same mun.	In neighbor mun.	In metropolitan area
	0.12	0.77	7.66

Data on farms comes the 2007 Agricultural Census at the crop-farm level. Agroexporter data from SADER, restricting to firms that were operating in 2007.


there is extensive work documenting the importance of these factors (e.g. Ganapati et al., 2021, Gáfaró and Pellegrina, 2022, and Zilberman et al., 2019), quantifying the importance of external economies of scale in agriculture remains elusive (see Bartelme et al. for estimates of other sectors). Yet, there is suggestive evidence that external economies of scale in agricultural production may be large. For instance, in order to export avocados to the US, a region in Mexico must petition both the USDA and SENASICA for export approval. This entails the creation of a USDA inspection station in the area, which surveys every farm in a municipality on a semi-annual basis, and strict sourcing requirements to ensure exports are only sourced from USDA approved regions. In these productive regions, local extension agents provide information to farmers on how to meet phytosanitary requirements, and local growers associations have been established to ensure uniform compliance with these rules, which need to be met by all farmers.

However, the fixed costs of entry for a region also come from the internal economies of scale needed to establish an agroexporting firm. This may come from a number of sources, such as the need to establish initial relationships with farms, construct a manufacturing facilities, and purchase machinery. Fortunately, these costs are observable in firm balance sheets. Although these internal economies of scale may well be smaller than external economies of scale (some of which are hard to measure), in Figure 1, I present suggestive evidence towards the size of internal fixed costs by crop, wherein I sum up the “active fixed costs” of all agroexporters producing a crop and average this across regions. Although the size of the internal total fixed costs will depend on the number of firms producing a crop as well as the size of each firms’ total fixed costs, this provides suggestive evidence particularly that external fixed costs are larger for certain crops, such as avocados, coffee, and oranges.

Observation 3. Areas nearby agroexporters have higher shares of land allocated to the crop exported by those firms

I draw from information on downstream firm locations from the Mexican Agroexporters Directory to understand the association between the presence of a downstream firm within a region and

Figure 1: Estimates of total internal EoS for agroexporters by crop



Fixed cost estimates are based on each category of fixed costs reported in the Economic Census, and converted to gross estimates, for a given crop across all agroexporters in a given region. Agroexporter and crop information from trade records are matched to DENU, fixed cost data comes from the 2019 Economic Census. Pest control spending is yearly and from SENASICA.

the share of land allocated to the production of the crop that downstream firm sells. In table 1, I alternatively regress the share of arable land within a municipality or an individual farm unit dedicated to a given crop on the presence of a downstream firm dedicated to the export of that same crop.

In column (1), I find a large association between the two – municipalities where an exporting firm is present have a higher land share of that crop by 6 percentage points at the intensive margin. These associations between packer presence and land share diminish as the location of the packing plant gets further away – packer presence in a neighboring municipality is associated with a 3.1 percentage point (p.p.) higher land share, and packer presence in the broader metropolitan commuting zone is associated with a 2.9 percentage point higher land share. Comparing this against a measure of rainfed agricultural suitability combining FAO GAEZ and Ecocrop measures, I find in column (1) that a movement from zero suitability to full suitability in a given crop at the municipality level is associated with a 0.7p.p. higher land share of that crop, a small but significant association. Comparing the relative influence of all of these factors in a horserace, the

association between various measures of packer/exporter presence and land share is stronger than that of suitability and land share, and the (combined) coefficient(s) on packer presence is (are) approximately nine (eighteen) times larger than that of suitability. In column (2), I re-examine this pattern for crops with larger than median phytosanitary regulations (crops with more than 7 lines of regulations in the US Code of Federal Regulations for imports) and find that both packer presence and suitability are relatively more important in explaining land share at the intensive margin, although their relative coefficients are similar. When I examine the effect of packer presence and suitability at the intensive and extensive margin in column (3), where I fill in for zero land shares to obtain observations for every crop-municipality-year triple, I find that the importance of packer presence decreases by 35% relative to column (1), although the coefficient on land suitability decreases by 78% vis-a-vis column (1). I take this to indicate that these suitability measures are not that informative towards explaining production patterns at the extensive margin at the municipality/regional level.

Table 1: Relationship between crop land share and presence of downstream plant specializing in it

Dependent variable: Crop share of land	Municipality level			Farm level	
	(1) (Intensive margin)	(2) (Intensive margin)	(3) (Extensive+Int. margins)	(4) (Intensive margin)	(5) (Ext.+Int. margins)
Mun. has crop exporter	0.0606 (0.00264)	0.0985 (0.00494)	0.0393 (0.00142)	0.167 (0.00201)	0.137 (0.000819)
Neighbor mun. has exporter	0.0310 (0.00128)	0.0495 (0.00213)	0.0140 (0.000374)	0.0858 (0.00158)	-0.00331 (0.000502)
Metro. area has exp.	0.0287 (0.000832)	0.00798 (0.000931)	0.00397 (0.0000978)	0.0524 (0.000779)	0.144 (0.000242)
Rainfed suit. $\in [0, 1]$	0.00676 (0.00118)	0.0127 (0.00171)	0.00149 (0.0000880)	-0.00362 (0.000630)	0.847 (0.000466)
Mun.-Year* FE	X	X	X	X	X
Crop-Year* FE	X	X	X	X	X
Data source	SIAP	SIAP	SIAP	CA-2007	CA-2007
Sample	Full, 151 crops	Hi. SPS crops	Full	Full	Top 17 crops
N	416,809	109,008	5,030,288	4,792,134	58,367,154

Data in columns (1),(2), & (3) is regional land share from SIAP at the crop-year-municipality level from 2003-2021 for 151 crops. Data in columns (4) and (5) is farm land share from the 2007 Agricultural Census at the crop-farm level, and municipality-year and crop-year fixed effects here are simply municipality and crop fixed effects.

Packers/exporter data from SADER, suitabilities from FAO GAEZ/EcoCrop. Column (2) restricts to high phytosanitary strictness crops; crops w/ > 7 lines in the US Code of Federal Regulations. Columns (3) and (5) fill in zeros for crops with zero production. Column (5) includes the 17 largest crops in terms of hectares planted due to computer memory restrictions. Standard errors in columns (4) and (5) clustered at the municipality-crop level.

In columns (4) and (5), I show that a similar relationship holds even at the farm level, where I regress the share of land a farm unit allocates to a given crop on indicators for the presence of downstream firms within a region. I find a largely similar pattern in column (4) to that observed

at the regional level, that having a packer within one's municipality is associated with a roughly 16p.p. higher share of the land of a farm unit dedicated to that crop at the intensive margin. This effect dissipates the further a agroexporter is away from a farm, roughly a 3 times smaller association if the agroexporter is simply located in the same region versus the same municipality. In these regressions, the coefficient on land suitability is even smaller than at the regional level at the intensive margin. The coefficient bears an unexpected negative sign, although the coefficient is small and close to zero, despite being statistically significant. In column (5), I fill in for zeros in farm level land shares dedicated to a given crop to obtain observations for every crop-farm pair to examine the effects of firm presence on both the intensive and extensive margins. Given the number of farms (more than 5 million), in order to run this on INEGI's servers without hitting computer memory constraints, I am forced to limit the sample to the top 17 crops in terms of country level production value, which results in more than 58 million observations. Similar to the regressions at the regional level, I find that the coefficient on agroexporter presence decreases by 18% when considering the extensive margin in column (5) relative to column (4), however the coefficient on the presence of an agroexporter present at the metropolitan area level triples relative to column (4). Interestingly, the coefficient on rainfed suitability in this specification is now larger than the combined effects of agroexporter presence at various levels of aggregation: for instance, on average, a fully suitable farm for a given crop would allocate 85p.p. of its land to that crop, whereas a farm with an agroexporter in both its municipality, neighboring municipality, and metropolitan area would on average allocate 28p.p. of its land to that crop. This suggests that at the extensive margin at the farm level, suitability is roughly three times more important than packer presence. Therefore, suitability is very important in the farmer's decision of what to plant and these measures likely reflect some observable signal and not simply noise, but when considering how much to plant – suitability is almost irrelevant, and the presence of an agroexporter whom the farmer can sell to (and subsequently receive potentially higher prices from) is far more important. This may explain the results at the municipality level – although suitability is important to farmers, the lack of an effect of suitability at their intensive margin may imply that when aggregating crop choice to the municipality level, the observed effect of suitability is much smaller.

In appendix table A.3, I also show that there is a strong relationship between packer presence and the probability that the farm grows only a single crop. This indicates that the presence of value chains contributes to areas of monoculture agriculture in Mexico.

Observation 4. Production clusters do not systematically display higher yields or land suitabilities

A frequent argument for the existence of clusters of specialized crop production is that there are only a number of locations that are ideally suited for the production of a given crop. I develop a procedure to identify geographically distinct production clusters using only information on hectares planted of each crop across 104 crops in Mexico (and the United States), which I

describe in more detail in Appendix Section A.7. Potential suitabilities for crops are unequally distributed across and within countries, and often specialized agricultural production areas are located in high suitability regions. However, limiting to areas that currently produce a given crop (i.e., conditioning on non-zero production), areas with more intensive production of that crop often do not feature higher agricultural suitabilities (or average yields) than areas that produce that crop with less intensity¹².

I can then further sort all production clusters based on their relative size in acreage planted within each country, and then compare their percentile in their size distribution to their percentile in their crop specific yield, suitability, or farm TFP distribution¹³. I report these results in Figure 1, where I plot on the x-axis the position of each crop-cluster in their size distribution, and on the y-axis the position of each crop-cluster in the variable of interest. Across all variables, I find a modest relationship between cluster size and various measures of productivity, with no regression exceeding an R^2 of above 50 percent, with many crop clusters located at the top-left (small, productive clusters) and bottom-right quadrants (large, unproductive clusters) of the figures. I find a relatively weak association between relative cluster size and their relative suitability within Mexico, with a R^2 of below 6 percent. In contrast, the association between relative cluster size and their relative suitability within the United States is much stronger, indicating frictions to Ricardian specialization in Mexico not present in the US. Likewise, I find a relatively weak association between cluster size and yields in Mexico, but a slightly stronger one between cluster size and average farm level TFP.


There are several reasons for why suitability may not be a strong influence for cluster location in Mexico as well as some concerns that can easily be ruled out. For instance, I find that the concern that some crops are not located in the right tail of the suitability (or yield) distribution because other, more suitable, crops are crowding them out is largely not true. For each of the largest clusters, I run an algorithm to determine whether all of the production that happens in a cluster could be moved to an area that would be more suitable (or have higher yields, etc.) *without* displacing production of more suitable crops in that region nor displacing another large cluster of another crop. I find that in 96 of the 104 total crops, I am able to perform such a rearrangement of production so that the cluster moves to the most suitable region. This results in an average 49.5 percentage point increase in the empirical suitability distribution for the largest crop cluster¹⁴.

¹²This finding is consistent with the findings of Costinot and Donaldson (2012), which finds high correlations between suitabilities and production at the country level, but worse predictions for land use at the grid level, selecting the highest comparative advantage crop for each grid cell.

¹³I describe the estimation of farm level TFP in section 6. Farm level TFP is averaged across all crop-producing farms in a municipality for a given crop.

¹⁴Upon examining the shifts in production necessary to achieve these results for each crop, 74 come from moving to a farther away region, and the remaining 22 come from reallocating production to more productive municipalities within the same cluster.

Figure 1: Correlation between size of cluster and cluster suitability, yields, and TFPs



An important explanation is that top agricultural clusters are driven by the locations of firms in the agricultural chain, who choose where to enter based on a number of factors. Although firms such as agroprocessors take factors such as crop suitability into account, they also consider the market access of the potential location, such as the effective distance their products must travel to markets, the degree to which an area will be able to supply their input, or competitive forces such as the presence of other firms in the value chain¹⁵.

¹⁵This finding is echoed by other work such as Wantchekon and Stanig (2015), who find poorer districts in Africa

Another explanation is that agricultural indices such as those produced by the FAO GAEZ project fail to fully account for adaptive investments undertaken by farms, particularly those in productive regions, to improve the suitability of the lands they cultivate. I attempt to mitigate this concern by using the suitability indices from GAEZ that correspond to high levels of intermediate input usage, the use of CO_2 fertilization, and irrigation, as well as construct FAO EcoCrop indices that relax some of the constraints on rain and soil suitability, assuming that these can be mitigated through irrigation and input use respectively. These indices corrected for high input usage generate the main results in Table 1. Perhaps farmers can adapt to poor measured suitability in ways beyond those modeled by these indices¹⁶. That said, these indices are reasonably informative, especially given their importance in explaining crop planting decisions at the extensive margin at the farm level as shown in Column (5) of Table 1. Furthermore, both the indices I generate from EcoCrop and the expert generated ones from GAEZ yield similar results for the crops where both indices are available.

Observation 5. Many farmers report that export rules are significant barriers to external sales Despite the observed patterns in suitability for export clusters, nonetheless some areas produce far more than other areas. One potential explanation for these patterns is that these clusters arise due to the necessity of downstream plants in agricultural value chain production. These firms are necessary to fulfill certain phytosanitary norms such as specific treatments, provide knowledge of other technical barriers to farmers, and to invest in marketing of domestic agricultural products abroad.

As evidence of this, in self-reported survey data, many farmers report difficulties to exporting, either in terms of a lack of knowledge of export rules or difficulties in fulfilling phytosanitary (SPS) or technical barriers to trade (TBT). In appendix figure A.6, I show the share of farms reporting difficulties in exporting due to SPS and TBT guidelines. On average, 11.1% of farmers report difficulties in selling their crops due to a lack of knowledge of export rules. However, this hides a considerable amount of heterogeneity in the strictness of phytosanitary and TBT guidelines across crops – only 5% of farm units growing soy report having difficulties exporting due to SPS and TBT guidelines, whereas 42% of farm units producing raspberries report the same. Of course, these survey measures may underestimate the degree to which TBT and SPS regulations prevent farmers from exporting – if farms cannot find buyers for their crops, they would likely report that as their main difficulty in selling abroad, even if the sparsity of buyers is driven partially by these

have higher soil suitabilities, but worse transport infrastructure, and suggest that investments are lower in places with more abundant suitability.

¹⁶Another potential issue with the GAEZ indices is that they suffer from measurement error, where in some countries potential yields are lower than actually reported yields (Rattalino Edreira et al., 2021). However, potentially more precise measures, such as those provided by the Global Yield Gap Atlas, lack the geographical and crop coverage that GAEZ and EcoCrop provide.

regulations.

In selected interviews with farmers, many farmers growing high-value crops such as avocados have identified these restrictions as crucial to their difficulties, particularly those from the United States. One avocado producer in the Valle de Atlixco, in the state of Puebla, when asked why they only supplied the regional market, responded: “As producers, we consider ourselves as not having the conditions to be able to export, even less to the United States, our neighbors to the north, who are very, very strict about the products that enter their country”.

Such downstream firms require sufficient scale to export abroad, and therefore locate only in regions with adequate capability to supply them. Likewise, the presence of firms that are exporting abroad or sell to domestic centers of demand induces farms to produce crops that the firms demand, since the farms will receive higher prices than what they would receive locally.

Observation 6. The majority of exporting farms are in close proximity to packing facilities

In Figure 2, using data on the quantity of avocado exports that matches exporting farms to the packinghouses they sell to, I display the distances between exporting farms and the packing firms they sell to. In appendix figures A.2 and A.3, I also report the total amount of avocado exports by distance between avocado farms and packing firms, as well as distances between mango farms and the closest mango exporting firms. For avocado exporting farms, almost 40% of farms are located nearby (within 50 kilometers) of packing plants, and almost 80% are located within 100km of such plants¹⁷. Mangos are even more spatially concentrated – more than 80% of farms are located within 50km of packing plants, and 94% are located within 100km of such plants. Only the largest farms are able to overcome these distances and continue exporting, leaving out smallholder farmers from receiving higher export prices in regions that do not feature exporting firms.

3.3 Modeling implications

Observation 1. suggests that only a few regions in Mexico feature agroexporters, and therefore the model should rationalize that only some areas will have these agricultural supply chains for a given crop. Observation 2. suggests that the fixed costs of entry for agricultural supply chains may be quite large and thus the model should feature these costs. It also shows that external economies of scale matter and the model should feature a way to estimate both internal and external economies of scale. This observation also suggests that firms will need to have market power in input markets in order to cover their fixed costs, assuming that output markets are perfectly competitive. Observation 3. suggests that areas in which agricultural exporters are present have much higher shares of land allocated to production of that crop, and highlights the need to have a production structure which can justify the importance of suitability at the farm level, and its

¹⁷Similar results are obtained weighting by the quantity of fruit supplied by avocado farms to exporting firms.




Figure 2: Distances between avocado exporting farms and the exporting firm they sell to

diminished importance at the regional level. Observation 4. suggests that suitability measures alone will not capture clustering patterns, and thus the model needs to incorporate more than fixed productivity measures to capture the determinants of agricultural land use. Observation 5. suggests that input costs may be much higher for farms that are exporting, but also that exporting firms are critical in lessening barriers to exporting, as many farms cannot overcome these barriers alone. Finally, observation 6 suggests that the size of the area of influence of agroexporters should be sufficiently small, as most exporting farms are located within 100 kilometers of the nearest agroexporting firm.

4 Event study: do agroexporters site or sort into areas of clustered crop production?

When examining the role that agointermediaries and, more specifically, agroexporting firms play in regional development, a question that arises is whether such firms enable farmers to realize higher prices and spur shifts toward export-oriented crops or whether those regions were already productive and oriented towards specialized crops, spurring the entry of such firms. To disentangle these factors, I analyze an event study design where I look of the effect of examining the first year of exports of a given crop in a certain region. The plausible exogeneity that I rely upon is the uncertainty of receiving approval to export to the United States, where more than 80% of Mexican agricultural exports are destined for. In order to export to the United States, agroexport

firms need to pass both Mexican and US approval procedures. Unfortunately, I know little about the stringency or uncertainty of Mexican export approval. However, on the US side, according to import permit information obtained from the USDA Animal and Plant Health Inspection Service (APHIS) through a Freedom of Information Act Request, the vast majority of requests are processed relatively quickly, with a median response time of 20 days. That said, some requests take much longer, with the 90% percentile response time of 3 months, and a maximum response time of around 45 months.

In the trade records, I exploit the information provided at the firm-crop level to determine which municipalities first begin exporting a given crop in a given year. In the majority of my estimations, since the trade records begin in 2011, I ignore “first observations” of crop exports in the years of 2011 and 2012, since these may not truly be the years a region begins exports of a crop, but simply reflect the first appearances in the trade records, and simply treat these as always treated. In certain estimations, I limit observations to agroexporters who continuously export the crop in the region between 2013 and 2023, or restrict to agroexporters exporting values of a crop sufficiently large.

To begin, I start by a difference-in-difference design using (three-way) fixed effects for crops, years, and municipalities. I examine separately the effect of having the first processor open on the outcomes listed above. I implement my results examining the average treatment effect on the treated of first starting exports of a given crop in a municipality using the estimator developed by Borusyak et al. (2021). In Figures 3(a) - 3(d), I present my results using this estimator in an event study design, where I examine the effect of the year of the first exports of a given crop by an agroexporter in a given crop-municipality pair on the subsequent municipality level performance of that crop. I do so while including municipality, crop-year, and other fixed effects to avoid the need to control for time-invariant characteristics of municipalities which may have contributed to their crop mix and export behavior, or global demand shocks for a given crop. To avoid SUTVA violations where unmodeled entrances of exports of other crops may downward bias effects on the modeled crops, I include fixed effects to account for the presence of other crop exports in a municipality.

My outcomes of interest are the share of arable land in the municipality dedicated to that crop, the hectares planted of that crop (in logs), and the value and quantity produced of the crop (also in logs). In these results for all variables, I observe generally insignificant, perhaps slightly pretrends in outcomes prior to plant openings, and significant increases in regional level production afterwards, with increasing effects over time. This suggests that these processors have large and increasing effects that do not fade over time, with reasonable lags in their effects which may allow for farm supply to increasingly come online, which may be expected for perennials which may take a number of years to bear fruit. Interpreting the results in figure 3(a), I conclude that ten years after the first exports of a given crop in a region, the share of the municipality’s agricultural acreage

planted with that crop has increased 2 p.p..




Figure 3: Event study results of first agroexporter entry

Of course, although in these regressions I control for municipality, crop-year, and other crop exporter fixed effects, some time varying characteristics that influence both agroexport presence and agricultural production may be driving these effects, as my results do not entirely rule out an increasing trend over time. To examine potential evidence against this, I run a placebo test where I falsely assign treatment of having a packer open in different crops and years, holding fixed the set of treatment municipalities (on the basis that these municipalities with processors may differ systematically from those who do not). In doing, so I sample from the same underlying distribution of years (in terms of mean, standard deviation, and minimum and maximum). I plot these results in appendix figure A.19. Although in these figures I see evidence of an upward increasing trend after the date of entrance of the first (false) packer (perhaps because I restrict to the same sample with agroexporter entry at some point), none of these effects are statistically significant.

5 Theoretical framework

5.1 Setup and Preferences

Based on these observations and evidence, I develop a model which can rationalize the patterns of clustering by incorporating specific features of agricultural value chains. The model features internal and external trade without gravity in the agricultural sector, similar to several recent papers in the international trade literature (Bergquist et al. (2025) and Sotelo (2020)). On the production side, I model a joint decision where farmers decide which crop varieties to plant and whether they will produce those crops with technology that enables them to be export eligible or not¹⁸. At the same time, crop-specific processing and packing firms (“downstream plants”) set prices paid to farmers¹⁹, taking into account the localized aggregate elasticity of crop supply. This joint decision affects the possibilities for agricultural trade: farmers can sell directly to domestic markets, or to downstream plants, but cannot sell directly to international markets. Downstream plants can sell the crop output in domestic and international markets, but their inputs (crops) must fulfill higher standards such as stricter phytosanitary requirements.

Model Setup

The model aims to capture the aggregate implications of farm-to-firm linkages in agricultural production and exporting. The model consists of several regions, where domestic regions within Mexico are indexed as $i \in \mathcal{M} \equiv \{1, \dots, I\}$ ²⁰, $i \in \mathcal{US} \equiv \{I + 1, \dots, I + J\}$ refers to regions in the United States, and $i = F$ refers to the rest of the world (ROW), modeled as one region for simplicity. I let $\mathcal{W} \equiv \mathcal{M} \cup \mathcal{US} \cup \{F\}$ refer to the full set of regions (i.e. the world). Within domestic regions $\mathcal{M}$, each region i is classified as either urban (u) or rural (r), with the set of urban regions (in Mexico) given by $\mathcal{U}_{\mathcal{M}}$ and rural regions $\mathcal{R}_{\mathcal{M}}$, with $\mathcal{M} = \mathcal{U}_{\mathcal{M}} \cup \mathcal{R}_{\mathcal{M}}$. Each rural region is associated with an urban region, and the associated urban region is given by the function I denote by $U(i)$ for $i \in \mathcal{R}$. I also define the set of rural regions associated with an urban region $i \in \mathcal{U}$ to be $R(i) : \{n : U(n) = i\}$.

There are several goods in the model. The first goods are crops, each of which is assumed to be homogenous and thus is not differentiated by origin, indexed by $k \in \mathcal{K} \equiv \{1, \dots, K\}$. Each crop k has a variety denoted by $e \in \{0, 1\}$, which refers to the export-eligibility of the variety. I denote $e = 0$ as non-export eligible varieties (crops grown with low phytosanitary standards) and $e = 1$ as export eligible varieties (with high phytosanitary standards). Export eligible varieties, however, do

¹⁸The model features a concave production possibilities frontier (PPF) across crops and technology levels, similar to Farrokhi and Pellegrina (2023).

¹⁹Similar to recent work (Zavala, 2025), I assume that these firms have market power, but in contrast, I assume that these firms use markdowns in order to cover the high fixed costs of entering export markets, which I estimate and validate using firm microdata.

²⁰Regions may also be denoted n as destination regions.

not have to be sold internationally, and will be demanded in domestic markets. Next, there is also a general traded manufacturing good M . Finally, there are also different intermediate inputs used in the production of each good k , $\{x_k\}_{k \in \mathcal{K}}$, all of which are imported from abroad. The set of all goods is referred to as $\mathcal{G} \equiv \mathcal{K} \cup \{M\} \cup \{x_k\}_{k \in \mathcal{K}}$, and an individual good $g \in \mathcal{G}$.

In rural regions $i \in \mathcal{R}$, there are two types of agents. First, the representative consumer owns land H_i and labor L_i , and rents out both factors and purchases consumption goods in local markets. The consumer supplies both land and labor inelastically, the former of which consists of a continuum of plots indexed by $\ell \in \Omega_i$, with $H_i = \int_{\Omega_i} \ell$. Next, the representative farmer hires factors in local markets, and sells their output either domestically or to downstream plants.

In urban regions $i \in \mathcal{U}$, there are three types of agents. The representative consumer here owns only labor L_i and no land, which she rents out at the prevailing wage w_i inelastically and purchases consumption goods in local markets. The representative downstream plant buys crop inputs from farmers, hires factors in local markets, and sells their output in domestic and international markets. Finally, there are firms in the traded manufacturing sector which hire labor in local markets.

Trade

In the model, trade is costly, which is modeled using iceberg trade costs $\tau_{ni,g} \geq 1$. The no-arbitrage rule applies, which means that if a good is cheaper to source from elsewhere rather than locally (inclusive of trade costs), consumers will source that good, lowering domestic prices to the trade cost inclusive price of the import. Formally, let p_{ng} be the price of good $g \in \mathcal{G}$ in home region n . Then, for any potential sourcing region i , the price of g in region n will be $p_{ng} \leq \tau_{ni,g} \times p_{ig}$, and this will hold with equality if there is trade of good g from region i to region n . Finally, Mexico is presumed to be “small” in the sense that it is a price taker, that is, it imports and exports at exogenously given goods prices p_{wg} .

Preferences

The representative consumer consumes two main aggregates in the upper tier: agricultural goods and manufactured goods. They have preferences over these aggregates which are Cobb-Douglas with expenditure shares ζ for agricultural consumption and $(1 - \zeta)$ manufacturing consumption, with $\zeta \in [0, 1]$. In the middle tier for agriculture, consumers consume agriculture as a constant elasticity (CES) aggregate of individual crops k . Consumers have an elasticity of substitution across crops k given by $\sigma_A > 0$ with preference shifters $a_k > 0$. In the middle tier for manufacturing, in a similar fashion consumers have CES preferences over varieties of manufacturing goods. These goods are differentiated by origin, and consumers have an elasticity of substitution across varieties given by $\sigma_M > 0$. In the lower tier, for each crop k consumers choose whether to purchase an individual crop k in a local market (with low phytosanitary standards $e = 0$) or from part of the agricultural value chain such as a supermarket ($e = 1$). Consumers have an elasticity of substitution $\sigma_e > 0$ across these types of stores, with preference shifters $a_{ke} > 0$.

5.2 Production

Agriculture I assume that agricultural production can only occur in rural regions $i \in \mathcal{R}$. Suppose that the production function of a crop k with export eligibility e in region i in plot ℓ is Cobb-Douglas:

$$q_{ike}(\ell) = \overline{q_{ike}} \times l_{ike}(\ell)^{\alpha_{ke}} \times x_{ik}(\ell)^{\beta_{ke}} \times [h_{ike}(\ell)z_{ike}(\ell)]^{\gamma_{ke}}. \quad (1)$$

Here, l_{ike} is the amount of labor employed in plot ℓ and $\alpha_{ke} \in [0, 1]$ is the cost share of labor. x_{ik} is an intermediate input used in plot ℓ , $\beta_{ke} \in [0, 1]$ is the cost share of the intermediate input. h_{ike} is the amount of land in plot ℓ allocated to k , and $\gamma_{ke} \equiv (1 - \alpha_{ke} - \beta_{ke}) \in [0, 1]$ is the corresponding cost share of land. I define $\overline{q_{ike}}$ to be $\equiv (\alpha_{ke})^{-\alpha_{ke}}(\beta_{ke})^{-\beta_{ke}}(\gamma_{ke})^{-\gamma_{ke}}$. The remaining term, $z_{ike}(\ell)$, is a productivity shifter whose realizations are drawn randomly from the following Fréchet distribution:

$$\Pr [\widehat{z}_i(\ell) < \widehat{z}_i] = \exp \left[-\phi \sum_{k \in \mathcal{K}} A_{ik}^{\vartheta} \left(\sum_{e \in \mathcal{E}} z_{ike}^{-\theta} \right)^{\vartheta/\theta} \right]. \quad (2)$$

²¹ If land is unsuitable for crop k , the scale parameter A_{ik} is set to zero²². I describe the roles of θ and ϑ and the implications of this productivity distribution in the following section.

Implications

Concave production possibilities frontier

The assumptions of productivity draws in agriculture introduced in equation 2 yield an aggregate two-tier constant elasticity of transformation (CET) production function at the regional level i . This aggregate production function can be characterized by a concave production possibility frontier across different types of export-eligible varieties and different crops. In the upper tier of this aggregate production function, a land use constraint governs the choice across crops. In the second tier, a lower tier land use constraint governs the choice to grow $e = 0$ or $e = 1$ type crops.

In the upper tier, in region i , land H_i can be used to produce crop k at efficiency units $\widetilde{H}_{ik}$ such that the following constraint holds:

$$\sum_{k \in \mathcal{K}} \left(\widetilde{H}_{ik} \right)^{\frac{\vartheta}{\vartheta-1}} = H_i^{\frac{\vartheta}{\vartheta-1}}. \quad (3)$$

Here, $\vartheta > 0$ governs the elasticity of transformation between different varieties of crop production. If ϑ is large, production will be skewed towards the crop the region is most productive in, and we will observe larger specialization. If ϑ is smaller, specialization will tend to be incomplete and more regions will produce a larger amount of crops.

²¹ Here, as standard, I define $\phi \equiv \left[\Gamma \left(1 - \frac{1}{\vartheta} \right) \right]^{-\vartheta}$, where $\Gamma(\cdot)$ is the gamma function.

²² Implicitly, I assume the Fréchet scale is common across e , this corresponds to $A_{ik0} = A_{ik1} \equiv A_{ik}$.

In turn, in the lower tier, I assume that efficiency units of land $\tilde{H}_{ik}$ allocated to the production of crop k can be used for the production of export eligible or export ineligible varieties e of that crop k using efficiency units of land $\tilde{H}_{ike}$. Then, the following constraint will hold:

$$\left(\tilde{H}_{ik0}/A_{ik0}\right)^{\frac{\theta}{\theta-1}} + \left(\tilde{H}_{ik1}/A_{ik1}\right)^{\frac{\theta}{\theta-1}} = \tilde{H}_{ik}^{\frac{\theta}{\theta-1}}. \quad (4)$$

Similarly, $\theta > 0$ governs the elasticity of transformation between varieties $e = 0$ and $e = 1$ of crop k . If θ is large, production will be skewed towards growing only for the domestic consumption or value chain (i.e. export-oriented) markets. If θ tends to be smaller, production will be skewed towards having some sales in both markets.

Next, I discuss the implications of these parameter values for the properties of the underlying statistical distribution. When $\theta > \vartheta > 1$, draws between export-eligible avocados and non-eligible avocados are more similar than draws between maize and avocados. In contrast, if $\theta = \vartheta > 1$, then draws between export eligible and non eligible varieties of a crop are just as similar as draws across crops. When $\vartheta > \theta > 1$, productivity draws across crop types are more similar than across $e = 0$ and $e = 1$ varieties of crops, which intuitively seems unlikely (Farrokhi and Pellegrina, 2023), so I assume (and find empirically) that $\theta > \vartheta > 1$.

Costs and land shares

Following from the agricultural production function introduced in equation 1, the unit cost of crop k in plot ℓ with standards e , $c_{ike}(\ell)$, is given by:

$$c_{ike}(\ell) = w_{iA}^{\alpha_{ke}} p_{ix_k}^{\beta_{ke}} \left(\frac{r_i}{z_{ike}(\ell)} \right)^{\gamma_{ke}}, \quad (5)$$

where w_{iA} is the prevailing agricultural wage in municipality i , p_{ix_k} is the price of the intermediate input (i.e. fertilizer) used for producing crop k , and r_i is the rental rate of land in region i .

Next, I can derive the shares of land devoted to each crop in equilibrium. Since all farmers are identical, the land share is equivalent to the probability a farmer chooses to grow crop k and variety e in their plot ℓ , and the probability arises from the underlying Fréchet distributed variable $z_{ike}(\ell)$. The share of land H_i allocated to production crop k with export eligibility e , η_{ike} , is

$$\eta_{ike} = \underbrace{\frac{\lambda_{ike}^{\theta}}{p_{ik}^{\theta}}}_{\eta_{ie|k}} \times \underbrace{\frac{A_{ik}^{\vartheta} p_{ik}^{\vartheta}}{P_i^{\vartheta}}}_{\eta_{ik}}. \quad (6)$$

As the product of two terms, η_{ike} is formed by the product of the share of total land allocated to crop k used to grow a crop with export eligibility e ($\eta_{ie|k}$) and the share of total agricultural land

used to grow both export eligible and ineligible varieties of crop k , η_{ik} . Above, I define

$$\lambda_{ike} \equiv p_{ike} \left(\frac{w_{iA}}{p_{ike}} \right)^{-\alpha_{ke}/\gamma_{ke}} \left(\frac{p_{ixk}}{p_{ike}} \right)^{-\beta_{ke}/\gamma_{ke}}, \quad (7)$$

where p_{ike} is the price of crop k with export status e in region i , and I define $p_{ik}^\theta \equiv \sum_{e \in \{0,1\}} \lambda_{ike}^\theta$, recalling that A_{ik} is the scale parameter of land productivity for crop k and $P_i^\vartheta \equiv \sum_{l \in \mathcal{K}} A_{il}^\vartheta p_{il}^\vartheta$.

Using competitive factor shares (land's revenue share is γ_{ke}) and the fact that the rental rate of efficient land use will equal P_i , the aggregate quantity produced of crop k with export eligibility e in region i will be given by:

$$q_{ike} = \eta_{ike} \times H_i \times \frac{P_i}{\gamma_{ke} \times p_{ike}}. \quad (8)$$

Next, I define the share of production of crop k with export eligibility e in region i to be $s_{ike} \equiv \frac{q_{ike}}{q_{ik}}$. Then, after inverting equation 8 (see Appendix for details), the inverse supply elasticity holding q_{ik} constant can be written as:

$$\left. \frac{\partial \log p_{ike}}{\partial \log q_{ike}} \right|_{q_{ik}} = \frac{\gamma_{ke}}{(\theta - \gamma_{ke})(1 - s_{ike})}. \quad (9)$$

When q_{ike} represents a relatively small share of a municipality's production, this reduces to $\gamma_{ke}/(\theta - \gamma_{ke})$, which will be useful to understand the entry stage of the plant problem.

What does this imply? When the cost share of land goes to zero, the regional crop supply becomes perfectly inelastic and fixed. When the cost share of land approaches 1, in contrast the elasticity of supply becomes larger. On average, I find the average land cost share across all crops k and regions i to be $\bar{\gamma}_{e=0}$ to be 0.57 for low phytosanitary quality crops in Section 6 and $\bar{\gamma}_{e=1} = 0.4$ for export eligible crops. In contrast, as $\theta > 1$ becomes larger, farmers are able to more easily shift between producing the export eligible and non-export eligible varieties as their relative prices change, which decreases the elasticity of supply for a given variety e of crop k since large price changes will now yield larger shifts into different varieties.

Downstream Sector: Plants

Here, I introduce the setup for the downstream packing plants v in sector P which source their inputs from nearby farms and are located in urban regions. Within the sparse matrix of trade in this model, however, the sourcing decision of these plants cannot yield analytical expressions as in Antràs et al. (2017). Therefore, for tractability I make the simplifying assumption that plants can only source from municipalities in their rural periphery, $R(i)$ (see Stylized Fact 4 for justification of this).

Each representative packing plant v is associated with a given crop k and is assumed to only be able to source and sell that specific crop. The location of each plant is given by i . Plants take

world prices $\tilde{p}_{Wk}$ as given, which locally is the price $\tilde{p}_{ik1}$ ²³, and given their location i , will only be able to sell their output at a price $\tilde{p}_{Wk}\tau_{Wik}$ (or lower), where τ_{Wik} are derived from the minimum distance from location i to a port or border. Here I use $\tilde{p}$ to represent the prices consumers pay or what intermediaries receive for their output, and regular p to represent the prices farmers receive. As there are no intermediaries for low phytosanitary crops ($e = 0$) this implies $\widetilde{p}_{ik0} = p_{ik0}$ for all regions i and crops k .

Plants can only purchase high phytosanitary quality standard ($e = 1$) crops, and only plants directly purchase high quality crops. To enter, firms need to pay a fixed cost F_k (these costs vary by crop k), but have perfect forward looking information as to regional crop supply and their productivity A_{vk} . They annualize these fixed costs over a given time horizon, and the annualized fixed cost is f_k . Within a given location, plants choose the quantity (q_{ivk}) to purchase of crop k of high quality $e = 1$ from farmers and draw their Hicks-neutral productivity A_{ivk} ²⁴. The firm's production function, y_{ivk} , is assumed to be linear and given by:

$$y_{ivk}(q_{ivk}) = A_{vk}q_{ivk}. \quad (10)$$

Recall that each firm v in region i holds some degree of oligopsonistic power in the market for crop k of high quality h , so the price paid to farmers at the factory gate p_{ivk} will be lower than the price the firm receives, $\tilde{p}_{ik1}$ ²⁵ (note that prices paid at the farm gate will be marked down further by trade costs²⁶). Therefore, the firm's profit maximization problem is as follows:

$$\max_{q_{ivk}} \pi_{ivk} = \max_{q_{ivk}} \tilde{p}_{ik1}A_{vk}q_{ivk} - q_{ivk}p_{ivk}(q_{ivk}) - f_k \quad (11)$$

Taking the first order condition with respect to q_{ivk} yields the markdown condition that determines the wedge between prices received by intermediaries in region i , $\tilde{p}_{ik1}$ and the marginal cost (factory-gate price) of a unit of a crop with high quality, p_{ivk} (net of productivity) as:

$$\frac{\tilde{p}_{ik1}}{p_{ivk}} = \frac{\left(1 + \frac{1}{\varepsilon_{vk}}\right)}{A_{vk}} \quad (12)$$

where $\frac{1}{\varepsilon_{ivk}} \equiv \frac{\partial \log p_{ivk}}{\partial \log q_{ivk}}$ ²⁷. Here, I define $m_{ivk} \equiv \frac{\left(1 + \frac{1}{\varepsilon_{vk}}\right)}{A_{vk}}$ to be the firm's (productivity inclusive)

²³If these packing plants sell abroad, by the no arbitrage condition, they must sell at $\tilde{p}_{ik1} = \frac{\tilde{p}_{Wk}}{\tau_{Wik}}$ or lower.

²⁴In the appendix, I consider an extension where I allow for labor in the production function.

²⁵To back out the farmgate price in rural regions corresponding to the firm located at urban region $i \in \mathcal{U}$, or the prices in regions $R(i)$, note that the no arbitrage condition yields that $p_{ik1} = \frac{p_{vk}}{\tau_{i(v),i,k}}$ for all $i \in R(i(v))$.

²⁶Jung et al. (2021) and Chatterjee (2019) conclude that market power of intermediaries is closely linked to their distance to sourcing farms.

²⁷Note that $q_{ivk} = \sum_{j \in R(i(v))} q_{jkh}$

markdown from the marginal revenue product of the crop k (i.e. the market price net of productivity)²⁸. Additionally, I assume that plants are able to purchase some of the intermediate inputs required for export eligible crop k , x_k at a bulk discount. These plants sell these intermediate inputs to their supplying farms at this reduced cost.

Plant entry condition I begin by using the first order condition from equation 12 to solve for the plant entry condition at the optimum price (i.e. when the firm posts a factory gate price to all farms of $p_{ivk} = \frac{\tilde{p}_{ik1}}{m_{ivk}}$) and by requiring the profit condition to be greater than zero:

$$A_{ivk}q_{ivk}(\tilde{p}_{ik1}/m_{ivk})[1 + \varepsilon_{ivk}]^{-1} \geq f_k. \quad (13)$$

That is, with knowledge of A_{ivk} , f_k , and how supply responds to prices ε_{ivk} and $q_{ivk}(\tilde{p}_{ik1}/m_{ivk})$, the exogenously determined price of the crop k , $\tilde{p}_{ik1}$, will determine whether or not a firm v will enter in region i and produce crop k .

To make progress on the condition for whether or not a firm will enter, I assume that it possesses perfect information and is able to calculate the regional (inverse) elasticity of supply²⁹ before it enters the market, using information on the equilibrium prices (quantities) in the previous period. The firm does not consider the effect of its actions on the equilibrium prices (quantities) of other crops. Note that when the firm is just deciding to enter, all municipalities $j \in R(i(v))$ have production $q_{jkh} = 0$, and so equation 27 implies that for these regions $\frac{\partial \log p_{ike}}{\partial \log q_{ike}} = \frac{\gamma_{ke}}{\theta - \gamma_{ke}}$. Therefore, $\frac{\partial \log q_{ivk}}{\partial \log p_{ivk}} = \sum_{n \in R(i(v))} \frac{\theta - \gamma_{nke}}{\gamma_{nke}} (\tau_{i(v),n,k})^{-1}$. In the pre-period, the firm decides if it will enter, pays its (annualized) fixed cost f_k , and then draws its productivity A_{vk} . If the firm enters, it becomes the only processing firm allowed to operate in that region to produce (the high-quality h variety of) crop k ³⁰. The firm then calculates its optimal markup based on the previous period's equilibrium. However, the farmers learn of its initial markup, and in the following period, the firm must charge the same markup.

Manufacturing

The manufacturing good is produced in urban regions $i \in \mathcal{U}$. Each location has manufacturing productivity A_{iM} and features a linear production function of labor alone. Therefore, the output of the manufacturing good in region $i \in \mathcal{U}$ is given by $q_{iM} = A_{iM}L_{iM}$.

²⁸It is important to note here that low productivity, in addition to the monopsony power of intermediaries, may lead to large wedges between the farm gate price and the price of the processed output.

²⁹Note here that the relevant elasticity of supply the firm needs to calculate is $\varepsilon_{ivk} = \frac{\partial \log q_{ivk}}{\partial \log p_{ivk}}$, where $q_{ivk}(p_{ivk}) = \sum_{i \in R(i(v))} q_{ik1}(\frac{p_{ivk}}{\tau_{i(v),i,k}})$.

³⁰In the appendix, I consider an extension of this where I allow for multiple entrants in a given region.

5.3 Equilibrium

I assume that all markets are competitive except for the market for high quality crops, in which farmers can only sell directly to the processing firm for its crop, who acts as a monopsonist in that local market for the crop. In equilibrium,

1. the representative consumer in both urban and rural regions solves their utility maximization problem optimally given income and prices,
2. the inputs and outputs solve the representative farmer's problem given prices,
3. the representative firm solves its profit maximization problem conditional on input costs and world prices,
4. goods markets clear,
5. labor markets clear,
6. and finally prices are given by the no-arbitrage conditions condition on measures of trade costs.

In appendix section A.6, I provide an extended description of the equilibrium conditions as well as counterfactual equations.

5.4 Multiple equilibria and conditions for uniqueness

I now discuss multiple equilibria and provide uniqueness conditions for the model. In the case of increasing returns to scale, the challenge of multiple equilibria can be vexing, as hysteresis can yield results where the first entrant to a market can become dominant and drive out entry of other, potentially more productive, entrants. Below, I demonstrate that the static equilibrium is unique for a given set of entrants/packers, but I argue that the dynamic one may not be. The main concern comes from the decision of which packer for a given region and crop chooses to enter first. After entry of the first packer to a region, then for potential entrants into the same region there is less land with which to source their crop from and wages will be higher for agricultural labor. For potential entrants producing the same crop as a pre-existing entrant, although the rest of world prices have not changed, since packers sell locally, the pre-existing entrant has depressed local prices, limiting their potential entry. For instance, if avocado packers choose to enter a given region, they will depress the entry of other crop packers to their region, as well as depressing the entry of avocado packers elsewhere as well. In order to determine uniqueness numerically, the computation of every possible equilibrium is generally infeasible as this is an NP-hard problem, where even with a limited

number of metropolitan regions (75 here) and crops requires running a simulation $2^{|\mathcal{U}| \times |\mathcal{K}|}$ times. Therefore, to make progress on determining an optimal equilibrium, I provide conditions below under which one can determine the unique set of entrants over time.

Uniqueness of equilibrium in each period t To demonstrate the uniqueness of the equilibrium in each period t , suppose that the matrix $\delta_{|\mathcal{M}| \times |\mathcal{K}|}^h$ representing whether an intermediary is operating in region $i \in \mathcal{M}$ and producing crop $k \in \mathcal{K}$ with corresponding firm markups m_{ik} is taken to be given. Then, by mapping m_{ik} into high phytosanitary quality ($e = 1$) good-specific ad valorem trade costs, uniqueness follows from a logic similar to the proof in Bergquist et al. (2025). Therefore, the challenge is to prove the uniqueness of packing firm entry and markups. Markups are set on the basis of regional production, and so markups should be unique. However, characterizing the uniqueness of firm entry is more difficult.

Uniqueness of firm entry over time To recover a unique path of firm entry over time, I assume that in the initial period $t = 1$, exogenous world prices $\{p_{Wk}\}_{k \in \mathcal{K}}$ (or trade costs, τ) are too low (or in the case of trade barriers, high) to support the entry of any packing plants in the domestic economy. Subsequently, the exogenous world prices faced domestically rise, and with each firm calculating a separate hypothetical operating profit π_{vk}^{op} (that depends mostly on factors such as trade costs and the regional production function for crops), in continuous time there will be (generically) one unique first entrant. I provide a proof of this in the appendix.

Without assuming this unique path, however, the equilibrium is not assured to be unique. Therefore, in most simulations I take the number of entrants to be fixed, and discuss counterfactual simulations without considering the entry and exit of firms.

6 Estimation

In this section, I describe how I estimate the primary parameters of interest needed for my model. These parameters govern the differences in input usage for subsistence and AGVC-based farming, the elasticity of transformation between supplying crops locally and to AGVCs, and the elasticity of transformation across crops.

Input shares of labor, land, and intermediate inputs

To begin, I study whether agriculture targeted at domestic consumption differs from that produced for agricultural value chains in terms of the costs of inputs required. Anecdotal evidence suggests that processing firms or extension agents will often transfer knowledge of special techniques or inputs to farms that are required to fulfill phytosanitary standards. They may also transfer knowledge that saves on labor costs, or manages land use better. All of these conclusions suggest that I must empirically study whether and how input usage varies across farms that participate in AGVCs and those who do not.

To do so, I estimate the coefficients of my hypothesized farm level production function to recover the farm level input shares of labor, land, and intermediate inputs. Applying logs to equation 1, I obtain the following equation:

$$\log q_{ike}(\ell) = \log \bar{q}_{ike} + \alpha_{ke} \log l_{ike}(\ell) + \beta_{ke} \log x_{ike}(\ell) + \gamma_{ke} \log t_{ike}(\ell) + \gamma_{ke} \log z_{ike}(\ell). \quad (14)$$

To estimate this equation, I use data from the 2007 Agricultural Census. This provides estimates of labor supply from the number of family members, hired workers, and migrant workers³¹, the amount of land allocated to the production of each crop (in hectares), and the amount of fertilizer used in the farm unit (in tons).

To take theory to data, I assume that export-eligible crops ($e = 1$) use fertilizer (i.e. $\beta_{ik1} > 0$), and that export-ineligible crops ($e = 0$) do not use fertilizer (i.e. $\beta_{ik0} = 0$)³². As many farms do not apply fertilizer, making this assumption in advance allows me to estimate equation 14 for farms with fertilizer use and omit the third term from the regression ($\log x_{ike}(\ell)$) for farms without fertilizer use.

While I would ideally like to estimate this equation separately for each region, crop, and export status (which here refers to an indicator for any fertilizer use), in practice this results in many underpowered regressions, especially for crops that are not as prominent as, say, maize. Therefore, in practice, I estimate this equation separately at the state-crop-fertilizer use status level. Since I estimate equation at this level, $\bar{q}_{ike}$, a function of the cost shares in the production function, becomes the intercept of these regressions.

Next, to proxy for the fundamental agricultural productivity of plot ℓ in crop k , $z_{ike}(\ell)$, I use precise geolocators at the level of the *área de control* which allow me to merge farm level information with measures of land suitabilities from FAO GAEZ and EcoCrop described earlier.

However, there is reason to believe that there may be other factors that affect farmers' input decisions outside of their productivities that enter the error term in the above equation, even if my suitability measures capture plot level productivity $z_{ike}(\ell)$ perfectly. For instance, factors such as early and late-season weather shocks, and farmer characteristics may affect farmer decisions, which would be unobserved by the econometrician in the above model. To instrument for these, I use information with early, middle, and late season temperature (measured as number of days with maximum (minimum) temperature above (below) the temperature maximum (minimum))

³¹Unfortunately, this is measured somewhat crudely, and labor inputs are not distinguished by crop. To provide checks of the accuracy of this procedure, in the main text I report results for farm units growing only one crop. In the appendix, I report results for farms with multiple crops, where I allocate input usage based on hectares planted.

³²This assumption closely matches the data, where the majority of exporting farms across all crops use inputs such as fertilizer, and average fertilizer use is much lower for non-contracting or exporting farms, and a large percentage do not use fertilizer at all.

thresholds established in the FAO EcoCrop database, daily rainfall from CONAGUA, and information on farmer characteristics from INEGI such as family vs non-family farming. I use these characteristics as instruments for the various inputs, and report the two-stage least squares (2SLS) coefficients of this estimation procedure³³.

I report summary statistics of this procedure in Table 2 below, where I take the median across all state-level estimates. On average across all crops, for farm units using fertilizer, labor costs represent 22% of unit costs in agricultural production, fertilizer and material usage represents 39% of costs, and land use represents the remaining 39% of costs. For farm units without fertilizer usage, labor represents 43% of costs, and land comprises the remaining 57%. There is some heterogeneity across crops, with some crops having input shares that suggest increasing returns to scale ($\alpha + \beta + \gamma > 1$) and some with decreasing returns to scale (especially for estimates of only α and γ in farms without fertilizer use). That said, a majority are close to one (i.e. within 0.3 of 1), or constant returns to scale, which I take as evidence against the assumption that firms directly have large fixed costs of entry (Gáfaró and Pellegrina, 2022). In appendix table A.4, I report the results of this procedure separately for farms that grow multiple crops. The results remain qualitatively similar, although results for certain crops change slightly. In particular, farms without intermediate input usage generally exhibit smaller returns to scale, suggesting that the inputs transferred in modern supply chains are crucial to achieving larger scale³⁴.

θ – Elasticity of transformation across export-eligible and ineligible varieties of a given crop

A key determinant into whether certain regions may benefit from agricultural value chain, or modern, production, is the degree to which a region is able to shift from producing only locally-consumed varieties to varieties eligible for sale in larger retailers (who generally have higher quality standards) or abroad. Therefore, I seek to estimate θ , which governs elasticity of transformation of land used in production of crops sold locally versus output sold in agricultural value chains. A low value of θ suggests that even after the entry of an agricultural value chain, farms will not shift their production towards techniques required for modern production and are less likely to benefit. In contrast, higher values represent a larger flexibility of land to respond to the higher prices that modern markets provide for farmers.

The estimation procedure for θ stems from rewriting equation 6 in logarithms. I make one simplifying assumption for calibration and estimation, that $A_{ike} = \bar{A}_{ik}, \forall i, k, e$ so that land productivity varies at the crop level k by region i , but that land is just as productive in producing export eligible vs. non-export eligible varieties. Then in $\eta_{ie|k}$ in equation 6 the term A_{ike}^{θ} disappears. Taking

³³I also estimate the above equation using the instrumental variables estimator for the linear correlated random coefficients (IVCRC) in the online appendix.

³⁴In ongoing work, I use the residuals of these separate regressions to measure farm level TFP in 1991 and 2007. I then regress this against measures for the duration of AGVC presence in a region to understand the effects of AGVCs on farm level productivity.

Table 2: Input cost shares by crop

Crop	Farms w/ int. inputs				Farms w.o. inputs		
	α	β	γ	$\alpha + \beta + \gamma$	α	γ	$\alpha + \gamma$
Avocados	0.093	0.339	0.618	1.050	0.200	0.918	1.118
Bananas	0.093	0.145	0.262	0.501	0.237	0.511	0.748
Barley	0.462	0.214	0.171	0.846	0.186	0.371	0.557
Beans	0.446	0.453	0.450	1.348	0.061	0.336	0.397
Coffee	0.148	0.476	0.548	1.172	0.002	0.323	0.325
Cotton	0.194	0.501	0.411	1.106	0.260	0.220	0.480
Lemons	0.084	0.093	0.997	1.174	0.020	0.947	0.967
Maize	0.683	0.418	0.533	1.634	0.142	0.448	0.590
Mango	0.153	0.053	0.510	0.716	0.107	0.881	0.987
Oats	0.302	0.353	0.383	1.038	0.034	0.299	0.333
Oranges	0.154	0.237	0.862	1.253	0.245	0.828	1.073
Sorghum	0.223	0.334	0.468	1.025	0.114	0.442	0.556
Soy beans	0.378	0.799	0.337	1.514	0.079	0.695	0.774
Sugar	0.130	0.134	0.640	0.905	0.025	0.420	0.445
Tomatoes	0.123	0.238	0.694	1.055	0.056	0.304	0.361
Wheat	0.184	0.411	0.523	1.118	0.071	0.329	0.400

Notes: Estimated using two-stage least squares with robust standard errors, using the 2007 Agricultural Census. I take the median of all state level estimates for each crop. α is the cost share of labor, β is the cost share of fertilizer, and γ is the cost share of land. Since I cannot estimate equation 14 if fertilizer use is zero, I split the sample into farm units with positive and zero fertilizer use and estimate the coefficients separately for both groups. Sample only consists of farms producing a single crop.

logarithms of the remaining equation yields:

$$\log \eta_{ie|k} = \theta \log p_{ike}^{\frac{1}{\gamma_{ke}}} + \theta \log w_i^{-\alpha_{ke}/\gamma_{ke}} + \theta \log p_{ix_{ke}}^{-\beta_{ke}/\gamma_{ke}} - \log \left[\lambda_{ik0}^{\theta} + \lambda_{ik1}^{\theta} \right]. \quad (15)$$

To simplify the equation further, I rewrite this in terms of differences between the share of land allocated to crop k in region i with export status $e = 1$ and that with export status $e = 0$. Doing so, I obtain

$$\log \left(\frac{\eta_{i1|k}}{1 - \eta_{i1|k}} \right) = \theta \log \overbrace{\frac{p_{ik1}^{\frac{1}{\gamma_{ik1}}}}{p_{ik0}^{\frac{1}{\gamma_{ik0}}}}}^{\text{Price term}} + \theta \times \underbrace{\frac{\alpha_{ik0}\gamma_{ik1} - \alpha_{ik1}\gamma_{ik0}}{\gamma_{ik0}\gamma_{ik1}}}_{\text{Wage term}} \log w_i + \theta \log \overbrace{\frac{p_{ix_{k1}}^{-\beta_{ik1}/\gamma_{ik1}}}{p_{ix_{k0}}^{-\beta_{ik0}/\gamma_{ik0}}}}^{\text{Input price term}}. \quad (16)$$

To understand the motivating intuition behind this equation, consider the case when the input shares are equal across export eligibility types (i.e. $\alpha_{ik0} = \alpha_{ik1} \forall i, k$, etc.). Then, the expression above simplifies to

$$\log \left(\frac{\eta_{i1|k}}{1 - \eta_{i1|k}} \right) = \theta \log \left(\frac{p_{ik1}}{p_{ik0}} \right)^{\frac{1}{\gamma_{ik}}} + \theta \log \left(\frac{p_{ix_{k1}}}{p_{ix_{k0}}} \right)^{-\beta_{ik}/\gamma_{ik}}. \quad (17)$$

In words, equation 17 states that the percent relative share of land allocated to crop k dedicated to growing export-eligible crops increases when the export premium does as well, at a rate given by θ/γ_{ik} . Likewise, the relative share of land allocated to export eligible crops is negatively affected by the premium in intermediate input costs required to fulfill phytosanitary costs at a rate $-\theta \times \frac{\beta_{ik}}{\gamma_{ik}}$, which is increasing as the input share of intermediate inputs, β_{ik} . Therefore, the elasticity θ can be identified using cross sectional variation in the shares of land used for crops destined for international markets, the export price premium, and the ratio of intermediate input costs used for exporting versus growing for domestic markets.

Table 3: θ estimation results

Dependent variable: $\log\left(\frac{\eta_{i1 k}}{1-\eta_{i1 k}}\right)$	(1)	(2)	(3)	(4)	(5)	(6)
Σ	0.0152 (0.00317)	-0.0701 (0.00284)	1.324 (0.106)	1.391 (0.377)		
Price term					1.039 (0.457)	1.050 (0.165)
Wage term					0.820 (0.0752)	1.416 (0.251)
Input price term					0.804 (0.148)	1.084 (1.018)
Type	OLS	PPML	IV	Poisson IV	IV	Poisson IV
N	9301	67037	9301	66997	9301	66997
First-stage/KP F. stat.			52.56		61.737	

Robust standard errors in parentheses

Notes: Estimated using land share and outut data from the 2007 Agricultural Census and prices and wages from the 2012 National Agricultural Survey (ENA). Σ is a shorthand for $\log \frac{\frac{1}{p_{ik1}^{\frac{1}{\gamma_{ik1}}}}}{\frac{1}{p_{ik0}^{\frac{1}{\gamma_{ik0}}}}} + \frac{\alpha_{ik0}\gamma_{ik1} - \alpha_{ik1}\gamma_{ik0}}{\gamma_{ik0}\gamma_{ik1}} \log w_i + \log \frac{p_{ixk1}^{-\beta_{ik1}/\gamma_{ik1}}}{p_{ixk0}^{-\beta_{ik0}/\gamma_{ik0}}}$, so here I estimate only one θ . Included instruments are distance from the municipality to the nearest port, as well as dummies that indicate whether a agroexporter is located nearby.

In Table 3, I report the results of estimating equation 16. I do so in two parts: in columns 1-4), I estimate the value for only θ alone, using the estimating equation

$$\log\left(\frac{\eta_{i1|k}}{1-\eta_{i1|k}}\right) = \theta \Sigma \equiv \theta \times \left[\log \frac{\frac{1}{p_{ik1}^{\frac{1}{\gamma_{ik1}}}}}{\frac{1}{p_{ik0}^{\frac{1}{\gamma_{ik0}}}}} + \frac{\alpha_{ik0}\gamma_{ik1} - \alpha_{ik1}\gamma_{ik0}}{\gamma_{ik0}\gamma_{ik1}} \log w_i + \log \frac{p_{ixk1}^{-\beta_{ik1}/\gamma_{ik1}}}{p_{ixk0}^{-\beta_{ik0}/\gamma_{ik0}}} \right]. \quad (18)$$

In my preferred estimation in column 4), I obtain an estimated θ of 1.39. This value is quite low, suggestive of high frictions in shifting into export-oriented agriculture. Note that the share of land allocated to crops destined for foreign markets is quite frequently zero (in more than 86% of observations), and so Poisson estimation is more suited to this context, as compared to columns 1), 3), and 5) where I report the results with OLS/IV, and am forced to drop these observations.

To address the potential endogeneity of remote regions facing *higher* export premiums due to depressed local prices for crops as well as classic simultaneity bias, I instrument the independent variables in columns 3)-6) with the road distance of the municipality to the closest port/border crossing, as well as indicators for the presence of an agroprocessor in the nearby vicinity (at different distances). This procedure increases the estimated θ from the regressions without instruments significantly. However, in columns 5-6), I estimate the coefficients separately, and do not impose restrictions that θ must be identical. In Column 6), I find that the estimated θ s are largely similar, although they are somewhat different in Column 5).

ϑ – elasticity of transformation between crops

A crucial input into the model is ϑ , which governs the elasticity of transformation between crops. In particular, a higher value of ϑ suggests more flexibility of land shares to changes in local crop prices. The parameter also governs the degree to which regions will be specialized in the most suitable and profitable crop. A lower value of ϑ implies that land specialization will be more incomplete, and even areas that are highly suitable for one crop will nonetheless grow others.

Equation 6 implies that in any given time period t , I have

$$\log \eta_{ikt} = \frac{\vartheta}{\theta} \log \left(\lambda_{ik0t}^{\theta} + \lambda_{ik1t}^{\theta} \right) + \vartheta \log A_{ikt} - \vartheta \log P_{it},$$

where η_{ikt} is the share of arable land in region i devoted to crop k , A_{ikt} is crop specific productivity, and P_{it} is a regional productivity index. Here, the regional productivity index can be controlled for through use of region-time fixed effects. However, the crop specific productivity term is more vexing. One possibility is to use potential suitability indices from GAEZ directly in the regression to proxy for A_{ikt} , however this may not be a perfect proxy for productivity, as it does not vary during my sample period. Another option, which I follow, is to assume $\log A_{ikt}$ can be written as $\log A_{ik} + \log A_{state(i)kt} + \log A_{it} + \log \xi_{ikt}$, so crop-municipality specific productivity is time-invariant, and $\log \xi_{ikt}$ are “errors” in cropping decisions that are systematically uncorrelated with productivity. Then, I can estimate ϑ without information on agricultural productivities (or wages, following the same assumption) through the usage of region-time, state-crop-time and region-crop fixed effects:

$$\log \eta_{ikt} = \frac{\vartheta}{\theta} \log \left(\lambda_{ik0t}^{\theta} + \lambda_{ik1t}^{\theta} \right) + \Lambda_{i \times t} + \Lambda_{s(i) \times k \times t} + \Lambda_{i \times k} + \varepsilon_{ikt}.$$

However, the expression above simplifies when the municipality only produces crop k for the domestic market $e = 0$. As seen in the case for the θ estimation procedure, over 80% of crop-municipality pairs in 2007 have no exporting, so this simplified case applies in most cases³⁵.

35

Without any exporting (i.e. $\lambda_{ik1t} = 0$), I obtain the following equation

$$\log \eta_{ikt} = \vartheta \log \left(p_{ikt}^{1/\gamma_{ke}} \right) + \vartheta \log \left(w_{it}^{-\frac{\alpha_{ke}}{\gamma_{ke}}} \right) + \Lambda_{i \times t} + \Lambda_{s(i) \times k \times t} + \Lambda_{i \times k} + \varepsilon_{ik,t}. \quad (19)$$

Conditional on estimates of the cost share of labor and land in the Cobb-Douglas agricultural production function $(\alpha_{ke}, \gamma_{ke})$ described below, I can capture the sensitivity of within-group land share changes to agricultural crop prices, ϑ . I present the results of my OLS regression in Table 4. Column 1 presents the results of estimating equation 19 using OLS. The estimated coefficient for ϑ is almost a true zero. One reason for this may be classic simultaneity bias inherent in supply and demand estimation. Therefore, I search for an instrument that exogenously shifts demand, and thus the prices of agricultural crops, which will yield an unbiased estimate of ϑ . To instrument for prices, I follow Roberts and Schlenker (2013) and use lagged yield shocks as an instrument for estimating supply elasticities. As Roberts and Schlenker (2013) note, the use of lagged yield shocks as an instrument is possible “because past weather-induced supply shocks affect inventories, and inventories affect the futures price in subsequent periods[, and the] key assumption for consistent identification of the supply elasticity is that past weather-induced supply shocks have zero covariance with unobserved supply shifters in the current period.” The authors argue that, to control for unobserved supply shifters, one should also control for contemporaneous yield shocks in the supply equation.

I present the instrumented results of such a procedure in Column 2. My estimated ϑ here is approximately 1.3. However, one concern is that the instrumental variables strategy of Roberts and Schlenker (2013) is intended for use only of annual crops, where planting decisions are made each year before the growing season, rather than annual crops. Furthermore, my sample from SIAP includes many artisanal crops likely intended only for domestic consumption, for which well defined markets may only be present locally (at best). Therefore, in column 4, I restrict my sample to only the top 5 annual crops by agricultural production value in Mexico, which includes crops like maize and sorghum, storable commodities for which futures markets are well defined and Mexico is a known exporter of. Upon performing this subsetting, my estimated ϑ is 1.6, which falls within the range of similar results in the literature. For instance, Bergquist et al. (2025) estimates a range of ϑ s from 1.8-2.9, Farrokhi and Pellegrina (2023) finds 2.05, Zavala (2025) finds 1.35, and Sotelo (2020) estimates a value most closely related to my results – 1.658.

$A_{i,k}$ – productivity of region i in growing crop k

Similarly to Costinot et al. (2016), I link the scale parameters $A_{i,k}$ to measures of fundamental crop suitability that come from the FAO (I use a mix of suitability indices either provided by the FAO Global Agro-Ecological Zones project or derived from FAO EcoCrop suitability indices, both rescaled to be mutually compatible, in a range where 0 is totally unsuitable and 1 is highly

suitable). The upside of this procedure is that it is transparent and illuminates how the geography of production depends upon fundamental agronomic and ecological conditions, relative to the importance of agricultural value chains in production. In contrast, an alternative in the literature is to use a maximum likelihood type approach which rationalizes equation 6 using nonlinear optimization methods and observed data. However, in this approach, one drawback is that any sources of productivity that arise from economies of scale will necessarily be loaded onto the estimates of A_{ik} . In appendix figure A.20, I compare my results with a maximum likelihood type approach to the fundamental agro-ecological suitability method displayed in the main text.

The primary drawback of matching agricultural productivities A_{ik} to FAO suitability indices is that any production region with productivities larger than zero will always have non-zero production. As a result, many regions feature low, but non-zero production of commodities where they actually have none. Although this degrades the model fit in the cross-section, this assumption is more ideal for computing counterfactuals. For instance, if $A_{ik} = 0$ in the model (which is necessary to rationalize any areas of zero production in the maximum likelihood approach), any municipality currently not producing a given crop will continue not to do so in the counterfactual. Therefore, it is important to use actual measures of crop suitability to examine other potential areas of production when examining counterfactuals such as a decrease in fixed costs.

f_k – sunk costs of entry for downstream plants

Next, I move onto estimating, f_k , the (annualized) fixed cost for downstream agointermediaries. Using the first order conditions following from equation 10, I can derive the firm's operating profits

$$\pi_{vk}^{op} = A_{vk} \tilde{p}_{ikh} q_{vk} (\tilde{p}_{ikh} / m_{vk}) [1 + \varepsilon_{vk}]^{-1}. \quad (20)$$

I can observe some outcomes that compose this equation from farm microdata. I can estimate ε_{vk} based on my model assumptions (and information on regional level crop production, combined with estimates of $\alpha_{ke}, \beta_{ke}, \gamma_{ke}, \theta, \vartheta$). However, there are some parameters that I do not have knowledge of. For instance, $\tilde{p}_{ikh}$ has no direct analogue in microdata³⁶, as well as $q_{vk} = \sum_{n \in R(i(v))} q_{nkh}$, the amount of crop k sold to processors and exporters. Further, I cannot observe the individual firm's productivity A_{vk} .

To make progress, I am unaware of each firm's draw of A_{vk} , but I assume the firm knows in advance the distribution of productivity $G(A_{vk})$ before deciding to enter. Therefore, the entry condition is modified here, and requires that the firms *expected* operating profit is greater than the

³⁶I do not observe the factory gate prices posted by processing firms in the economic census, although I do know this information for avocado processors in my setting that I can use as a sanity check.

Table 4: ϑ estimation results

	log_share			
	(1)	(2)	(3)	(4)
$\log_{price_t}ransform$	0.010***			
(0.001)	0.307***			
(0.028)	0.002			
(0.003)	0.441***			
(0.066)				
$\log_{wage_t}ransform$	-0.035***			
(0.001)	0.139***			
(0.017)	-0.008**			
(0.003)	0.052***			
(0.010)				
$\log_{yield_{shock}}$	-0.081***			
(0.004)		-0.124***		
(0.006)				
$muncode_{year}$	x	x	x	x
$munc_{crop}$	x	x	x	x
$state_{crop_{yr}}$	x	x	x	x
Observations	375434	375434	104766	104766
S.E. type	hetero	hetero	hetero	hetero
R^2	0.963	-	0.985	-

Significance levels: * p < 0.05, ** p < 0.01,

*** p < 0.001. Format of coefficient cell: Coefficient (Std. Error)

Notes: Estimated using robust SEs.

annualized sunk cost: $\mathbf{E}_{A_{vk}} \left[\pi_{vk}^{\text{operating profit}} \right] =$

$$\mathbf{E}_{A_{vk}} \left[A_{vk} \tilde{p}_{ikh} q_{vk} (A_{vk} \tilde{p}_{ikh} / m_{vk}) (1 + \varepsilon_{vk})^{-1} \right] \geq f_k. \quad (21)$$

There is only one firm per crop per region, and so without free entry, there is no reason to think the inequality above goes to zero. However, I know the matrix $\delta_{|\mathcal{M}| \times |\mathcal{K}|}^h$ of pre-existing entrants (and region-crop pairs without entrants) from my data in Mexico. Then, given a certain draw of $\mathbf{A}$ for all entrants and non-entrants, one estimator is to find the minimum of operating profits across all entrants. Next, conditional on the same draw of $\mathbf{A}$, one can do the same for non-entrants and take the maximum of operating profits across all non-entrants. Then, averaging the minimum of operating profits across all entrants and the maximum of operating profits across all non-entrants should approach f_k from either side of the inequality.

Based on above description, one estimator for $\hat{f}_k$ is $\hat{f}_k =$

$$\left[\pi_{vk}^{\text{Avg. of min entrant} + \text{max non.entrant}} \right] =$$

Table 5: Fixed cost estimates from procedure

Crop	Avocados	Beans	Maize	Sugar	Tomatoes
Mean est. of f_k	170,263,586	129,464,871	42,230,978	179,057,318	132,074,815
95% C.I. of f_k	[159,858,034; 181,759,458]	[123,967,953; 136,632,891]	[40,371,865; 45,132,291]	[171,190,777; 189,209,592]	[126,572,576; 139,381,058]

Mean and 95% confidence interval for fixed effects for each crop calculated using 250 bootstrap simulations of the model, each time drawing different values of A_{vk} from a Weibull distribution with shape 1.5 (chosen to fit observed $G(A_{vk})$).

$$\left[\frac{\pi_{vk}^{\text{min. operating profit among entrants}}}{2} + \frac{\pi_{vk}^{\text{max. operating profit among non-entrants}}}{2} \right]. \quad (22)$$

To generate the operating profits of non-entrant, I simulate their entry by holding all other agrointermediaries fixed and running this through the model. Let $\hat{\psi}^d$ to represent the following moments in the data. First, we have $\tilde{p}_{ikh}q_{vk}$, which is the sum total of revenue for all agrointermediaries in a given region-crop combination, derived from firm microdata. Next, $E[G(A_{vk})]$ represents the expectation of the productivity distribution for each crop, where the mean comes from TFP estimation of agroexporters, also informed by microdata. Finally, I know the number of entrants and non-entrants, and so can count the number of non-entrants misclassified as having higher operating profits than entrants.

Comparing these moments with those from the simulation, denoted as $\hat{\psi}^s$, I employ the method of simulated moments (MSM) to estimate both fixed costs f_k and agrointermediary productivity. This estimator for the pair can be expressed as:

$$\{\hat{f}_k, \mathbf{A}\}^{MSM}(\mathbf{W}) = \arg \min \left[\hat{\psi}^d - \hat{\psi}^s(\hat{f}_k, \mathbf{A}) \right]^T \mathbf{W} \left[\hat{\psi}^d - \hat{\psi}^s(\hat{f}_k, \mathbf{A}) \right] \quad (23)$$

In practical application, I first find $\mathbf{A}$ for entrants as it is exactly identified, then iterate over draws of $G(A_{vk})$ for non-entrants and various values of f_k .

I report the estimates for the procedure described above in Table 5. I find relatively large estimates of fixed costs that correspond with both the clustering of production of that crop as well as the phytosanitary requirements for the crop. For crops such as Avocados and Sugar, which are heavily clustered and have large barriers to entry as well as trade, fixed costs are relatively large, at 170 and 180 million pesos respectively. However, for Maize, the estimated fixed costs are much smaller, at roughly 40 million pesos. Next, I perform some back of the envelope calculations in Table 6 to understand the magnitude of these estimates. Using unmatched data on yearly municipality revenue for each crop from SIAP, I compute the ratio between the estimated fixed

Table 6: Comparison between $\hat{f}_k$ estimate above regional yearly revenue of crop from SIAP

	Avocados	Beans	Maize	Sugar	Tomatoes
Minimum	∞	∞	86.32	∞	∞
25%	∞	21.70	0.32	∞	35.02
50%	352.70	7.64	0.12	217.04	4.40
75%	28.05	1.94	0.06	0.84	1.36
Maximum	0.02	0.11	0.00	0.07	0.07

Table reports ratio of $\hat{f}_k$ over sum of yearly revenue across all municipalities that correspond to a region. Source: SIAP data, 2007.

cost by crop³⁷ and yearly revenue for a given crop in a region. In 2007, the median region for avocado production would need to pay almost 353 times its regional revenue from that crop to establish a supply chain in the region, given the estimated fixed cost of 170 million pesos. In contrast, the median maize producing municipality would need only 12% of its regional revenue to establish a supply chain for maize, and establishing a supply chain appears feasible for all but the smallest of regions. This can explain the high concentration of avocado production compared to more uniformly distributed maize production.

$\tau_{ni,S}$ – Iceberg/ad-valorem trade costs

To calculate the value of ad-valorem trade costs in my setting, I draw upon wholesale agricultural market price information provided by the National System of Information and Market Integration (*Sistema Nacional de Información e Integración de Mercados*, or SNIIM). SNIIM reports monthly prices of various crops (differentiated by state of origin), in 48 wholesale food markets (*mercados de abasto*), in some cases back until 1989.

This data allows me to use information on the differences of crop prices across markets to inform me about trade costs in Mexico, with the possibility to examine how these costs differ across crops and over time. To do so, I can only infer that differences in prices of a given crop (and presentation) are reflective of trade costs at a given point in time t if that crop is sold in a *mercado de abasto* that corresponds to the urban center of the region in which it is produced (i.e. $i = U(j)$ for any producing region j), as well as in some destination region n . Then the difference in the price of crop k , with presentation q ³⁸, at time t between the origin market i and destination market n , would be reflective of the costs of moving the good between regions i and n ³⁹.

³⁷Note that this does not vary at the region level, only by crop.


³⁸Here, reflective of second degree price discrimination, crops are often sold at quantity discounts. In the data, a crop may be listed at a price per kilogram, but also a price per 20 kilogram box or basket.

³⁹Of course, this is only true of perfectly competitive settings where there is no processing occurring. This setting may alleviate some of these concerns, as these markets are wholesale markets where processing is less likely to occur. However, if intermediaries operating in these wholesale markets hold market power, the price gap would identify a combination of trade costs and intermediary markdowns, and would naively inflate these trade costs. That said, these wholesale markets are much more likely to consist of atomistic traders and sellers than larger agricultural value chains who source to supermarkets and grocery stores or foreign countries. Although I cannot rule out the existence of

Carrying out such a process would generate estimates of transportation costs between markets (which, without averaging, potentially vary by crop k , presentation q , or time t). However, to extrapolate these estimates to obtain estimates of trade costs between different municipalities, I relate these price differences to the travel distances required to ship these goods. Then using the estimated elasticity of trade costs to travel distance, I can back out trade costs for all regions (again, potentially varying across different crops k). In appendix section A.6, I describe the estimation process in further detail.

7 Results

Figure 4: Simulation results for avocado



Using my calibrated model⁴⁰, the first set of results I display represents the geography of crop production, where I take the set of agrointermediaries for each crop directly from the data (i.e. ⁴⁰I use the Pyomo optimization package in Python to solve the model (Bynum et al. (2021) and Hart et al. (2011)), in this section I present the results of several counterfactuals. A quickstart tutorial I contributed to can be found [here](#).

the agroexporter directly). Although I can produce maps of production for any crop in the model, I plot maps of avocado and maize production in Figures 4 and 5, given their status as the top two crops in Mexico by production value. Comparing to the plots of recent avocado and maize production in Appendix Figures A.11 and A.12, I find that the results match well to the currently observed allocation of production of both crops in Mexico. These total production plots are formed as the sum of production of both low ($e = 0$) and high ($e = 1$) phytosanitary quality crops, which I report for avocados and maize in Figures 5(b), 5(c), 5(f), and 5(g). Examining the plots of only high phytosanitary quality production that is processed through agrointermediaries against the plots of low quality production, in these plots the high concentration of production comes through the limited number of regions with processing firms, whereas the low quality production is more localized. However, high quality production is more of a prominent feature for avocados than maize, as avocados produced in agricultural value chains exceed 100 times low quality production. In contrast, for maize, low quality production is almost twice as large as high quality production, which reflects the fact that most maize in Mexico is produced for local, domestic consumption. These relative magnitudes influence the strength that moving to an equilibrium featuring more processing plants plays on the concentration of production.




Figure 5: Simulation results for maize

Table 7: HHI of production from SIAP data compared to model output

Crop	Data	Model	Model w/o intermeds	Model with share fit
Avocados	0.50	0.23	0.03	0.20
Maize	0.06	0.07	0.02	0.04


In Figures 5(a) and 5(e), I plot the equilibrium outcomes when there are no agrointermediaries. In general, these plots look very similar to the distribution of low ($e = 0$) phytosanitary quality production, which is intuitive because here only low quality production can meet domestic demand. As with low quality production, production without agrointermediaries is much less concentrated, and in particular, the distribution of avocado production is more similar to that displayed in the 1950's (see Figure A.9). Without agrointermediaries for avocados, the largest two states in terms of avocado production in my simulation results produce about $\approx 12\%$ of the total quantity produced. With the actual distribution of agrointermediaries, the same two states represented 44% of the total quantity of avocados produced in the country. This suggests the addition of fixed costs in my model is necessary to capture observed outcomes in clustering of agricultural production. Comparing to 2007 data, 88% of avocado production came from the top two largest producing states (Michoacán and Jalisco, respectively), so my model fails to capture the full extent of concentration seen in practice, although it matches observed concentration patterns more closely than a model that fails to model agricultural intermediaries. Even if we relax the assumption that A_{ik} should be matched to an empirical counterpart derived from agronomic suitability measures, the fit hardly improves. In figure A.20, I use nonlinear optimization to select the set of A_{ik} 's to best fit land shares. After running these calculations, the top two states still only produce 24% of the total avocados in Mexico, so merely assuming that the productivity terms capture all of the relevant dimensions of clustering by relaxing their dependence on suitability measures is not enough.

I summarize the results of this for avocados and maize in Table 7. I compute the Herfindahl-Hirschman index (HHI) of production for each crop across all regions using SIAP data, my model with agrointermediaries (packers), my model without agrointermediaries, and a model without agrointermediaries, with A_{ik} chosen to best fit land shares. For both avocados and maize, my model with agrointermediaries fits the data best, although it does not fully capture the full pattern of concentration in avocados. The model without agrointermediaries or packers performs considerably worse, and although the model with productivities fit to land shares does better, it underestimates concentration patterns relative to my preferred specification.

Finally, I show results using my full calibrated model, where I use my estimated fixed costs for crops and my model to sequentially determine entry of agrointermediaries into urban regions. In order to guarantee uniqueness, I perform the sequential entry procedure described earlier, where I

gradually remove barriers to trade for Mexico, thus increasing the effective world price received by agroprocessing firms. As the world price continuously increases, generically there is only one firm willing to enter, which is the firm with the highest hypothetical operating profit, a function of its draw of productivity. These firms enter, and then barriers continue to be removed until there is free trade with the rest of the world. In this equilibrium, there is no more entry or exit.

Figure 6: Hectares planted with avocados only of high quality ($e = 1$), with sequential agrointermediary entry as described in text




Therefore, I do not rely upon my data of locations of these agrointermediaries (aside from using this information to determine fixed costs); the location of agrointermediaries is solely a result of the model. In figure 6 I plot the equilibrium outcomes for high quality production (i.e. that which is carried out through agrointermediaries) of avocados. Although the regional patterns of production do not match perfectly to observed production of avocados in agricultural value chains, the regions where production does occur are in places with high agricultural suitability. Measuring the overall accuracy of the procedure to predict where avocado supply chains are located, the model overall has 80% accuracy, where 93.5% of the regions where avocado firms did not enter in the sequential model are regions in which no agrointermediaries exist in the agroprocessors data for 2007 (i.e. a low rate of false negatives), and 15.4% of the regions where avocado firms enter are regions in which agrointermediaries do actually exist in 2007 (a relatively high rate of false positives).

Of course, the distribution of agrointermediaries in this sequential equilibrium may be quite different than the observed distribution of agrointermediaries in the data. This highlights the role


of hysteresis in firm placement, as given the competitive effects of entry, pre-existing entrants can prevent entry in other regions. The large difference between the distribution in agointermediaries in the data and sequential game suggest that there may be a very large number of equilibria in this setting. Assuming spillovers of productivity from entrants to other locations (in terms of A_{vk} in other regions) or assuming firms subsidize input costs for contracting farms could improve the fit of the model to the data, where most avocado supply chains are located in the southwest of Mexico and few other regions.

Figure 7: Cluster size and suitability, model output vs. data



(a) Size and suitability relationship in model output

(b) Size and suitability (real distribution)



(c) Size and suitability in comparison model output without agointermediaries

In figure 7, I revisit some of the descriptive statistics earlier, and demonstrate that my model can rationalize some of my observed findings. In particular, I show that my model generates patterns of clustering that are not solely determined by agricultural suitability, as shown in Figure 1. For instance, if agricultural suitability solely explained where agricultural clusters were located, the plot below would have a coefficient close to one, and large R^2 . Instead, when I run the output of my model through the clustering algorithm presented in section 3.2, I find both a low coefficient and R^2 in both the model derived patterns of clustering (on the left) and the patterns of clustering observed in real data (right hand figure). This suggests my model does well in capturing the complexities in AGVC presence in Mexico. For point of comparison, I also present the same figure for the model output of when I do not include intermediaries in my model. Here, the lack of fixed costs allows for agricultural production to largely be a factor of the total land in a municipality and their relative suitabilities, and thus the relationship between cluster size and suitability is much closer, far closer than what I see in the actual data from SIAP.

8 Agricultural Supply Chains and Climate Change

In this section, I examine the effects of spatially fixed investments in agriculture on the ability to respond to climate change using my quantitative model. Climate change is expected to not only impact absolute crop productivities⁴¹, but also relative productivities within countries⁴². Therefore, the ability to switch to different crops as they become relatively more productive can potentially play an important role in reducing projected damages. Work such as Costinot et al. (2016) and Gouel and Laborde (2021) concludes that crop choice and switching will reduce global welfare losses due to climate change by 67% and 46%, respectively.

In my model, the fixed nature of locations of agricultural supply chains is quite important, as it reduces the ability for farmers to switch into other crops. Supply chains are able to provide greater market access to local farmers in a given crop, potentially keeping them specialized in that crop through higher prices, even as their land productivity declines. Unless productivity declines are large enough to lead to supply chain exit, these fixed investments will mute the adaptation responses of farmers (which would otherwise be governed solely via equation 3).

To quantify these forces empirically, I use current and projected agricultural productivities under climate change that come from the FAO GAEZ project⁴³. I use an ensemble of models

⁴¹ Across all crops found in the FAO GAEZ data and regions (either municipalities in Mexico and counties in the US), I find that under the IPCC AR5 CMIP5 RCP4.5 (Intergovernmental Panel on Climate Change Fifth Assessment Report Coupled Model Intercomparison Project Phase 5 Representative Concentration Pathway) experiment from 2041-2070, average crop productivity in Mexico will decrease 39.9% and in the United States will increase by 3.9%.

⁴² See figure 8 below, which displays relative changes in maize productivity within the US and Mexico.

⁴³ For more information on these indices and the climate change projections, see [online appendix \(section B\)](#)




Figure 8: Expected relative suitability changes for maize from 2010 to 2041
Source: FAO GAEZ data. The plot represents the change in maize suitability under the RCP 4.5 experiment (i.e. a middle-of-the-road climate projection) taken as the ensemble of several model predictions from 1980-2010 to 2041-2070, assuming CO_2 fertilization, high input levels, and no irrigation. Areas with missing data shown in gray. The legend is truncated at 2.

from the Coupled Model Intercomparison Project Phase 5 (CMIP5) under the Representative Concentration Pathway (RCP) experiment 4.5, which represents a "middle-of-the-road" climate experiment. I consider changes in productivities from the baseline period of 1981-2010 (current agricultural productivity) to the modeled period of 2041-2070.

To understand which crops will be subject to the largest relative productivity changes in the future, I select the top production clusters for a number of prominent crops grown in both the United States and Mexico using the algorithm described in Section 3.2. For each cluster, I compute the ratio of their predicted future suitability to their present suitability. I plot these results in Appendix Table A.2. I observe that agricultural productivity for most top production clusters is set to decrease, on average 44% for top Mexican agricultural clusters and 15% for US agricultural clusters. However, examining the ratio of relative productivity changes between the US and Mexico, I find that Mexican agricultural clusters will become relatively more productive crops such as Carrots, Onions, Sunflower, and Tomatoes, whereas the US will become increasingly productive in Maize, Potatoes, Sugar, and Wheat, suggesting that future cross-border agricultural trade may intensify in those products.

Next, I use my calibrated model to quantify the ability for agriculture to shift in response to climate shocks. I do so both using a model which includes packers and their fixed costs and one that does not feature fixed costs and intermediaries. I examine such effects in a partial equilibrium framework, where I hold world prices fixed to examine only the margins of adaptation that come

via reallocation of land⁴⁴. I find that crop switching occurs in 62% of land under the RCP4.5 projections in a model with current packer presence, but occurs in 78% of land in a model that does not feature fixed costs. This results in a reduction in crop switching of 21% in a model with fixed costs, suggesting the ability to reallocate production in response to climate shocks is more limited vis-a-vis models that do not incorporate these features. Despite the large changes in agroclimatic suitabilities due to climate change, I find that there is not an outsized response of packers to exit these markets – only 6 packing firms lose profitability such that it exceeds their fixed costs of production. As many firms have operating profits that are several times larger than their fixed costs at baseline, the market power of each firm within its own region and their profitability insulates these firms to a certain degree from climate shocks, thus limiting the response of entry and exit in response to climate shocks.

Relative to studies such as Costinot et al. (2016), Gouel and Laborde (2021), and Nath (2025) which examine the effects of climate change worldwide, I only examine Mexico and the United States. A key challenge to expanding this analysis geographically is that although global information on agricultural productivity and production is relatively easily accessible, similar information on the locations of agricultural supply chains is much more difficult to obtain. Still, this regional exercise suggests that these effects will be important more broadly. Compared to the literature which suggests that the brunt of climate change will be borne by countries with relatively low trade openness, in contrast, these supply chains mechanisms depress crop reallocations the most in areas with higher rates of supply chain penetration, which generally are more open to trade.

Of course, in practice, agricultural supply chains are formed by individual producers and actors. Many individual actors may not be able to profitably move alone, even if through coordination they could jointly achieve a more profitable outcome by moving together to a more suitable region. In this sense, by modeling the individual firms and actors that compose supply chains as one unit, my results may be considered to be an upper bound on the ability for agricultural supply chains to shift in response to changes in agricultural productivity. At the same time, untangling the effects of productivity decreases in agriculture and where tipping points exist at which these agricultural supply chains cease to become profitable is a crucial question to further study empirically. In particular, the increasing returns to scale of these value chains may ensure their current and medium-term profitability even in the face of a changing climate, but whether they can weather longer term shocks through adaptive investments remains unclear.

⁴⁴As Gouel (2025) highlights, doing so can underestimate welfare losses due to climate change due to imperfect substitution on the demand side, so this exercise serves to examine the role of supply chains in mitigating responses to shocks, rather than to estimate welfare losses.

9 Concluding remarks

In this paper, I study the factors driving the spatial clustering of agriculture. I argue that the presence of firms such as processors, packers, and exporters in the agricultural value chain, and ultimately their role in facilitating intra- and international trade, drives regional crop specialization. Further, I show that farms are limited in their opportunities to export primarily due to the granularity and concentration of these firms in only a handful of locations. I argue that the establishment of modern agricultural supply chains involves spatially fixed and large costs of entry which generate regions of high specialization in crop cultivation, but also leave many regions that are unable to pay these fixed costs unable to specialize. Given the link between specialization and rural incomes, these fixed costs associated with modern production generate winners and losers across rural regions that are not entirely driven by fundamental agricultural suitability. Finally, I show that explicit modeling of agointermediaries is necessary to explain the distribution of agricultural production in spatial models.

This paper also raises a number of questions to be addressed in future work. A full cross country accounting of patterns of clustering in agriculture would help reveal the barriers to entry into modern agricultural supply chains across regions. Most importantly, given the estimated effects of climate change on agriculture, an important exercise would be to understand how the granularity of agricultural supply chains and how their spatially fixed investments may affect the ability of crop choice to mitigate climate change worldwide, building upon my regional, partial equilibrium findings. In particular, more empirical evidence needs to be obtained as to the speed at which supply chains can move in response to climatic shocks, which my model suggests should be faster for supply chains with lower fixed costs of entry. Finally, my work highlights the importance of more fully documenting the role that governments and other actors play in helping establish supply chains. More research here will be crucial to understand the potential for (agro)-industrial policy to affect regional incomes. Ultimately, this paper shows that in order to understand the macroeconomic implications of agriculture, one must take stock of its underlying supply chains.

References

- Adamopoulos, T. and D. Restuccia (2021). Geography and agricultural productivity: Cross-country evidence from micro plot-level data. *Review of Economic Studies*.
- Alfaro-Ureña, A., I. Manelici, and J. P. Vasquez (2021). The effects of joining multinational supply chains: New evidence from firm-to-firm linkages. *Quarterly Journal of Economics*.
- Allen, T. and D. Donaldson (2020). Persistence and path dependence in the spatial economy. *NBER Working Paper No. 28059*.

- Ambikapathi, R., K. A. Schneider, B. Davis, M. Herrero, P. Winters, and J. C. Fanzo (2022). Global food systems transitions have enabled affordable diets but had less favourable outcomes for nutrition, environmental health, inclusion and equity. *Nature Food*.
- Antràs, P. and D. Chor (2013). Organizing the global value chain. *Econometrica* 81(6), 2127–2204.
- Antràs, P. and A. De Gortari (2020). On the geography of global value chains. *Econometrica* 88(4), 1553–1598.
- Antràs, P., T. C. Fort, and F. Tintelnot (2017). The margins of global sourcing: Theory and evidence from us firms. *American Economic Review* 107(9), 2514–64.
- Atkin, D., B. Faber, and M. Gonzalez-Navarro (2018). Retail globalization and household welfare: Evidence from Mexico. *Journal of Political Economy* 126(1), 1–73.
- Autor, D., D. Dorn, and G. H. Hanson (2013). The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review* 103(6), 2121–68.
- Bartelme, D. G., A. Costinot, D. Donaldson, and A. Rodríguez-Clare (2019). The textbook case for industrial policy: Theory meets data. *NBER Working Paper No. 26193*.
- Bellemare, M. F. and J. R. Bloem (2018). Does contract farming improve welfare? a review. *World Development* 112, 259–271.
- Bergquist, L., B. Faber, T. Fally, M. Hoelzlein, E. Miguel, and A. Rodríguez-Clare (2025). Scaling Agricultural Policy Interventions: Theory and Evidence from Uganda. *Working Paper*.
- Bergquist, L. F. and M. Dinerstein (2020). Competition and entry in agricultural markets: Experimental evidence from kenya. *American Economic Review* 110(12).
- Bleakley, H. and J. Lin (2012). Portage and path dependence. *The Quarterly Journal of Economics* 127(2), 587–644.
- Bleakley, H. and J. Lin (2015). History and the sizes of cities. *American Economic Review* 105(5), 558–63.
- Blyde, J. S., M. Busso, and D. Romero Fonseca (2020). Labor market adjustment to import competition: Long-run evidence from establishment data. *IDB Working Paper Series*.
- Borusyak, K., X. Jaravel, and J. Spiess (2021). Revisiting event study designs: Robust and efficient estimation. *arXiv preprint arXiv:2108.12419*.
- Bynum, M. L., G. A. Hackebeitl, W. E. Hart, C. D. Laird, B. L. Nicholson, J. D. Sirola, J.-P. Watson, and D. L. Woodruff (2021). *Pyomo—Optimization modeling in Python* (Third ed.), Volume 67. Springer Science & Business Media.
- Chatterjee, S. (2019). Market power and spatial competition in rural India.
- Comité Sistema Producto Aguacate (2005, June). Plan Rector Sistema Nacional Aguacate.
- Costinot, A. and D. Donaldson (2012). Ricardo’s theory of comparative advantage: old idea, new evidence. *American Economic Review: Papers & Proceedings* 102(3), 453–58.
- Costinot, A., D. Donaldson, and C. Smith (2016). Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world. *Journal of Political Economy* 124(1), 205–248.
- de Gortari, A. (2020). Global value chains and increasing returns. Technical report, Working Paper.
- de Janvry, A. and E. Sadoulet (2019). Transforming developing country agriculture: Removing adoption constraints and promoting inclusive value chain development. *Working Paper*.

- Dhingra, S. and S. Tenreyro (2020). The rise of agribusiness and the distributional consequences of policies on intermediated trade. *CEPR Discussion Paper No. DP14384*.
- Domínguez-Iino, T. (2025). Efficiency and Redistribution in Environmental Policy: An Equilibrium Analysis of Agricultural Supply Chains. *Journal of Political Economy*.
- Donaldson, D. (2018). Railroads of the Raj: Estimating the impact of transportation infrastructure. *American Economic Review* 108(4-5), 899–934.
- Fajgelbaum, P. and S. J. Redding (2022). Trade, structural transformation, and development: Evidence from argentina 1869–1914. *Journal of Political Economy* 130(5), 1249–1318.
- Fally, T. and J. Sayre (2018). Commodity Trade Matters. *NBER Working Paper No. 24965*.
- Farrokhi, F. and H. S. Pellegrina (2023). Trade, Technology, and Agricultural Productivity. *Journal of Political Economy*.
- Fischer, G., F. O. Nachtergaele, H. van Velthuisen, F. Chiozza, G. Francheschini, M. Henry, D. Muchoney, and S. Tramberend (2021). Global Agro-ecological Zones (GAEZ v4)-Model Documentation.
- Gáfaró, M. and H. S. Pellegrina (2022). Trade, farmers’ heterogeneity, and agricultural productivity: Evidence from Colombia. *Journal of International Economics* 137.
- Ganapati, S., W. F. Wong, and O. Ziv (2021). Entrepôt: Hubs, scale, and trade costs. *NBER Working Paper No. 29015*.
- Gollin, D., D. Lagakos, and M. E. Waugh (2014). Agricultural productivity differences across countries. *American Economic Review* 104(5), 165–70.
- Gouel, C. (2025). Measuring climate change impacts on agriculture: An equilibrium perspective on supply-side approaches. *Journal of the Association of Environmental and Resource Economists* 12(1), 181–220.
- Gouel, C. and D. Laborde (2021). The crucial role of domestic and international market-mediated adaptation to climate change. *Journal of Environmental Economics and Management* 106, 102408.
- Grant, M. and M. Startz (2025). Cutting out the middleman: The structure of chains of intermediation. *Working paper*.
- Grossman, G. M. and E. Rossi-Hansberg (2010). External economies and international trade redux. *The Quarterly Journal of Economics* 125(2), 829–858.
- Hart, W. E., J.-P. Watson, and D. L. Woodruff (2011). Pyomo: Modeling and solving mathematical programs in Python. *Mathematical Programming Computation* 3(3), 219–260.
- Hijmans, R. J., L. Guarino, M. Cruz, and E. Rojas (2001). Computer tools for spatial analysis of plant genetic resources data: 1. diva-gis. *Plant genetic resources newsletter*, 15–19.
- Jung, J., J. Sesmero, and R. Siebert (2021). A structural estimation of spatial differentiation and market power in input procurement. *American Journal of Agricultural Economics*.
- Krugman, P. (1979). Increasing returns, monopolistic competition, and international trade. *Journal of International Economics* 9(4), 469–479.
- Krugman, P. (1991). Increasing returns and economic geography. *Journal of Political Economy* 99(3), 483–499.
- Kucheryavyy, K., G. Lyn, and A. Rodríguez-Clare (2016). Grounded by gravity: A well-behaved trade model with industry-level economies of scale. *NBER Working Paper No. 22484*.
- Kucheryavyy, K., G. Lyn, and A. Rodríguez-Clare (2021). Spatial equilibria: The case of two regions.

- Lence, S. H. (2016). Modeling the Market and Welfare Effects of Mexico's "Agriculture by Contract" Program. *American Journal of Agricultural Economics* 98(3), 925–945.
- Lim, S. (2023). Global agricultural value chains and structural transformation. *Risks in Agricultural Supply Chains, National Bureau of Economic Research*.
- Lin, J. and F. Rauch (2020). What future for history dependence in spatial economics? *Regional Science and Urban Economics*, 103628.
- Marshall, A. (1890). *Principles of Economics*. London: MacMillan.
- Meemken, E.-M. and M. F. Bellemare (2020). Smallholder farmers and contract farming in developing countries. *Proceedings of the National Academy of Sciences* 117(1), 259–264.
- Nath, I. (2025). Climate change, the food problem, and the challenge of adaptation through sectoral reallocation. *Journal of Political Economy*.
- Nuhu, A. S., L. S. O. Liverpool-Tasie, T. Awokuse, and S. Kabwe (2021). Do benefits of expanded midstream activities in crop value chains accrue to smallholder farmers? Evidence from Zambia. *World Development* 143.
- Olver, R. and D. Zilberman (2022). Why soil fumigation changed the strawberry industry. *ARE Update, University of California Giannini Foundation of Agricultural Economics* 25(3), 5–8.
- Pellegrina, H. S. (2022). Trade, productivity, and the spatial organization of agriculture: Evidence from Brazil. *Journal of Development Economics* 156, 102816.
- Rao, P. (1993). Intercountry comparisons of agricultural output and productivity. *FAO Economic and Social Development Paper* (112).
- Rattalino Edreira, J. I., J. F. Andrade, K. G. Cassman, M. K. van Ittersum, M. P. van Loon, and P. Grassini (2021). Spatial frameworks for robust estimation of yield gaps. *Nature Food* 2(10), 773–779.
- Reardon, T., R. Echeverria, J. Berdegue, B. Minten, S. Liverpool-Tasie, D. Tschirley, and D. Zilberman (2019). Rapid transformation of food systems in developing regions: Highlighting the role of agricultural research and innovations. *Agricultural Systems* 172, 47–59.
- Rivera-Padilla, A. (2020). Crop choice, trade costs, and agricultural productivity. *Journal of Development Economics* 146.
- Roberts, M. J. and W. Schlenker (2013). Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate. *American Economic Review* 103(6), 2265–95.
- Sotelo, S. (2020). Domestic trade frictions and agriculture. *Journal of Political Economy* 128(7), 2690–2738.
- Tombe, T. (2015). The missing food problem: Trade, agriculture, and international productivity differences. *American Economic Journal: Macroeconomics* 7(3), 226–58.
- Wantchekon, L. and P. Stanig (2015). The curse of good soil? Land fertility, roads, and rural poverty in Africa. *Unpublished manuscript*, 1–65.
- Zavala, L. (2025). Unfair Trade? Market Power in Agricultural Value Chains. *Working Paper*.
- Zilberman, D., L. Lu, and T. Reardon (2019). Innovation-induced food supply chain design. *Food Policy* 83, 289–297.

A Appendix

A.1 Mathematical Derivations

Extensions to allowing labor in the production function of packing firms

To include labor in the firm's production function, I can write their production function as follows:

$$y_{ivk}(q_{ivk}) = A_{ivk}(q_{ivk}^\rho l_{ivk}^{1-\rho})$$

A firm v operating in region i which produces crop k has the following PMP:

$$\max_{q_{ivk}} \pi_{ivk} = \max_{q_{ivk}} \tilde{p}_{ik1} y_{ivk}(q_{ivk}) - q_{ivk} p_{ivk}(q_{ivk}, q_{-ivk}, q_{ik}, Q_i) - w_i l_{ivk} - f_k. \quad (24)$$

Substituting the firms' optimal demand for labor (from the canonical properties of Cobb Douglas production functions, the total wage bill of the firm will be $(1-\rho)\%$ of revenues) into its profits, I obtain that the firm's profit maximization can be written as:

$$\max_{q_{ivk}} \rho \left[\frac{(1-\rho)}{w_i} \right]^{(1-\rho)/\rho} A_{ivk}^{1/\rho} \tilde{p}_{ik1}^{1/\rho} q_{ivk} - q_{ivk} p_{ivk}(q_{ivk}, q_{-ivk}) - f_k = \max_{q_{ivk}} \tilde{A}_{ivk} q_{ivk} - q_{ivk} p_{ivk}(q_{ivk}, q_{-ivk}) - f_k, \quad (25)$$

where I define $\tilde{A}_{ivk} \equiv \rho \left[\frac{(1-\rho)}{w_i} \right]^{(1-\rho)/\rho} A_{ivk}^{1/\rho} \tilde{p}_{ik1}^{1/\rho}$ (noting that everything in $\tilde{A}_{ivk}$ is taken to be exogenous with respect to the firm v). In this case, I assume that packing plants only represent a small share of labor employment compared to manufacturing employment and so take the wage to be given, and equal to the wage in the manufacturing sector.

Allowing for multiple entrants

I can allow for extensions of the model where I allow for multiple packing firms in a region. In each region i , there are $n_{ik} \in \mathbb{N}_0$ processing firms producing crop k , where if no firms choose to enter, $n_{ik} = 0$. The firms act as Cournot competitors, and so choose optimal quantities, taking their competitors' actions as given. Each firm draws from the same distribution for A_{vk} , and chooses to enter if it's calculated hypothetical operating profits are greater than its fixed costs (which depends on its markdown). If I make assumptions about the distribution of A_{vk} in the future, I can characterize the number of firms n_{ik} in equilibrium, but for now I can simulate this analytically. Here, I will solve for the optimal markdown based upon Cournot competition between the n_{ik} entrants in a region.

Recalling that farms can only sell their high quality output to processors then for supply to equal demand we must have:

$$\sum_{j \in R(i)} q_{jk1} = \sum_{v=1}^{n_{ik}} q_{ivk},$$

For notational simplicity here, I assume that the firms only source from farms in their same region i , or

$$q_{ik1} = \sum_{v=1}^{n_{ik}} q_{ivk}.$$

Taking the first order condition with respect to q_{ivk} yields the markdown condition that determines the wedge between prices received by intermediaries in region i , $\tilde{p}_{ik1}$ and the marginal cost (factory-gate price) of a unit of a crop with high quality, p_{ivk} (net of productivity) as:

$$\frac{\tilde{p}_{ik1}}{p_{ivk}} = \frac{\left(1 + \frac{1}{\varepsilon_{vk}}\right)}{\tilde{A}_{ivk}} \quad (26)$$

where $\frac{1}{\varepsilon_{vk}} \equiv \frac{\partial \log p_{ik1}}{\partial \log q_{ivk}}$ ⁴⁵. For simplicity, note that $p_{ivk} = p_{ik1}$, and so $\frac{1}{\varepsilon_{vk}} = \frac{\partial \log p_{ik1}}{\partial \log q_{ivk}}$.

But we have that

$$\frac{\partial \log p_{ik1}}{\partial \log q_{ivk}} = \frac{\partial \log p_{ik1}}{\partial \log q_{ik1}} \times \frac{\partial \log q_{ik1}}{\partial \log q_{ivk}} \quad (27)$$

⁴⁵Note that $q_{ivk} = \sum_{j \in R(i(v))} q_{jk1}$

To begin, note that

$$\frac{\partial \log q_{ik1}}{\partial \log q_{ivk}} = \frac{q_{ivk}}{q_{ik1}},$$

which is the share that of crop k with high quality h firm v purchases from region i (out of all purchasing firms v). Therefore, larger firms (i.e. those with larger A_{vk} 's) have larger markdowns. The remaining term is the standard inverse supply elasticity for crop k with export-eligibility $e = 1$ in region i .

Regional crop supply By competitive factor pricing, the revenue share of land for crop k with eligibility e in region i is γ_{ke} , so $\gamma_{ke} p_{ike} q_{ike}$ equals land payments for that tranche. Since the rental rate of an efficiency unit of land is P_i , total land income is $P_i H_i$, and the share of the rental rate for crop k with eligibility e in region i is thus $\eta_{ike} P_i H_i$. Hence

$$\gamma_{ke} p_{ike} q_{ike} = \eta_{ike} P_i H_i \Rightarrow q_{ike} = \eta_{ike} H_i \frac{P_i}{\gamma_{ke} p_{ike}} \quad (28)$$

Plugging in for η_{ike} from equation 6, I obtain

$$q_{ike} = \gamma_{ike}^{-1} P_i^{1-\vartheta} \lambda_{ike}^\vartheta A_{ik}^\vartheta p_{ik}^{\vartheta-\theta} H_i p_{ike}^{-1}. \quad (29)$$

Next, define $b_{ike} \equiv w_{iA}^{-\theta \alpha_{ke} / \gamma_{ke}} p_{ix_k}^{-\theta \beta_{ke} / \gamma_{ke}}$, so equation 7 can be rewritten as

$$\lambda_{ike} \equiv p_{ike} \left(\frac{w_{iA}}{p_{ike}} \right)^{-\alpha_{ke} / \gamma_{ke}} \left(\frac{p_{ix_k}}{p_{ike}} \right)^{-\beta_{ke} / \gamma_{ke}} = \underbrace{w_{iA}^{-\alpha_{ke} / \gamma_{ke}} p_{ix_k}^{-\beta_{ke} / \gamma_{ke}}}_{b_{ike}^{1/\theta}} p_{ike}^{1/\gamma_{ke}}. \quad (30)$$

Continuing, define $\Xi_{ik} \equiv A_{ik}^\vartheta p_{ik}^{\vartheta-\theta} P_i^{1-\vartheta} H_i$, so that equation 29 can be written as

$$q_{ike} = \frac{1}{\gamma_{ke}} b_{ike} p_{ike}^{\theta/\gamma_{ke}-1} \Xi_{ik}. \quad (31)$$

Define the sum

$$S_{ik} \equiv \sum_{e \in \{0,1\}} \frac{b_{ike}}{\gamma_{ke}} p_{ike}^{\theta/\gamma_{ke}-1}. \quad (32)$$

so that summing 31 over e gives the crop-level quantity

$$q_{ik} = \Xi_{ik} S_{ik}. \quad (33)$$

Derivation of elasticity of regional crop supply

To obtain the inverse supply function, from (??) we can invert directly given the upper-tier term Ξ_{ik} :

$$p_{ike} = \left(\frac{\gamma_{ke} q_{ike}}{b_{ike} \Xi_{ik}} \right)^{\gamma_{ke} / (\theta - \gamma_{ke})}, \quad (34)$$

Eliminating Ξ_{ik} via equation 33 yields a form depending only on q_{ike} , the crop total q_{ik} , and the composition term S_{ik} :

$$p_{ike} = \gamma_{ke}^{\frac{\gamma_{ke}}{\theta - \gamma_{ke}}} b_{ike}^{\frac{\gamma_{ke}}{\gamma_{ke} - \theta}} q_{ike}^{\frac{\gamma_{ke}}{\theta - \gamma_{ke}}} q_{ik}^{-\frac{\gamma_{ke}}{\theta - \gamma_{ke}}} S_{ik}^{\frac{\gamma_{ke}}{\theta - \gamma_{ke}}}. \quad (35)$$

Taking logs of (35) and ignoring the b_{ike} and γ_{ke} terms,

$$\log p_{ike} = \underbrace{\frac{\gamma_{ke}}{\theta - \gamma_{ke}}}_{>0} \log q_{ike} + \underbrace{\frac{\gamma_{ke}}{\gamma_{ke} - \theta}}_{<0} \log q_{ik} - \underbrace{\frac{\gamma_{ke}}{\gamma_{ke} - \theta}}_{<0} \log S_{ik} + \text{constant}(b_{ike}, \gamma_{ke}, \theta). \quad (36)$$

Before proceeding, I define the share of production of crop k with export eligibility e in region i to be

$$s_{ike} \equiv \frac{q_{ike}}{q_{ik}} = \frac{b_{ike}^{\theta/\gamma_{ke}-1} p_{ike}}{S_{ik}}. \quad (37)$$

Then, the own inverse-supply elasticity holding the crop total q_{ik} fixed is:

$$\left. \frac{\partial \log p_{ike}}{\partial \log q_{ike}} \right|_{q_{ik}} = \frac{\gamma_{ke}}{(\theta - \gamma_{ke})(1 - s_{ike})}. \quad (38)$$

When s_{ike} is small (for instance when there is no production of $e = 1$ before firms enter), ?? reduces to

$$\frac{\partial \log p_{ike}}{\partial \log q_{ike}} \approx \frac{\gamma_{ke}}{(\theta - \gamma_{ke})}. \quad (39)$$

Land and revenue shares Define the share of revenue of region i from crop k with export eligibility e as the following: $\pi_{ike} \equiv \frac{V_{ike}}{V_i} \equiv \frac{p_{ike}q_{ike}}{\sum_{k \in \mathcal{K}} \sum_{e \in \mathcal{E}} p_{ike}q_{ike}}$. Then, the share of land allocated to each crop k with status e within region i and the share of revenue from the same triplet can be related to the revenue share by

$$\eta_{ike} = \frac{\gamma_{ike}\pi_{ike}}{\sum_{k \in \mathcal{K}} \sum_{e \in \mathcal{E}} \gamma_{ike}\pi_{ike}}. \quad (40)$$

To see this, note that equation 8 tells us that $V_{ike} = p_{ike}q_{ike} = \gamma_{ike}^{-1} P_i \eta_{ike} H_i$. Using this and the definition $\pi_{ike} \equiv \frac{V_{ike}}{V_i}$, the right side of equation 40 becomes $\frac{\eta_{ike}}{\sum_{k \in \mathcal{K}} \sum_{e \in \mathcal{E}} \eta_{ike}}$, which reduces down to η_{ike} , proving the identity.

A.2 Additional derivations of model details

Preferences

The representative consumer consumes two main aggregates in the upper tier: agricultural goods and manufactured goods. They have preferences over these aggregates which are Cobb-Douglas with shares ζ for agricultural consumption and $(1 - \zeta)$ manufacturing consumption, with $\zeta \in [0, 1]$. In the lower tiers, consumers consume agriculture as a constant elasticity aggregate given by

$$C_{i,A} = \left(\sum_{k=1}^K a_k^{1/\sigma_A} C_{i,k}^{\frac{\sigma_A-1}{\sigma_A}} \right)^{\frac{\sigma_A}{\sigma_A-1}}, \quad (41)$$

where $\sigma_A > 0$ is the elasticity of substitution across crops, and $\sum_{k=1}^K a_k = 1$, with $a_k > 0$. With this assumption, the CES price index for agriculture is⁴⁶

$$\tilde{P}_{i,A} = \left(\sum_{k=1}^K a_k^{-1} \tilde{p}_{ik}^{1-\sigma_A} \right)^{\frac{1}{1-\sigma_A}}. \quad (42)$$

Above, I use $\tilde{p}_{ik}$ to represent the prices paid for crop k by consumers, to differentiate these prices from those received by farmers. In a similar fashion, consumers have a taste for variety in manufacturing goods, which are differentiated by origin, with an elasticity of substitution across varieties given by $\sigma_M > 0$.

To differentiate each crop k , consumers choose whether to purchase their crop either in a local market or from part of the agricultural value chain such as a supermarket, where

$$C_{i,k} = \left(a_{k0}^{1/\sigma_e} C_{ik0}^{\frac{\sigma_e-1}{\sigma_e}} + a_{k1}^{1/\sigma_e} C_{ik1}^{\frac{\sigma_e-1}{\sigma_e}} \right)^{\frac{\sigma_e}{\sigma_e-1}}, \quad (43)$$

and $\sigma_e > 0$ is the elasticity of substitution across crops with varying phytosanitary standards, and $a_{k0} + a_{k1} = 1$, with

⁴⁶Combining these assumptions, this implies that the share of income spent on crop k is given by $\xi_{i,k} = \frac{a_k^{-1} \tilde{p}_{i,k}^{1-\sigma_A}}{\tilde{P}_{i,A}^{1-\sigma_A}} \times \zeta$ for $k \in \mathcal{K}$ and the share of income spent on manufacturing goods from origin i' is $\xi_{ii',M} = \left(\tau_{ii',M} \frac{w_{i'M}}{A_{i'M}} \right)^{1-\sigma_M} \tilde{P}_{i,M}^{\sigma_M-1} (1 - \zeta)$, with $\tilde{P}_{i,M} = \left[\sum_{n \in \mathcal{W}} \left(\tau_{in,M} \frac{w_{nM}}{A_{nM}} \right)^{1-\sigma_M} \right]^{\frac{1}{1-\sigma_M}}$.

$a_{ke} > 0$. The corresponding price index for crop k , in turn, is given by⁴⁷:

$$\tilde{p}_{ik} = \left(a_{k0}^{-1} \tilde{p}_{ik0}^{1-\sigma_e} + a_{k1}^{-1} \tilde{p}_{ik1}^{1-\sigma_e} \right)^{\frac{1}{1-\sigma_e}}. \quad (44)$$

A.3 Equilibrium conditions and counterfactual equations

Competitive Equilibrium The game is played dynamically for periods $t \in [1, T]$ ⁴⁸. In period $t = 1$ it is assumed that the vector of foreign prices for all crops k , $\{p_{Wk}\}_{k \in \mathcal{K}}$, is too low to support the entry of any packing plants in the domestic economy. Then, conditional on the non-existence of packing plants, the equilibrium is achieved via the solution described above. In subsequent periods, firms decide to enter based upon the exogenously determined price paths for crops, and set their markups based on the information available to them in that period, m_{ik} . After the firm entry decisions have been made, then the equilibrium is solved for. In future periods, the game repeats in the same way.

A competitive equilibrium in each period t consists of, for each region $i \in \mathcal{W}$:

1. final prices $\tilde{p}_{ike}$ for all crops k , and farmgate prices p_{ike} (note that $p_{ik0} = \tilde{p}_{ik0}$, and $p_{ik1} = \frac{\tilde{p}_{ik1}}{m_{ik}} \times \delta_{ik}$). Note that if an agointermediary is not present ($\delta_{ik} = 0$) in region i and crop k , then the farmgate price for high phytosanitary varieties is zero, $p_{ik1} = 0$.
2. wages in each region w_i
3. final consumption C_{ike} for all crops k with quality e and final goods expenditure $E_{ni,M}$ in manufacturing
4. for the representative farmer, input demands $l_{ike, x_{ik}}$ and output q_{ike} for all crops k with qualities e
5. for processing plants, q_{vk} and output y_{vk} for all crops k
6. for the manufacturing sector, input demands l_{iM}
7. trade flows:
 - Domestic trade flows $z_{ni,g}$ for all regions $i, n \in \mathcal{M}$ for all goods $g \in \mathcal{G}$
 - International trade flows with ROW: $z_{Fi,g}$ and $z_{iF,g}$ for all regions $i \in \mathcal{M}$ and $g \in \mathcal{G}$ (where $z_{Fi, x_k} = 0$ for all intermediate inputs into crops k)
 - International trade flows with the US: $z_{US,i,g}$ and $z_{i,US,g}$ for all regions $i \in \mathcal{M}$ and $g \in \mathcal{G}$ (where $z_{US,i, x_k} = 0$ for all intermediate inputs into crops k)
8. The matrix $\delta_{|\mathcal{M}| \times |\mathcal{K}|}$ with binary entries $\in \{0, 1\}$ which represents whether a packing firm is operating in region $i \in \mathcal{M}$ and producing crop $k \in \mathcal{K}$ and corresponding firm markups m_{ik}

such that

- a) The quantities in (3) solve the consumer's problem, given income and prices

$$C_{ike} = \frac{\xi_{ike} E_i}{\tilde{p}_{ike}} = (a_{ike} a_{i,k})^{-1} \tilde{p}_{ike}^{-\sigma_e} \tilde{p}_{i,k}^{\sigma_e - \sigma_A} \tilde{P}_{i,A}^{\sigma_A - 1} \times \zeta \times E_i. \quad (45)$$

$$E_{ii',M} = \left(\tau_{ii',M} \frac{w_{i',M}}{A_{i',M}} + \tau_{ii',M} \pi_{ii',M} \right)^{1-\sigma_M} \tilde{P}_{i,M}^{\sigma_M - 1} (1 - \zeta) E_i = \lambda_{ii',M} (1 - \zeta) E_i,$$

$$\text{with } \tilde{P}_{i,M} = \left[\sum_{n \in \mathcal{W}} \left(\tau_{in,M} \frac{w_{nM}}{A_{nM}} + \tau_{in,M} \pi_{in,M} \right)^{1-\sigma_M} \right]^{\frac{1}{1-\sigma_M}}.$$

- b) The inputs and outputs in 4) solve the representative farmer's problem, given prices;

⁴⁷With this assumption, the share of expenditure in region i devoted to a crop k with low ($e = 0$) or high ($e = 1$) phytosanitary standards is given by $\xi_{ke} = (a_{ke} a_k)^{-1} \tilde{p}_{ike}^{1-\sigma_e} \tilde{p}_{i,k}^{\sigma_e - \sigma_A} \tilde{P}_{i,A}^{\sigma_A - 1} \times \zeta$.

⁴⁸Since the equilibrium is solved in each period as it would in a static model, I suppress the t notation in the main text.

- c) The inputs and outputs in 5) solve the representative downstream firm's problem, given prices;
- d) The no-arbitrage conditions hold, or $\tau_{ni,g} (p_{ig} + \pi_{ni,g}) \geq p_{ng} \perp z_{ni,g} \forall n, i \in \mathcal{W}$ and $\forall g \in \mathcal{G}$
- e) The labor demand in urban regions $i \in \mathcal{U}$ in manufacturing is given by:

$$w_{iM} L_{iM} = \sum_{n \in \mathcal{M}} \left(\tau_{ni,M} \frac{w_{iM}}{A_{iM}} + \tau_{ni,M} \pi_{ni,M} \right)^{1-\sigma_M} \tilde{P}_{nM}^{\sigma_M-1} E_{nM} +$$

$$\sum_{n \in \{F, US\}} \left(\tau_{ni,M} \frac{w_{iM}}{A_{iM}} + \tau_{ni,M} \pi_{ni,M} \right)^{1-\sigma_M} \tilde{P}_{nM}^{\sigma_M-1} X_{n,MEX,M} = \lambda_{ni,M} E_{nM} + \sum_{n \in \{F, US\}} \lambda_{ni,M} X_{n,MEX,M}$$

- F) The labor demand in rural regions $i \in \mathcal{R}$ in agriculture is given by⁴⁹:

$$w_{iA} L_{iA} = \sum_k (\alpha_{ik0} p_{ik0} q_{ik0} + \alpha_{ik1} p_{ik1} q_{ik1}).$$

- g) In all urban regions $i \in \mathcal{U}$, local markets clear for labor and crops for final consumption (both for low and high quality varieties – low type crops are only consumed from outlying areas, high type crops can be consumed from AGVCs anywhere in the country or from abroad), and the intermediate market clears for high type crops.

- $L_i = L_{iM} + L_{iP} = L_{iM} + \sum_{k \in \mathcal{K}} l_{vk}$
- $E_{ik0} = \tilde{p}_{ik0} C_{ik0} = \tilde{p}_{ik0} \sum_{i' \in \mathcal{R}(i)} z_{ii',k,l}, \forall k \in \mathcal{K}$
(low phytosanitary standard crop consumption must come from the local rural periphery)
- $p_{vk} q_{vk} = p_{vk} \sum_{i \in \mathcal{R}(n)} q_{ik1}, \forall k \in \mathcal{K}$
(processing crop inputs must equal supply from outlying regions)
- $E_{ik1} = \tilde{p}_{ik1} C_{ik1} = \tilde{p}_{ik1} (y_{vk} - \sum_{n \in \mathcal{W}} z_{ni,k,h} + \sum_{i' \in \mathcal{W}/\mathcal{R}} z_{ii',k,h}), \forall k \in \mathcal{K}$
(expenditure on crop k with high phytosanitary standards (h) must be equal to the value of the urban region's production from packing plants, less exports to the world, plus imports from the world)

- h) In all rural regions $i \in \mathcal{R}$, local markets clear for labor and crops for final consumption

- $L_i = L_{iA} = \sum_{k \in \mathcal{K}} l_{vk} = \sum_{k \in \mathcal{K}} \int_{\ell \in \Omega_{ik}} l_{vk}(\ell) d\ell$
(labor is fully employed by agriculture)
- $E_{ik0} = \tilde{p}_{ik0} C_{ik0} = \tilde{p}_{ik0} (q_{ik0} - \sum_{n \in \mathcal{U}(i)} z_{ni,k,l}), \forall k \in \mathcal{K}$
(low phytosanitary standard crop consumption must come from self production, less what is sold in the urban market)
- $E_{ik1} = \tilde{p}_{ik1} C_{ik1} = \tilde{p}_{ik1} (\sum_{i' \in \mathcal{W}/\mathcal{R}} z_{ii',k,h}), \forall k \in \mathcal{K}$ (high phytosanitary standard crop consumption must come from urban regions or abroad)

- i) The representative consumer's expenditure, equals the household's income from all sources:

- For urban regions $i \in \mathcal{U}$: $E_i = w_{iM} L_{iM} + \sum_{k \in \mathcal{K}} (\pi_{i(v)k}^{op} - f_k)$
- For rural regions $i \in \mathcal{R}$: $E_i = w_{iA} L_{iA} + \int_{\Omega_i} r_i(\ell) d\ell$

Conditions g)-i) imply that trade is balanced between Mexico, the US, and Foreign.

⁴⁹Here, I define the price the processing plant offers farms in rural regions to be p_{ik1} for $i \in \mathcal{R}$. Taking a given rural region $i \in \mathcal{R}$ and the no arbitrage condition in d) yields that $p_{ik1} = \frac{p_{U(i),v,k}}{\tau_{U(i),i,k}} - \pi_{U(i),i,k}$.

A.4 Appendix: Generic uniqueness of the entry path

Next, I move to proving the generic uniqueness of the dynamic path. Let $\mathcal{V} \equiv \{(i, k)\}$ denote the finite set of potential plants (one crop k per urban region i). Time is continuous, $t \in [0, \infty)$. Let $S \subseteq \mathcal{M}$ be the set of plants that have already entered by time t . For each $v \in \mathcal{V}$ and state S , let $p_v(\cdot; S)$ denote the residual farmgate inverse supply faced by v (aggregated over its sourcing municipalities) and let

$$C_v(q; S) \equiv \int_0^q p_v(u; S) du$$

be the induced (convex) sourcing cost. Let $\tilde{p}_v(t) \equiv \tilde{p}_{ik1}(t) = p_{wk}(t)/\tau_{wik}$ be the exogenous price faced by v and $A_{vk} > 0$ its productivity as in 10. The plant's *operating* profit at time t in state S is the value of the static problem

$$\Pi_v^{\text{op}}(t; S) \equiv \max_{q \geq 0} \left\{ A_v \tilde{p}_v(t) q - C_v(q; S) \right\} - f_k,$$

which is well defined by the static uniqueness result in the text.⁵⁰ v is *viable* at (t, S) if $\Pi_v^{\text{op}}(t; S) \geq 0$.

To make progress on generic uniqueness, I make the following assumptions:

- (A0) **Static uniqueness.** For every $S \subseteq \mathcal{V}$, the static block has a unique equilibrium that yields continuous, strictly increasing residual inverse supply $p_v(\cdot; S)$ for each v .
- (A1) **Price path.** For each crop k , $p_{wk}(\cdot)$ is continuous and strictly increasing, and initially too low for any entry: $\Pi_v^{\text{op}}(0; \emptyset) < 0$ for all v .
- (A2) **Generic no ties.** For any state S realized along the entry path, the set of *hitting times*

$$T_v(S) \equiv \inf\{t \geq 0 : \Pi_v^{\text{op}}(t; S) \geq 0\}$$

are distinct across $v \notin S$. Equivalently, equalities $T_v(S) = T_{v'}(S)$ occur on a measure-zero set of primitives $(A_v, f_k, \tau_{wik}, \dots)$ and are excluded.

Entry tightens residual supply via (a) higher agricultural wages w_{iA} and (b) higher $p_v(\cdot; S)$ through the composition terms in the inverse supply. Both raise $C_v(\cdot; S)$.

Lemma A1 (Residual supply shrinks when another plant enters). Fix a potential plant v and a set of incumbents $S \subseteq \mathcal{V}$. Let p_v be a posted factory-gate price and let $\phi_v^j(p_v)$ denote the delivered farmgate price in municipality j offered by v (e.g. $\phi_v^j(p_v) = (p_v - \pi_{ij,k})/\tau_{ij,k}$). Let $q_{jk1}(\cdot)$ be the strictly increasing local supply of the high-quality crop k in j . Define v 's residual supply at price p_v as

$$q_v(p_v; S) = \sum_{j \in R(i)} \mathbf{1}\left\{ \phi_v^j(p_v) \geq \max_{s \in S} \phi_s^j(p_s) \right\} q_{jk1}(\phi_v^j(p_v)).$$

Then for any additional entrant $v' \notin S$ we have

$$q_v(p_v; S \cup \{v'\}) \leq q_v(p_v; S) \quad \text{for all } p_v \geq 0.$$

Consequently, the residual inverse supply is weakly higher,

$$p_v(q; S \cup \{v'\}) \geq p_v(q; S) \quad \text{for all } q \geq 0,$$

and the sourcing cost $C_v(q; S) = \int_0^q p_v(u; S) du$ is nondecreasing in S .

Proof. Fix p_v . Adding v' augments the set of delivered prices available to farmers. For each municipality j , the indicator $\mathbf{1}\{\phi_v^j(p_v) \geq \max_{s \in S} \phi_s^j(p_s)\}$ can only weakly fall (if v' offers a higher delivered price in j). Since $q_{jk1}(\cdot)$ is increasing, each term weakly falls, hence $q_v(p_v; S \cup \{v'\}) \leq q_v(p_v; S)$ pointwise in p_v . Because $q_v(\cdot; S)$ is increasing, its right inverse is order-reversing, yielding $p_v(\cdot; S \cup \{v'\}) \geq p_v(\cdot; S)$, and integrating gives $C_v(q; S \cup \{v'\}) \geq C_v(q; S)$.

⁵⁰Static uniqueness delivers well-defined residual supply schedules $p_v(\cdot; S)$ for every S and hence well-defined $C_v(\cdot; S)$.

Entry can also raise local wages w_{jA} and input prices p_{jx_k} and reallocate land via the two-tier CET; from the inverse-supply representation, these channels further shift up p_{jk1} at given quantities.

Lemma 2 (Profit increases in time). Under (A0)-(A1), for any fixed S , the map $t \mapsto \Pi_v^{\text{op}}(t; S)$ is continuous and strictly increasing.

Proof. For fixed S , $\Pi_v^{\text{op}}(t; S) = \phi_v(A_v \tilde{p}_v(t); S) - f_k$ where $\phi_v(r; S) \equiv \sup_{q \geq 0} \{rq - C_v(q; S)\}$ is the convex conjugate of $C_v(\cdot; S)$. Since $C_v(\cdot; S)$ is convex and not identically $+\infty$, $\phi_v(\cdot; S)$ is finite, continuous and strictly increasing. Because $\tilde{p}_v(\cdot)$ is strictly increasing and continuous by (A1), the composition is strictly increasing and continuous.

Lemma 3 (Profit decreases in the set of entrants). Under (A0)-(A1), for any fixed t , if $S' \supseteq S$ then $\Pi_v^{\text{op}}(t; S') \leq \Pi_v^{\text{op}}(t; S)$.

Proof. $C_v(\cdot; S') \geq C_v(\cdot; S)$ pointwise implies $\phi_v(r; S') \leq \phi_v(r; S)$ for all r . Thus, $\Pi_v^{\text{op}}(t; S') \leq \Pi_v^{\text{op}}(t; S)$.

Thresholds and their monotonicity. For $S \subseteq \mathcal{V}$, define the *entry threshold* $T_v(S)$ as above. By Lemma 2 and (A1) each $T_v(S)$ is well defined and unique. By Lemma 3, $T_v(S)$ is nondecreasing in S : if $S' \supseteq S$ then $T_v(S') \geq T_v(S)$. With more entrants, the sourcing costs shift up, which delays the break-even threshold of all other potential entrants.

Recursive construction of the path. Start at $S_0 = \emptyset$. Define recursively

$$t_j \equiv \min_{v \notin S_{j-1}} T_v(S_{j-1}), \quad v_j \equiv \arg \min_{v \notin S_{j-1}} T_v(S_{j-1}), \quad S_j \equiv S_{j-1} \cup \{v_j\}.$$

By (A3) the arg min is a singleton at each step. Let $\{(v_j, t_j)\}_{j=1}^J$ be the resulting sequence until no further v has finite threshold (or until the finite set $\mathcal{V}$ is exhausted).

Theorem Under (A0)-(A3), the sequence $\{(v_j, t_j)\}_{j=1}^J$ constructed above is the *unique* entry path: (i) for each j , plant v_j is the only plant that can enter at time t_j ; (ii) no plant enters before its assigned time; and (iii) any other purported path must coincide with $\{(v_j, t_j)\}_{j=1}^J$.

Proof. I argue by induction.

Base step ($j = 1$): By Lemma 2 each $\Pi_v^{\text{op}}(t; \emptyset)$ is strictly increasing in t and negative at $t = 0$, so each $T_v(\emptyset)$ exists and is unique. By definition, $t_1 = \min_v T_v(\emptyset)$ and by (A3) there is a unique minimizer v_1 . For $t < t_1$ we have $\Pi_v^{\text{op}}(t; \emptyset) < 0$ for all v , so no entry can occur; at $t = t_1$ only v_1 is viable.

Inductive step: Suppose the statement holds for the first $j - 1$ entrants, yielding state S_{j-1} . Consider the thresholds $\{T_v(S_{j-1})\}_{v \notin S_{j-1}}$. By Lemma 2 each is unique; by (A3) the minimum t_j is achieved by a unique v_j . For $t < t_j$ and any $v \notin S_{j-1}$, $\Pi_v^{\text{op}}(t; S_{j-1}) < 0$, so entry before t_j is impossible. At t_j only v_j is viable. After v_j enters, the state becomes S_j . By Lemma 3, for any remaining v and any $t \geq t_j$, $\Pi_v^{\text{op}}(t; S_j) \leq \Pi_v^{\text{op}}(t; S_{j-1})$, hence their thresholds weakly shift to the right: $T_v(S_j) \geq T_v(S_{j-1}) > t_j$. Thus no two plants enter at t_j . Iterating completes the construction. Uniqueness follows because at each step the next entrant and time are uniquely pinned down by the (state-dependent) minimal threshold.

A.5 Counterfactual analysis

To study the effects of changes to the agricultural productivity of regions, as well as changes the fixed costs of entry for downstream processors, I write the main equilibrium equations of my model in terms of changes, rather than in levels. Using the exact hat/calibrated share form notation (i.e. $\hat{X} = X'/X$, where X' is the new outcome and X is the initial outcome), these shocks can be written as $\{\hat{a}_{ik}, \hat{a}_{ike}\}$, $\{\hat{F}_k\}$, $\{\hat{\tau}_{ni,g}, \hat{\pi}_{ni,g}\}$, and $\{\hat{\delta}_{ik}\}$. This exercise requires a number of moments in data, as well as structural parameters. These are:

- Initial moments $\mathcal{D} \equiv \{\xi_{ike}(\tilde{\mathbf{p}}_i), \xi_{ik}(\tilde{\mathbf{p}}_i), \eta_{ie|k}, \eta_{ik}, \lambda_{in,M}\}$,
- the parameters which govern supply $\Delta_S \equiv \{\theta, \vartheta, \alpha_{ike}, \beta_{ike}, \gamma_{ike}, m_{vk}, A_{vk}, A_{ik}, A_{iM}\}$,
- parameters which govern demand $\Delta_D \equiv \{\sigma_A, \sigma_M, \sigma_e, \zeta, \{a_k\}_{k \in \mathcal{K}}, \{a_{ke}\}_{k \in \mathcal{K}, e \in \mathcal{E}}\}$
- endowments of labor and land $\mathcal{E} \equiv \{L_i, H_i\}$,
- and exogenously determined prices and demand from abroad
 $\Delta_X \equiv \{\{\tilde{p}_{Fg}\}_{g \in \mathcal{G}}, \{p_{Fx_k}\}_{k \in \mathcal{K}}, X_{n \in \{F, US\}, MEX, M}\}$.

The following counterfactual equations will determine the new outcomes:

$$1. \quad \hat{q}_{ike} = \hat{\eta}_{ike} \times \hat{P}_i \times \hat{p}_{ike}^{-1}.$$

2. $\widehat{C}_{ike} = (\widehat{a}_{ike}\widehat{a}_{i,k})^{-1}\widehat{p}_{ike}^{-\sigma_e}\widehat{p}_{i,k}^{\sigma_e-\sigma_A}\widehat{P}_{i,A}^{\sigma_A-1} \times \widehat{E}_i.$
3. $\widehat{P}_{i,A} = \left(\sum_{k=1}^K \frac{\xi_{ik}}{\zeta} \widehat{a}_k^{-1} \widehat{p}_{ik}^{1-\sigma_A} \right)^{\frac{1}{1-\sigma_A}}.$
4. $\widehat{p}_{ik} = \left(\frac{\xi_{ik0}}{\xi_{ik}} \widehat{a}_{k0}^{-1} \widehat{p}_{ik0}^{1-\sigma_e} + \frac{\xi_{ik1}}{\xi_{ik}} \widehat{a}_{k1}^{-1} \widehat{p}_{ik1}^{1-\sigma_e} \right)^{\frac{1}{1-\sigma_e}}.$
5. $\widehat{P}_{i,M} = \left[\sum_{n \in \mathcal{W}} \lambda_{in,M} \left(\tau'_{ii',M} \frac{w'_{nm}}{A'_{nm}} + \tau'_{ii',M} \pi'_{in,M} \right)^{1-\sigma_M} \left(\tau_{ii',M} \frac{w_{nm}}{A_{nm}} + \tau_{ii',M} \pi_{in,M} \right)^{\sigma_M-1} \right]^{\frac{1}{1-\sigma_M}}$
6. $\widehat{p}_{ik}^\theta = \eta_{il|k} \widehat{\lambda}_{ik0}^\theta + \eta_{i1|k} \widehat{\lambda}_{ik1}^\theta$
7. $\widehat{\lambda}_{ike} = \widehat{w}_{iA}^{-\alpha_{ke}/\gamma_{ke}} \widehat{p}_{ix_k}^{-\beta_{ke}/\gamma_{ke}} \widehat{p}_{ike}^{1/\gamma_{ke}}$
8. $\widehat{P}_i^\vartheta = \sum_{l \in \mathcal{K}} \eta_{il} \widehat{A}_{il}^\vartheta \widehat{p}_{il}^\vartheta$
9. $\widehat{\eta}_{ike} = \widehat{\lambda}_{ike}^\theta \widehat{p}_{ik}^{\vartheta-\theta} \widehat{A}_{ik}^\vartheta \widehat{P}_i^{-\vartheta},.$
10. $\widehat{E}_i = (1-\zeta) \widehat{w}_{iM} \widehat{L}_{iM}$ for $i \in \mathcal{U}$
11. $\widehat{E}_i = \frac{\bar{\alpha}_i V_i}{E_i} \widehat{w}_{iA} \widehat{L}_{iA} + \frac{\int_{\Omega_i} r'_i(\ell)}{E_i} d\ell$ for $i \in \mathcal{R}$
12. $\widehat{p}_{ik0} = \widehat{p}_{ik0}$ and $\widehat{p}_{ik1} = \widehat{p}_{ik1} \widehat{m}_{ik}^{-1} \times \widehat{\delta}_{ik}$, where $\widehat{m}_{ik} = \frac{\left(1 + \frac{1}{\varepsilon_{ik}}\right)}{\left(1 + \frac{1}{\varepsilon_{ik}}\right)}.$

A.6 Additional details regarding parameter estimation

$\tau_{ni,S}$ – Iceberg/ad-valorem trade costs

I begin by constructing a measure of the “effective” distance of the transport network between regions or municipalities based on information from the National Network of Roads (RNC). To do so, let a road, or edge in graph theory terminology, be given by e and the effective distance of that edge be given by $t(e)$, where $t(e) = \left(\frac{\text{maximum speed limit}=110\text{km/h}}{\text{speed limit}_e} \right) \times \text{length}_e \times \text{surface type}_e$. If the road is paved, I set $\text{surface type}_e = 1$, and if dirt, I set $\text{surface type}_e = 5$. Letting p denote a given path (a set of connected edges), $E(p)$ the set of edges that path comprises, and $\mathcal{P}_{ni}$ denote all feasible paths between region i and region n , the total effective distance is given by

$$\min_{p \in \mathcal{P}_{ni}} t(p) = \min_{p \in \mathcal{P}_{ni}} \sum_{e \in E(p)} t(e).$$

The solution to the above problem is canonically given by Dijkstra’s algorithm, which I calculate for all pairs of municipalities in Mexico to obtain effective distances as a 2,454 by 2,454 matrix⁵¹.

To obtain $\tau_{ni,k}$, I estimate the following equation:

$$\log(p_{nqkt}/p_{iqkt} - 1) = \beta_0 + \beta_1 \log \text{effective distance}_{ni} + \varepsilon_{it}. \quad (46)$$

To recall, k indexes crop, t represents the year-month (observations are provided at this level), q represents the quantity presentation (usually either in kilograms or a box/basket of a given weight), i represents the urban market which corresponds to origin municipality i , and n represents the destination market. In many cases, prices are either missing for the destination or origin market, and I simply drop these observations. In some cases, I estimate this equation including origin-time, destination-time, and crop-presentation-time fixed effects to control for seasonal fluctuations which may affect these price differences.

⁵¹This problem is reduced in complexity by noting that the graph is undirected, and so effective distances are symmetric between any set of municipalities n and i out of those 2,454. I note that two municipalities are wholly-contained islands, such as Cozumel in Quintana Roo, and are dropped from the analysis due to this complication.

Table A.1: Effective distance regressions

Dependent variable: $\log(p_{nkqt}/p_{ikqt} - 1)$	(1)	(2)	(3)	(4)
log effective distance _{ni}	0.274*** (0.0900)	0.355*** (0.0543)	0.427*** (0.0466)	0.421*** (0.0471)
Constant	-5.079*** (1.223)			
Origin-Time FE	No	Yes	Yes	Yes
Destination-Time FE	No	Yes	Yes	Yes
Crop FE	No	No	Yes	No
Quantity FE	No	No	Yes	No
Crop-Quantity-Time FE	No	No	No	Yes
Observations	69327	67521	67503	59507
R ²	0.0376	0.405	0.479	0.692

Standard errors two-way clustered at level of crops and bilateral pairs.

Constant not reported for regressions with fixed effects.

* $p < .10$, ** $p < .05$, *** $p < .01$

Notes: Estimated using OLS with standard errors two-way clustered at the bilateral pair and crop levels.

I present the results in Table A.1. In columns 1, I obtain my preferred estimates of $\hat{\beta}_0 = -5.079$ and $\hat{\beta}_1 = 0.274$. I extrapolate these estimates to calculate trade costs in the whole of Mexico using the formula $\tau_{nik} = \exp(\beta_0 + \beta_1 \log \text{effective distance}_{nik})$. The reader may note that these coefficients do not assure that $\tau_{nik} \geq 1$, so I set it to $\tau_{nik} = 1$ if not.

I can use these estimates to also estimate the expected trade costs of shipping agricultural products abroad. Based on information I obtain from the US Census Bureau and the Secretariat of Communications and Transport, there are only a few ports/border crossings that account for the vast majority of the exportation of agricultural goods abroad. In Figure A.1, I display the ports and border crossings which account for the vast majority of agricultural exports abroad⁵². Therefore, to calculate the (domestic component of) trade costs required to ship a good abroad, for each municipality/region i , I calculate the closest port and border crossing, using my effective distance matrix. Then the domestic component of trade costs is calculated using the distance elasticity estimated above.

Trade costs with the United States I use information from the US Census Tiger shapefile database to compute effective distance between nodes in the United States in a similar manner. Although tariffs for agricultural trade between the United States and Mexico are largely zero due to NAFTA/USMCA, crossing the border involves large costs, mostly due to the requirement to establish infrastructure on both sides of the border, labor requirements that require changing drivers inland, as well as the time costs of crossing the border. I abstract from the first two considerations, however I obtain the time required to cross the border for commercial vehicles from a Freedom of Information Act Request during my study period. I convert the average waiting time into an effective distance equivalent and calibrate this as my cost of crossing the US-Mexico border.

Other Parameters

σ_A – Elasticity of substitution across crops


In future work I aim to estimate elasticities of substitution using the National Survey of Household Income and Expenditures (ENIGH), but for now I use $\sigma_A = 1$ from the in progress work of Bergquist et al. (2025) (i.e. I assume the utility function is Cobb-Douglas).

σ_e – Elasticity of substitution across supermarkets vs. local markets

I use a related estimate, the elasticity of substitution across store types in Mexico, reported by Atkin et al. (2018),

⁵²For instance, to examine a specific agricultural good: 97% of American tomato imports from Mexico pass through five southern border crossings: Nogales, Arizona, Laredo, Texas, Hidalgo, Texas, Otay Mesa, California, and El Paso, Texas, all of which are displayed in red in Figure A.1.

Figure A.1: Visualization of National Network of Roads



who estimate their parameter to be in the range between 2.28 and 4.36. I set $\sigma_e = 4$ in my calibration.

σ_M – Elasticity of substitution across traded manufacturing goods

I use an estimate of 5 based upon Head and Mayer (2014).

ζ – Sectoral expenditure shares

I calibrate aggregate agricultural spending from production and trade data using the identity $Q - X + I$. For shares of expenditure on traded manufactured goods, I target information on output from domestic product tables.

a_k, a_{ke} – consumption shares of crops k at the national level

As before, I calibrate aggregate agricultural spending on individual crop varieties using production and trade data⁵³ using the identity $Q - X + I$. I calibrate the shares of expenditure across low and high quality crop types a_{ik0} and a_{ik1} to be equal across all crops (i.e. $a_{ik0} = a_{i0}$ and $a_{ik1} = a_{i1}$).

$R(i), U(i)$ – corresponding rural and urban regions

I rely on the definition of metropolitan zones in Mexico from INEGI provided in cleaned format by Blyde et al. (2020), which group together large urban municipalities into commuting zones (for instance, the Valle de México metropolitan zone encompasses Mexico City, which comprises 76 municipalities from 3 states). I modify this definition slightly by adding any localities with more than 100,000 people to this list of metropolitan zones, taking care to merge contiguous “zones” together into one metropolitan zone. Then, to compute the rural outlying municipalities for each urban zone, I use my estimates of effective distance from the road network to compute the urban zone closest to each rural municipality. For the United States, I use the definition of commuting zones made popular by Autor et al. (2013) and provided by the USDA Economic Research Service (ERS).

⁵³I match individual commodities in the production data using a correspondence between commodities and trade data from Fally and Sayre (2018). Sometimes, I am forced to aggregate commodities in the production data to match the trade data, in which cases I divide the aggregated expenditure share into smaller variety-specific expenditure shares based on their production values.

A.7 Additional details regarding definition of clusters

To identify geographically distinct production clusters for each crop, I begin by drawing a circle of fixed radius (200 km., but I experiment with different radii) around each municipality that produces a given crop and summing up the total production of all other municipalities that overlap with the circle encompassing the original municipality. I then rank the production of all such circles centered around different municipalities to determine the largest circle in terms of production (measured as hectares planted of the given crop), or cluster. Then, examining the municipalities contained in the largest cluster, I drop all circles centered around municipalities contained in the largest cluster, as these overlapping circles will often be the second or third largest clusters in terms of production. After dropping these observations, I move on the second largest cluster. After ensuring that none of the municipalities in the second largest cluster overlap with the first (and dropping them if they do), I now have the two largest distinct clusters for a given crop. I then drop all circles centered around municipalities contained in the second largest cluster, and repeat the same process until I recover the full set of non-overlapping circles/clusters covering the entire map of municipalities. I plot these distinct clusters in Appendix Figures A.14 and A.15, where I include the set of producing municipalities contained in each cluster as well as the non-producing ones for the given crop⁵⁴.

I compute the average suitability, yield, and average farm level total factor productivity (TFP) for each crop cluster in all circles (taking the average only over areas where production of that crop is present). Doing so provides me an empirical distribution of yields, TFPs, and suitabilities for all clusters, or similarly sized production regions, I can define in the same way. Using the empirical distribution for each crop, I can then report where in the empirical cumulative distribution function (CDF) the given cluster's average yield, suitability, or TFP falls.

To fix ideas, in appendix figures A.16 and A.17 I display the distribution of suitabilities for all production clusters for avocados and beans, respectively. In these figures the yellow bar indicates the location of the largest production cluster – the largest production cluster for avocados is quite suitable among all other producing regions, whereas the largest production cluster for beans is relatively less suitable. Repeating this analysis over all of my crops, I report the percentile of the largest production cluster for each crop in its respective empirical CDF for yields, suitabilities, and TFPs in Appendix Figure A.5.

In appendix figure A.18, I contrast the results of figure 1 by splitting my sample into crops that have low and high phytosanitary barriers. The relationship between cluster size and productivity (broadly defined) is stronger for crops with higher phytosanitary restrictions to trade (as proxied by the lines of restrictions listed in the US Code of Federal Regulations for their importation) than for crops with fewer restrictions in both Mexico and the United States, although the relationship for all crops is stronger in the United States. One reason for this may be differences in the types of commodities that are likely to have such regulations – cereals and grains are less likely to have such restrictions, whereas specialty fruits and vegetables are more likely to have them, the growth of which may be more dependent on higher levels of suitability. The restrictions themselves may also drive this finding, wherein needing to meet higher standards of crop quality may necessitate locating in higher suitability areas. This finding also holds for associations between cluster size and average farm level TFP, which is much stronger for crops with more phytosanitary restrictions – which suggests that in settings with higher barriers to exportation, only the most productive farmers are able to do so (or that they may learn from agrointermediaries more in these settings).

A.8 Additional figures

⁵⁴Note: I do not visualize these circles of 200km., but rather the administrative boundaries of the municipalities that overlap with them.




Figure A.2: Total avocado exports, by distance between avocado exporting farms and the exporting firm they sell to





Figure A.3: Distances between mango exporting farms and the nearest packing firm


Figure A.4: Agricultural wages for workers

Daily wage for hired agricultural workers in USD, 2022




Source: 2022 Agricultural Census. Areas with no wage data in gray. According to the 2022 agricultural census, 26,984,247 jobs were filled in agricultural sector (if each job was filled by one person this would represent 21.4% of the population, although this measure double counts jobs).


Figure A.5: Top cluster position in empirical yield CDF



(a) Largest cluster position in empirical suitability CDF



(b) Largest cluster position in empirical yield CDF



(c) Largest cluster position in empirical farm TFP CDF

For each crop, I plot in these histograms the location of the top production crop cluster in its empirical CDF. For instance, if a crop is in the right hand bar, its top production cluster falls within the 90% to 100% percent of its empirical CDF for either yields, suitabilities, or TFP. The top crop cluster is defined as the circle of fixed width (i.e. 50km) which contains the most production of that crop within. Suitability comes from the FAO GAEZ project, and if not available, the FAO EcoCrop project. Yield comes from the 2007 Agricultural Census and is averaged across farms to the municipality level. I estimate TFP (more details within the text) at the farm level using the 2007 Agricultural Census. Names plotted for at most 7 top value crops.




Figure A.6: Share of farm units in 2017 National Agricultural Survey (ENA17) reporting exporting difficulties due to phytosanitary (SPS) rules and technical barriers to trade (TBT) by crop




Figure A.7: Share of farm units in ENA17 reporting lack of knowledge of export rules by crop




Figure A.8: Avocado value chain

Source: Comité Sistema Producto Aguacate (2005), translated from Spanish.




Figure A.9: Hectares planted with avocados in 1950 (log)

Source: 1950 Agricultural Census. Areas with zero production are displayed in light gray, with missing municipalities displayed in darker gray. The legend is left censored at 1 hectare to avoid negative log values.




Figure A.10: Hectares planted with maize in 1950

Source: 1950 Agricultural Census. Areas with zero production are displayed in light gray, with missing municipalities displayed in darker gray. The legend is left censored at 1 hectare to avoid negative log values.




Figure A.11: Hectares planted with avocados in 2020

Source: SIAP. Areas with zero production are displayed in light gray. The legend is left censored at 1 hectare to avoid negative log values.




Figure A.12: Hectares planted with maize in 2020

Source: SIAP. Areas with zero production are displayed in light gray. The legend is left censored at 1 hectare to avoid negative log values.




Figure A.13: Total exports for agroexporters with and without US FDI

Source: Digitized trade records from SENASICA phytosanitary certificates for Mexico 2011–2024, matched to a sample of firms found in the Economic Census data using the DENU. Foreign investment data provided by the Economic Census.

Crop	(1) Mexican suit. ratio	(2) American suit. ratio	(3) Mex./Am. suit. ratio	(4) Mex./Am. Clust. size
Beans	0.224	0.800	0.280	2.858
Carrots	0.836	0.004	230.834	1.245
Maize	0.816	1.145	0.713	0.375
Oats	0.000	0.950	0.000	1.538
Olive	0.006	0.619	0.009	0.854
Onions	0.774	0.316	2.451	1.816
Oranges	0.533	0.972	0.548	11.553
Peanuts	0.574	1.046	0.549	0.041
Potatoes	0.081	0.392	0.207	0.278
Rapeseed	1.026	0.921	1.114	0.000
Sorghum	0.585	0.743	0.787	0.895
Soy beans	0.572	1.012	0.565	0.022
Sugar	0.676	1.235	0.547	1.846
Sunflower seeds	1.160	1.017	1.140	17.602
Sweet potato	0.936	1.850	0.506	0.036
Tobacco	0.281	1.362	0.207	38.922
Tomatoes	0.866	0.419	2.067	0.236
Wheat	0.139	0.546	0.255	0.082

Table A.2: Ratio of agricultural productivity 2041-2070 over 1981-2010 for top Mexican and American ag. clusters

In column 1, I compute the predicted agricultural productivity of the top Mexican agricultural production cluster in 2041-2070 over their current agricultural productivity (i.e. in 1981-2010), where I define top clusters as described in Section 3.2. In column 2, I plot the same for the top American cluster. Column 3 reports the ratio of column 1 over column 2. Column 4 reports the ratio of the current size of the top Mexican production cluster over the top American one (in terms of number of hectares planted of each crop in both “clusters”).




Figure A.14: Largest non-overlapping clusters for avocados




Figure A.15: Largest non-overlapping clusters for beans




Figure A.16: Distribution of observed average suitabilities for all similarly defined avocado clusters




Figure A.17: Distribution of observed average suitabilities for all similarly defined bean clusters

Table A.3: Relationship between farm growing only one crop and presence of downstream plant specializing in it

Dependent variable: Farm produces only 1 crop	(1)	(2)	(3)	(4)
Mun. has packer	0.244 (0.00282)	0.213 (0.00285)	0.210 (0.00284)	0.210 (0.00284)
Neighbor mun. has packer		0.130 (0.00204)	0.0648 (0.00217)	0.0652 (0.00217)
Metro. area has packer			0.0917 (0.00107)	0.0908 (0.00108)
Rainfed suitability $\in [0, 1]$				-0.0110 (0.000932)
Municipality FE	X	X	X	X
Crop FE	X	X	X	X
Sample	Full	Full	Full	Full
N	4,792,134	4,792,134	4,792,134	4,792,134

Robust standard errors in parentheses


Data is from by the 2007 Agricultural Census. Dependent variable is whether farm produces only one crop. Agroiintermediary/agroexporter data comes from SADER, suitabilities from FAO GAEZ/EcoCrop. Municipality and crop fixed effects are included in all columns.

Table A.4: Input cost shares by crop, only for multi-crop farms


Crop	Farms w/ int. inputs				Farms w.o. inputs		
	α	β	γ	$\alpha + \beta + \gamma$	α	γ	$\alpha + \gamma$
Avocados	0.123	0.675	0.087	0.885	0.123	0.087	0.210
Bananas	0.235	0.152	0.301	0.687	0.237	0.262	0.499
Barley	0.280	0.381	0.235	0.895	0.280	0.278	0.558
Beans	0.248	0.655	0.574	1.476	0.248	0.574	0.821
Coffee	0.112	0.501	0.129	0.742	0.149	0.208	0.357
Cotton	0.188	0.622	0.116	0.927	0.186	0.121	0.307
Lemons	0.044	0.764	0.105	0.913	0.044	0.105	0.149
Maize	0.251	0.725	0.291	1.266	0.263	0.487	0.750
Mango	0.145	0.006	0.275	0.425	0.177	0.290	0.467
Oats	0.078	0.432	0.420	0.930	0.108	0.429	0.537
Oranges	0.231	0.658	0.107	0.997	0.231	0.107	0.338
Sorghum	0.168	0.557	0.206	0.931	0.200	0.291	0.491
Soy beans	0.272	0.450	0.475	1.197	0.208	0.518	0.726
Sugar	0.063	0.186	0.400	0.649	0.065	0.424	0.489
Tomatoes	0.109	0.428	0.209	0.746	0.109	0.209	0.319
Wheat	0.233	0.631	0.317	1.182	0.225	0.338	0.563

Notes: Estimated using two-stage least squares with robust standard errors, using the 2007 Agricultural Census. I take the median of all state level estimates for each crop. α is the cost share of labor, β is the cost share of fertilizer, and γ is the cost share of land. Since I cannot estimate equation 14 if fertilizer use is zero, I split the sample into farm units with positive and zero fertilizer use and estimate the coefficients separately for both groups. Sample consists only of farm units growing multiple crops, where labor and fertilizer are allocated according to the hectares planted of each crop.


Figure A.18: Correlation between size of cluster and cluster suitability, yields, and TFPs



(a) Size and suitability in Mex. (low SPS-rule crops) (b) Size and suitability in Mex. (high SPS-rule crops) (c) Size and suitability in US (low SPS)



(d) Size and suitability in US (high SPS) (e) Size and yield in Mexico (low SPS) (f) Size and yield in Mexico (high SPS)



(g) Size and TFP in Mexico (low SPS) (h) Size and TFP in Mexico (high SPS)

High phytosanitary strictness (SPS) crops are crops with more than 7 lines in the US Code of Federal Regulations; low SPS crops are those with 7 lines or less.




Figure A.19: Placebo check: Event study using false year and crop packer assignment
Crop share

Figure A.20: Estimation results for avocado, using A_{ik} from maximum likelihood estimation rather than fundamental agro-ecological suitabilities

