

Unfair Trade?

Monopsony Power in Agricultural Value Chains

Lucas Zavala*

March 4, 2022

Abstract

Exporters exercise monopsony power in agricultural value chains. I quantify the consequences for farmers. I link exporters to farmers across the universe of cash crops in Ecuador and document that farmers earn a small share of export revenues. I develop a model in which exporter monopsony power driven by two elasticities that govern farmer costs of switching crops and switching exporters. The estimated elasticities imply that farmers are paid only half of their marginal revenue products. A modest counterfactual Fair Trade policy would have the same effect as an unreasonably high universal price floor, increasing increase farmer income by 12%.

*Princeton University and The World Bank, lzavala@princeton.edu.

I thank my advisors, Peter Schott, Michael Peters, Samuel Kortum, and Lorenzo Caliendo for support throughout the project. I thank Pol Antras, Costas Arkolakis, Juan Manuel Castro Vincenzi, Daisy Chen, Swati Dhingra, Dave Donaldson, Pablo Fajgelbaum, Manuel Garcia-Santana, Penny Goldberg, Alejandro Graziano, Brian Greaney, Gene Grossman, Kyle Herkenhoff, Federico Huneus, Amit Khandelwal, Florence Kondylis, Martin Mattsson, Joan Monras, Monica Morlacco, Eduardo Morales, Tristan Reed, Andres Rodriguez-Clare, Michael Rubens, and Diana Van Patten for helpful comments. I also thank Paul Carrillo, Dina Pomeranz, and the Center for Fiscal Studies at the Internal Revenue Service of Ecuador for assistance acquiring the data. This work was supported by the National Science Foundation GRFP and the Yale MacMillan Center for International and Area Studies.

Exporters are critical links in agricultural value chains. They connect small-holder farms in emerging economies to lucrative international markets. At the same time, export markets for cash crops are often dominated by a few large intermediaries. These intermediary exporters can use their bargaining power to depress farmgate prices, limiting the benefits of globalization for low-income farmers. Despite the crucial role played by exporters, research has focused on the market power of local intermediaries. This paper quantifies the effect of exporter monopsony power on farmer income across the universe of exported cash crops in Ecuador.

There are two main challenges to measuring the market power of intermediaries over farmers. First, it requires information on prices at the farm gate as well as in retail markets. Second, a demand shock is needed to separate monopsony power from other determinants of farmer prices, such as value added by intermediaries. Prior work has overcome these challenges by focusing on a specific crop market. Some studies rely on variation induced by policy ([Chatterjee 2020](#); [Dhingra and Tenreyro 2020](#); [Van Patten and Mendez-Chacon 2021](#)), while others induce this variation experimentally ([Bartkus, Brooks, Kaboski and Pelnik 2021](#); [Bergquist and Dinerstein 2020](#); [Casaburi and Reed forthcoming](#)).

I follow a different approach, combining firm-level export data with domestic firm-to-firm transactions across the entire agricultural sector in Ecuador. This allows me to trace every dollar earned by exporters all the way back to farms for 157 products as diverse as banana and shrimp. These data also allow me to construct demand shocks from variation in international prices, which are widely available across products and over time. Finally, the matched nature of my data allows me to infer monopsony power directly from the firm-level pass-through of price shocks using a discrete choice model of farmer cropping decisions.

Exporter monopsony power in the model is determined by the ability of farmers to substitute (a) across crops and (b) across exporters within a crop. The model specifies how variation in pass-through to farmer income as a function of exporter size in the data can be used to estimate these two elasticities of substitution. Given the elasticities, the model further allows me to conduct two sets of counterfactual exercises. The first measures the effect of monopsony power by calculating how much higher farmer income would be in a first-best world with

competitive exporters. The second compares farmer income under two popular second-best policies: a Fair Trade program and a universal price floor.

I document three new facts about agricultural value chains in Ecuador. First, agricultural exporters are highly concentrated, with just a few intermediaries in each crop purchasing the entire value produced by farmers. Second, for every dollar earned by a given exporter on international markets, the farmers that supply him receive less than a quarter on average. Third, this “farmer share” is lower when the exporter purchases a greater share of domestic output of a given crop, even after controlling for observable measures of value added. These facts are consistent with exporters exercising monopsony power over farmers.

To infer monopsony power from these patterns, I extend a frontier model of monopsony in labor markets ([Berger, Herkenhoff and Mongey forthcoming](#)) to the context of crop markets. Farmers choose which crop to produce and which exporter to supply. They trade off the price offered by each exporter with their idiosyncratic productivity shock for growing that crop and idiosyncratic cost shock for supplying that exporter. Through these shocks, the model captures the land’s suitability for different crops and proximity to different exporters, two key dimensions of heterogeneity in the agricultural trade literature ([Costinot, Donaldson and Smith, 2016](#); [Sotelo, 2020](#); [Farrokhi and Pellegrina, 2021](#); [Bergquist, Faber, Fally, Hoelzlein, Miguel and Rodriguez-Clare, 2019](#)). The more costly it is for farmers to switch from growing coffee to growing cocoa, or to switch from supplying one coffee exporter to supplying another, the greater the scope for monopsony power.

Exporters act strategically when purchasing crops in the model, internalizing their influence over prices. The optimal price they offer farmers is marked down from the price they receive in competitive international markets. Assuming Cournot competition among exporters in domestic crop markets, this price can be written in closed form as a function of exporter size and the elasticities of substitution within and across crops. The average farmgate price is low when the elasticity of substitution across exporters within a crop is low, and it declines with exporter size when the elasticity of substitution across crops is low – consistent with the cross-sectional facts. One contribution relative to [Berger et al. \(forthcoming\)](#), whose model features a closed form for worker wages as a

function of firm size and the elasticities of substitution within and across labor markets, is that I can structurally interpret the elasticities of substitution within and across crops in terms of the heterogeneity in trade costs and land quality discussed above.

A challenge to inferring these elasticities of substitution purely from the cross-sectional facts is that the average farmer share reflects both monopsony power and value added. To overcome this, I exploit the fact that Ecuador is a small open economy and estimate the elasticities from the pass-through of shocks to international prices.¹ This approach draws on a large literature of using variation in the pass-through of *cost* shocks, such as exchange rate changes, to estimate *demand* elasticities and measure *seller* market power (Bergquist and Dinerstein, 2020; Atkin and Donaldson, 2015; Nakamura and Zerom, 2010; Atkeson and Burstein, 2008; Rubens, 2021). A second contribution relative to Berger et al. (forthcoming), who estimate their model using aggregate tax shocks, is that changes in international prices are widely available, arguably exogenous, and idiosyncratic shocks.²

Following a shock to the international price faced by exporters, the domestic price they offer to farmers only partially adjusts. This pass-through is also lower when the exporter purchases a greater share of domestic output of a given crop. In the model, the average pass-through of international price shocks to domestic producer prices is low when the elasticity of substitution across exporters is low, and declines with exporter size when the elasticity of substitution across crops is low. The elasticities can be recovered via indirect inference by matching pass-through magnitudes in the model with those in the data. The implied elasticities are small, indicating that farmers face substantial monopsony power.

I use the estimated model to quantify this monopsony power across the universe of exported agricultural products. I find that farmer prices are marked down from their marginal revenue products by 51%, implying large gains simply from eliminating markdowns and redistributing exporter profits to farmers. Indeed,

¹An alternative approach is to infer the elasticities indirectly by measuring value added via production function estimation (De Loecker and Warzynski, 2012; Morlacco, 2020; Rubens, 2021; De Loecker, Goldberg, Khandelwal and Pavcnik, 2016).

²A recent paper, released after the first draft of this paper, uses a similar insight to estimate the labor market power of Brazilian exporters (Felix 2021).

a counterfactual economy with perfectly competitive exporters would see a 77% increase in farmer income, two thirds of which is explained by redistribution. The remaining third are efficiency gains from farmers reallocating across crops and across exporters within crops.

In the final part of the paper, I use the estimated model to study the impact of two popular agricultural support policies: Fair Trade programs and broad price floors. Fair Trade is the fastest-growing certification program for sustainable farming. Buyers pay higher prices to promote the economic well-being of certified farmers, which they recover by selling a differentiated Fair Trade product to consumers who care about farmer well-being. Several studies evaluate the impact of Fair Trade programs in the coffee sector ([Dragusanu and Nunn, 2018](#); [De Janvry, McIntosh and Sadoulet, 2015](#)). Instead, I focus on its potential impact across the entire agricultural sector and in comparison with other pro-farmer policies.

Following [Podhorsky \(2015\)](#), I model Fair Trade by introducing an exporter in each crop who behaves competitively and therefore pays a premium relative to other exporters. This has a positive direct effect on the farmers who supply the Fair Trade exporter. It also has a positive *indirect* effect, since the Fair Trade exporter reduces the market power of other exporters, forcing them to raise prices. Together, these effects can raise farmer income up to 25%, depending on the market share of the Fair Trade exporter.

To highlight the effectiveness of Fair Trade, I consider a second policy in which the government sets a broad price floor in each crop. This also has a positive direct effect on prices, since exporters can no longer offer prices below the floor. Unlike Fair Trade, however, it has a *negative* indirect effect. The smallest exporters contract, increasing the market power of larger exporters who can afford to pay the minimum price. Because of these offsetting effects, high price floors are required to realize the income gains from even a modest Fair Trade program.

The paper is organized as follows. In Section [I](#), I provide an overview of agriculture exports in Ecuador, discuss the construction of my value chain dataset, and present key facts. In Section [II](#), I develop a model of farmer crop choice and exporter strategic pricing to quantify market power. In Section [III](#), I estimate the model and validate it. In Section [IV](#), I use the estimated model to measure the market power faced by farmers. In Section [V](#), I conduct counterfactual analyses

of Fair Trade and other agricultural support policies. I conclude in Section VI by discussing the limitations of the current study and the directions for future research.

I Data

In this section, I construct the entire value chain across the universe of exported crops in Ecuador. To do so, I combine administrative microdata on firm-product exports from Customs declarations, firm-to-firm transactions from VAT declarations, and firm characteristics from a national registry. I document three new facts about value chains using this dataset, which together point to the importance of exporter market power.

I.A Ecuador: an ideal setting

Ecuador is a microcosm of the issues surrounding agricultural trade in emerging economies. GDP per capita in Ecuador is a little over \$6,000, close to the global median. Agriculture employs almost 30% of the workforce and accounts for over half of export revenues. Across all developing countries, agriculture employs 40% of the workforce and generates a third of export revenues (Cheong, Jansen and Peters 2013).

Despite its small size, Ecuador is an important producer of cash crops such as cocoa, coffee, bananas, palm, shrimp, tuna, and cut flowers. More generally, developing countries account for more than a third of agricultural trade, and more than half of seafood trade (Aksoy and Beghin 2004). Cash crops are typically produced by many small farms, and exported by only a handful of large firms. Domestic consumption of cash crops is low, as they command much higher prices in international markets. Across South America, the largest 5% of exporting firms receive 80% of export revenue (Cunha, Reyes and Pienknagura 2019). In contrast, most crops are produced on small farms, and average farm size has been decreasing over time (Lowder, Scoet and Raney 2016). Even in the banana sector, which has historically been dominated by vertically-integrated, multinational giants like Chiquita and Dole, there has been a trend toward divestment from plantations (FAO 2014). In Ecuador, these multinationals control less than 20% of the export market, and most of the remaining exporters do not produce bananas themselves, but instead source from thousands of producers (Wong 2008).

A disproportionate share of the poor work in agriculture, both in Ecuador and across developing countries. Income gains in the agricultural sector are therefore crucial for reducing poverty. Ecuador offers an ideal setting for studying an important barrier to such gains: the lack of competition among exporters. In field interviews, producers cited low bargaining power as a barrier to receiving higher prices. To examine this barrier on a large scale, I partner with the Tax Authority of Ecuador (*Servicio de Rentas Internas*, henceforth SRI) to access several administrative databases, which together allow me to trace the value of crops all the way from farm to port.

I.B Constructing agricultural value chains

A key challenge to tracing the value of crops is that farmers typically do not export directly. To overcome this challenge, I proceed in several steps: (1) calculate the value received by exporters, (2) match exporters to their suppliers, (3) calculate the value received by each supplier, and (4) identify which suppliers are farmers. To do so, I combine several administrative datasets obtained in collaboration with the SRI. I outline the procedure here and provide details in the appendix.

The first dataset covers the universe of export transactions from 2008-2011. The data are compiled from Customs declarations and contain the value (in USD) and quantity (in kg) traded internationally for each firm, product, and year. Products are classified at the HS 6-digit level. I use these data to calculate the price and value received by exporters (step 1). I restrict my attention to 157 animal products, vegetable products, and foodstuffs (HS 2-digit codes 01-24), which represent roughly half of all exports from Ecuador. For convenience, I will refer to all of these products as “crops.”

The second dataset captures the universe of domestic firm-to-firm transactions from 2008-2011. The data are derived from value added tax (VAT) declarations and measure the value transacted for each buyer-seller pair in each year. Using these data, I construct the network of suppliers for each exporter in (step 2). I can then calculate the value paid by each exporter to each of his suppliers (step 3).

The third dataset contains basic characteristics for all firms active in 2011. The data are pulled from a national register and include the industry and location

of each firm. Industries are classified at the ISIC 5-digit level. I use these data to identify which suppliers are farmers (step 4). Taxpayers in the agriculture, forestry, and fishing industries (ISIC 2-digit codes 01-03) are classified as farmers.

My novel agricultural value chain dataset comprises over 800 exporters selling 157 agricultural products sourced from almost 50,000 farmers. Table 1 summarizes the farmers and exporters in my dataset. The median exporter earns over \$1 million in revenue and employing more than 20 people. Exporters report large payments to workers and domestic suppliers, but earn an average profit rate of 25%. In contrast, the median farm is tiny, earning less than \$9,000 annually. Furthermore, 94% of farmers are self-employed. Almost three quarters of exporters are in the wholesale sector, implying that few farmers export directly.³ However, 75% of farmer sales are indirectly exported, indicating the importance of constructing the value chain.

A few important concerns arise when using tax information to study agricultural value chains. First, information may be missing due to informal labor in the agricultural sector. Several factors mitigate this concern. The VAT records underlying my dataset are filed by the purchasing firm, in this case an exporter. Firms have an incentive to *over-report* the value they pay to farmers, as their tax liability is assessed on profits (Pomeranz, 2015). Instead, I show in Section I.C.3 that larger exporters pay farmers disproportionately *less*. Exporters may have an incentive to under-report costs if the probability of fraud detection depends on their profit rate (Carrillo, Pomeranz and Singhal, 2017).⁴ In this case, the optimal reported profit rate should not change following a shock to the exporter. Instead, I show in Section III.B that following a demand shock, larger exporters adjust payments to farmers disproportionately *less*.

³An exception is the cut flower industry, where many small farms export directly. I exclude direct exporters from the analysis and exclude the cut flower industry entirely.

⁴Exporters with negligible market share often report no costs, consistent with under-reporting. To control for this measurement error, I include an indicator for these firms in my cross-sectional regressions.

Table 1: Farmer and exporter statistics

(a) Exporters		(b) Farmers	
\$ Sales	1,177,543	\$ Sales	8,678
\$ Purchases	543,053	\$ Purchases	0
\$ Wage Bill	108,246	\$ Wage Bill	0
# Employees	21	# Employees	0
% Wholesale	74	% Self-employed	94
% Single-product	76	% Export Intensity	75
Observations	804	Observations	49,475

Notes: Table summarizes agricultural exporters and the farmers who supply them. Panel A shows summary statistics across exporters. Panel B shows summary statistics across farmers. Rows 1-4 show medians. Rows 5-6 show means. “Wholesale” indicates exporters in ISIC 2-digit sector 46-47; “Single-product” indicate exporters sell a single HS 6-digit product; “Self-employed” indicates farmers filed the simplified F102 tax form; “Export Intensity” indicates the share of farmer sales indirectly exported.

A second concern is that the data may not be capturing small family farms, but rather large factory farms. The median farm does not report any employees or wages, consistent with the high rate of self-employment. SRI officials confirmed that family farms appear in the data as self-employed taxpayers. For larger farms, I could calculate farmer income using wage bill and employment data. However, some farm owners may be farmers, and some farm employees may not be farmers. To avoid distributing farm profits and arbitrarily deciding who is a farmer, I measure farmer income as sales, making no distinction between farms and farmers. This overestimates farmer income and underestimates the number of farmers. Later, I estimate the model without using any information on the number of farmers or the size of farms.

A final limitation is that VAT records measure trade between firms in general rather than trade of a particular *product* between firms. A few features of agricultural value chains in Ecuador allow me to overcome this limitation. First, unlike in more complex value chains, where firms in different industries produce important components of the final product, the key producers in agricultural value chains are farmers and fishers. They are the ones who harvest fruits from plants and fish from water, and since I observe them in my dataset, I can pin down both ends of the value chain. If the exporter at one end only exports coffee, I assume

that the product he purchases from the farmer at the other end is coffee. This is a reasonable assumption for Ecuador, where (a) the majority of exported crops are produced exclusively for the international market and (b) the majority of exporters export a single crop. Table 1 shows that 76% of exporters fall into this category. I assign multi-product exporters to their top product, which accounts for 93% of exports for these firms. Finally, farmers typically sell to a single exporter, so it is unlikely that farmers produce multiple different crops for export. Together, these facts imply that I can infer the product being traded between farmers and exporters in my dataset.

Table 2 summarizes the funnel-like structure of agricultural value chains.⁵ The median exporter buys from 24 farmers (exporter indegree), but the median farmer only sells to a single exporter (farmer outdegree). This is a key feature of the data which points to the existence of monopsony power. It holds both on aggregate and within many of the top exported products. For example, shrimp is the second most important product, with over 2 billion dollars in export sales. There are almost 6,000 shrimp producers along the coast, but only 50 shrimp exporters. In the next section, I leverage the micro-structure of the data to establish further evidence of market power.

Table 2: Exporter-farmer relationships

	\$ Exports (Millions)	# Exporters	# Farmers	Exporter Indegree	Farmer Outdegree
All Crops	16,954	804	49,745	24	1
Bananas	6,038	188	9,685	81	3
Shrimp	2,208	50	5,729	77	1
Tuna	2,043	22	1,825	54	1
Cocoa	1,314	56	17,686	363	2
Palm oil	616	13	7,821	1,640	2
Coffee	110	17	1,611	28	1

Notes: Table summarizes exporter-farmer relationships across 157 crops. Crops are defined as HS 6-digit products in chapters 01-24. Row 2 shows all crops. Rows 3-8 show a selection of important crops. Columns 2-4 show totals. Column 5 shows the median number of farm suppliers across exporters. Column 6 shows the median number of exporting customers across farmers.

⁵See the appendix for a breakdown by crop category.

I.C Exporter concentration and the farmer share

I document three new facts about supply chains of agricultural exports from Ecuador. Together, they suggest that exporters exercise monopsony power in crop markets, and they motivate the development of a model to explore the consequences for small farmers.

I.C.1 Crop markets are highly concentrated

The number of exporters in a given crop market may understate the degree of concentration. For example, the cocoa market has 56 exporters in Table 2, but the top 4 cocoa exporters control almost the entire export market. I take advantage of the micro-structure of my dataset and define the *effective number of exporters* as the number of exporters required to control 90% of the market for a given crop.

To examine the potential for market power across a broad range of crops, I divide crops into six bins based on the effective number of exporters: 1, 2, 3, 4, 5-9, 10+. Figure 1a plots the distribution across these bins for more than 100 crops. The majority of crop markets are highly concentrated: the median crop is dominated by a single firm, and almost all crops have fewer than 10 exporters. However, concentration on its own does not imply market power. To provide further evidence, I take advantage of matched exporter-farmer nature of my dataset in the next fact.

I.C.2 Farmers receive a small share of the export value of their crops

Exporters exercise market power over farmers by forcing them to accept lower prices. To investigate this in my dataset, I compute the value that each exporter pays to farmers as a share of the value he earns from selling their crops on the international market.⁶ I refer to this as the *farmer share* for exporter i of crop j :

$$\text{farmer share}_{ij} \equiv \frac{\text{exporter } i\text{'s purchases of crop } j}{\text{exporter } i\text{'s sales of crop } j}$$

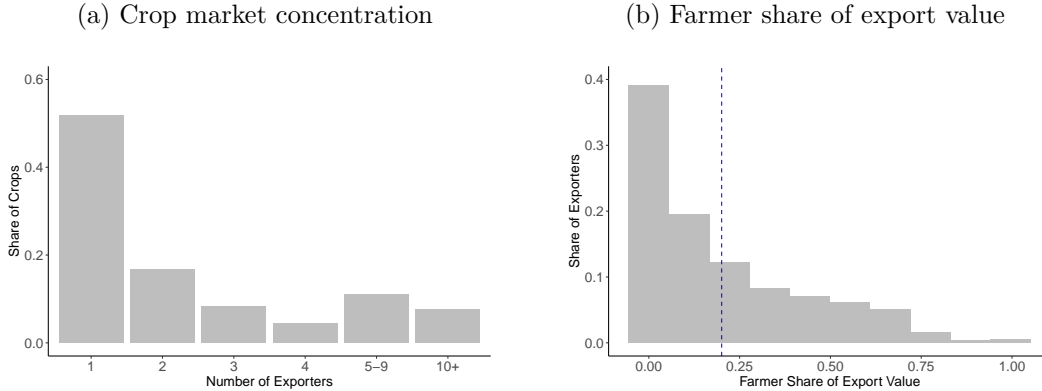
Figure 1b shows the distribution of the farmer share across all exporters. The blue line indicates an average farmer share of 0.24, meaning that for every dollar of agricultural products exported from Ecuador, farmers earn 24 cents. Many

⁶Because I do not observe the quantity purchased from each farmer, I cannot compute this share at the farmer level. Larger exporters purchase from many more farmers, but only purchase slightly more value per farmer, .

exporters have farmer shares lower than 10%, while very few have shares above 50%.

An alternative explanation for the low farmer shares depicted in Figure 1b is that exporters add value to crops by transforming or transporting them. For example, a cocoa exporter may re-package the beans he purchases from farmers before selling them internationally, or ship them from the eastern Amazon provinces, where a substantial share of cocoa, is grown to the coastal port of Guayaquil. In my dataset, this could appear as wages or payments to suppliers who are *not* classified as farmers. I exploit this dimension of the data to establish the next fact.

Figure 1: Exporter concentration and farmer shares



Notes: Panel A plots the distribution of the effective number of exporters across 157 exported crops. “Effective number of exporters” is defined as the minimum number of exporters required to reach 90% market share in the domestic market for crop purchases. Bars indicate the proportion of crops with 1, 2, 3, 4, 5-9, and 10 or more exporters. Panel B plots the distribution of the farmer share across exporters. “Farmer share” is defined as an exporter’s purchases of a crop from farmers divided by his sales of the same crop on international markets. The dashed blue lines indicates that the average farmer share is 0.24.

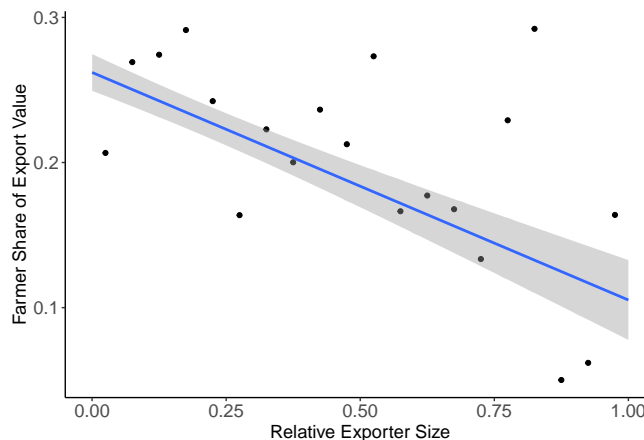
I.C.3 Farmer shares are lower when exporters are more concentrated

Neither the high exporter concentration in fact 1 nor the low farmer shares in fact 2 alone are sufficient evidence of market power. To establish a connection between them, I define the *relative size* of exporter i in crop j as the value purchased by exporter i as a share of the total market for crop j .

$$\text{exporter size}_{ij} \equiv \frac{\text{exporter } i\text{'s purchases of crop } j}{\text{total purchases of crop } j}$$

An exporter with relative size near 1 controls the entire market for a crop and is therefore a *monopsonist*, while an exporter with relative size near 0 exerts little control. If the relative size of an exporter measures his potential for market power, and he realizes this potential by forcing farmers to accept lower prices, then we should see a negative relationship between farmer shares and relative exporter size. Figure 2 confirms this: on average, an exporter who controls all of the market pays 20 percentage points less to farmers than an exporter who controls none of it. At the mean farmer share of 0.25 in Figure 1b, this represents an 80% decrease.

Figure 2: Relationship between farmer share and exporter concentration



Notes: Figure plots exporter size on the x-axis and farmer shares on the y-axis. “Exporter size” is defined as the share of the domestic market for a given crop purchased by a given exporter. “Farmer share” is defined as an exporter’s purchases of a crop divided by his exports of the same crop. Each dot indicates the average farmer share within bins of 5% market share. Solid blue line indicates predictions from a linear regression where each observation is an exporter-crop-year. Grey area indicates a 95% confidence interval.

Figure 2 pools exporters across all crops. However, farmer shares should be lower in crops that require extensive transformation or transportation. If this in turn requires large fixed investments in machines or vehicles, such crops may have fewer exporters in equilibrium. For example, the shrimp market may have more exporters and larger farmer shares than the cocoa market simply because shrimp is sourced along the coast, whereas cocoa is sourced as far as the Amazon, removed from major ports. In this case, farmer shares and relative exporter size

would be negatively correlated, even if exporters did not exercise market power. A similar phenomenon may play out within crops. For example, 80% of cocoa is grown in coastal provinces. If sourcing the remaining 20% from inland provinces requires large fixed investments that only large exporters can afford, the same spurious correlation would arise.

To explore the negative relationship between farmer shares and relative exporter size in more detail, I estimate a series of regressions:

$$\log(\text{farmer share}_{ijt}) = \beta \text{exporter size}_{ijt} + X'_{ijt} \Gamma + \delta_{jt} + u_{ijt}$$

where X is a vector of controls, δ is a crop-year fixed effect, u is an error term, and t indexes the year. The coefficient of interest, β , measures the relationship between exporter size and farmer shares. Table 3 displays the results. Column 1 shows the baseline specification with no controls or fixed effects, consistent with Figure 2. Column 2 includes product-year fixed effects to control for unobserved differences in processing and fixed costs of sourcing across crops. Because many 6-digit products (crops) are effectively controlled by a single exporter, fixed effects are at the 2-digit product level. Column 3 controls for systematic differences across exporters by adding wages, payments to non-farm suppliers, and log export prices. Wage payments help capture exporter-specific value added, while non-farm payments help measure observable costs of sourcing.⁷ Log export prices help control for quality differences across exporters.

Table 3: Relationship between farmer share and exporter concentration

	Log Farmer Share	Log Farmer Share	Log Farmer Share
	(1)	(2)	(3)
Exporter Size	-0.823 (0.158)	-0.681 (0.185)	-0.530 (0.180)
FE	No	Yes	Yes
Controls	No	No	Yes
Observations	1,923	1,923	1,923
R ²	0.014	0.355	0.397

Notes: Table summarizes OLS regressions of log farmer shares on relative exporter size. Each observation is an exporter-crop-year. Controls include: log wage bill, log payments to non-farmer suppliers, log export unit values, and an indicator for exporters with market share smaller than 1%. Clustered standard errors are shown in parentheses.

⁷This includes all 2-digit ISIC sectors except agriculture (01-03) and domestic wholesale (45-47). The next largest sector is transportation and storage (49-53).

My preferred specification in Column 3 effectively compares two exporters selling the same crop internationally with similar value added. If one purchases 99% of production in Ecuador and the other purchases the remaining 1%, the coefficient indicates that the former pays farmers a 53% smaller share of his export revenue. This suggests that larger exporters have market power over farmers. To quantify the importance of market power, I develop a framework in the next section that links exporter size to farmgate prices via farmer substitution patterns across crops and across exporters within a crop. Unobserved exporter-specific fixed costs, such as branding, remain an alternative to market power. To help rule this, I will rely on variation within exporters over time to estimate the model.

II Theory

In this section, I develop a model of imperfect competition among exporters in the market for crops. Farmers choose a crop to produce and sell to exporters, who have market power. The concentration of exporters, and hence their market power, differs across and within crops and impacts farmer well-being. The formulation of the model builds on the work of [Atkeson and Burstein \(2008\)](#) and [Berger et al. \(forthcoming\)](#). I model the farmer’s choice of crop and exporter as a discrete choice problem, which yields a nested CES supply curve for crops. Given this supply curve and Cournot (or Bertrand) competition among exporters, the equilibrium farmer share is a decreasing function of relative exporter size, consistent with Section [I.C.3](#). The shape of this function is determined by two key elasticities which govern the heterogeneity of costs in the farmer’s choice problem. Intuitively, the more heterogeneous are farmer costs, the greater the consequences of exporter market power. In this way, the model also connects to the work of [Costinot et al. \(2016\)](#) and [Sotelo \(2020\)](#).

II.A The value chain

The value chain consists of two agents: a continuum of farmers and a finite number of exporters. Crops such as shrimp and cocoa are indexed by $j \in \{1, \dots, M\}$. Each crop is sold by an exogenous, finite number of exporters, indexed by $i(j) \in \{1, \dots, N(j)\}$. Each exporter purchases the crop from farmers,

adds some value, and sells it internationally. For example, cocoa exporters may pack beans into bags or ship them across the country before selling them abroad. Crops are produced by a continuum of farmers, indexed by $f \in [0, 1]$. Consistent with the funnel-like nature of value chains in Section I, farmers choose a single crop to produce and a single exporter to supply, while exporters buy a single crop from a measure of farmers. Figure A1 in the appendix summarizes the structure of the model.

II.B Farmer crop choices

Farmer f is endowed with a unit of land, which she farms inelastically with efficiency $q_f \sim G$. This is the only source of ex-ante heterogeneity among farmers and reflects differences in farmer productivity and land quality. The farmer makes two decisions: which crop to produce and which exporter to supply. She receives an idiosyncratic shock ν_{fj}^c for producing each crop j and an idiosyncratic shock $\nu_{fi(j)}^e$ for supplying each exporter $i(j)$. Since each exporter buys and sells a single crop, $i(j)$ uniquely identifies an exporter. For convenience, I drop the parentheses in subscripts, so that ν_{fij}^e becomes shorthand for $\nu_{fi(j)}^e$.

A farmer with efficiency q_f can supply q_{fij} units of crop j to exporter i :

$$q_{fij} = e^{\frac{\nu_{fj}^c}{1+\theta}} e^{\frac{\nu_{fij}^e}{1+\eta}} q_f$$

where η and θ are two key elasticities discussed in detail below. The idiosyncratic shocks determine her yield: the higher are ν_{fj}^c and ν_{fij}^e , the more she can supply if she chooses crop j and exporter i . One way to interpret these shocks is that ν_{fj}^c models the land's suitability for growing crop j in a stochastic way, while ν_{fij}^e models geographic proximity to exporter i in a stochastic way.

Each exporter buys and sells a single product, offering price p_{ij} to all farmers. Farmers trade off higher prices with lower idiosyncratic shocks: a shrimp exporter in the coastal port of Guayaquil may pay a high price, but it does them little good if they happen to live far away in the Ecuadorian Amazon, where the shock for producing shrimp and reaching Guayaquil is prohibitively low. If the farmer chooses crop j and exporter i , she earns profits $p_{ij}q_{fij}$. She chooses a crop and exporter by solving:

$$\arg \max_{i,j} p_{ij}q_{fij} = \arg \max_{i,j} \left\{ \log p_{ij} + \log q_f + \frac{\nu_{fj}^c}{1+\theta} + \frac{\nu_{fij}^e}{1+\eta} \right\}$$

The probability that farmer f chooses crop j and exporter i , $\Pr(fij)$, is independent of her efficiency, q_f . This implies that the model can accommodate any distribution of land quality or farmer productivity. I assume ν_{fij}^e follows an extreme value distribution, and ν_{fj}^c is distributed such that the sum $\nu_{fij} = \frac{\nu_{fj}^c}{1+\theta} + \frac{\nu_{fij}^e}{1+\eta}$ follows a Gumbell distribution (Cardell 1997).⁸ Under this assumption, $\Pr(fij)$ follows a nested logit structure: it can be written as a product of the marginal probability of choosing crop j and the conditional probability of choosing exporter i , conditional on choosing crop j :

$$\Pr(f \text{ chooses exporter } i, \text{ crop } j) = \underbrace{\frac{p_{ij}^{1+\eta}}{\sum_{i(j)} p_{ij}^{1+\eta}}}_{\Pr(f \text{ chooses exporter } i|j)} \times \underbrace{\frac{(\sum_{i(j)} p_{ij}^{1+\eta})^{\frac{1+\theta}{1+\eta}}}{\sum_j (\sum_{i(j)} p_{ij}^{1+\eta})^{\frac{1+\theta}{1+\eta}}}}_{\Pr(f \text{ chooses crop } j)}$$

This expression has an intuitive interpretation: conditional on choosing crop j , the probability of choosing exporter i , $\Pr(i|j)$ depends on how large the price of exporter i (numerator) is relative to the price index of crop j (denominator), which is a CES aggregate of prices across exporters within a crop. The unconditional probability of choosing crop j , $\Pr(j)$, then depends on how large the price index of crop j (numerator) is relative to the overall price index (denominator), which is a CES aggregate of price indexes across crops.

If $\eta > \theta$, the nested logit shocks have the interpretation that farmers maximize profits by choosing a crop and an exporter conditional on each crop (McFadden 1978). Although the theory does not require $\eta > \theta$, the data will turn out to satisfy this condition.

As η increases, the price becomes more important in determining whether a farmer chooses exporter i , conditional on choosing crop j . In the limit, as $\eta \rightarrow \infty$, the entire market goes to the exporter with an infinitesimally higher price than the other exporters. As η decreases, the price becomes less important. In the limit, as $\eta \rightarrow 0$, the entire market only goes to an exporter with an *infinitely* higher price. Similarly, as θ decreases, the price index becomes less important in determining whether a farmer chooses crop j . As $\theta \rightarrow 0$, even a crop with a low

⁸The joint distribution is $F(\nu_{11}, \dots, \nu_{N(M)M}) = \exp \left[- \sum_j \left(\sum_{i(j)} e^{-(1+\eta)\nu_{ij}} \right)^{\frac{1+\theta}{1+\eta}} \right]$.

price index will attract some farmers. As θ increases, the price index becomes more important.

The discrete choice framework offers intuitive interpretations of the parameters θ and η . First, θ governs the correlation of crop-specific shocks. The higher is θ , the more correlated are the farmer’s productivity draws across crops. Since her idiosyncratic productivity for two different crops is likely to be similar, the prices of the crops will determine her choice. This will be the case the land is suitable for growing many different crops, so that there is little heterogeneity in productivity across crops. In Section III.C, I relate my estimates of θ to a large literature that estimates this heterogeneity directly, e.g. Costinot et al. (2016).

Similarly, η governs the correlation of exporter-specific shocks. The higher is η , the more correlated are the farmer’s draws across exporters within a crop. Since her idiosyncratic proximity to two different exporters is likely to be similar, the prices they offer will be more important. If η is high, farmers will be able to reach many different exporters, and there will be little heterogeneity in the cost of accessing exporters. In Section III.C, I relate my estimates of η to a large literature that estimates trade costs directly, e.g. Sotelo (2020).

II.C Exporter price setting

Aggregating across farmers allows us to derive the nested CES supply curve faced by exporter i of crop j :

$$(1) \quad q_{ij} = \left(\frac{p_{ij}}{p_j}\right)^\eta \left(\frac{p_j}{P}\right)^\theta \frac{Y}{P}$$

where $p_j = \left(\sum_{i(j)} p_{ij}^{1+\eta}\right)^{\frac{1}{1+\eta}}$ is the price index for crop j , $P = \left(\sum_j p_j^{1+\theta}\right)^{\frac{1}{1+\theta}}$ is the overall price index, and $Y = \sum_{i,j} p_{ij} q_{ij}$ is total farmer income.⁹

The CES aggregation offers additional interpretations of θ and η . First, θ is the elasticity of substitution across crops in the supply function. The higher is θ , the more substitutable are different crops from the point of view of farmers. In a dynamic setting, this would correspond to higher rates of farmer switching across crops in response to price changes. Similarly, η is the elasticity of substitution across exporters within a crop. The higher is η , the more substitutable are

⁹See the appendix for a full derivation.

exporters from a farmer's point of view, and the more frequently a farmer would switch exporters within a crop. Finally, $\eta > \theta$ has the natural interpretation that exporters are more substitutable within crops than across crops from a farmer's point of view.

Each product j is exported by a set of exporters, which I take to be exogenous. Exporter i purchases q_{ij} units of crop j from farmers, combines them with m_{ij} units of other inputs, and exports x_{ij} units of the finished product. His production function is

$$x_{ij} = z_{ij} q_{ij}^\alpha m_{ij}^{1-\alpha}$$

where $z_{ij} \sim H$ is an idiosyncratic productivity term. This is the only source of ex-ante heterogeneity across exporters within a given product.

Exporters of product j exert market power over farmers, which I model as Cournot or Bertrand competition for crops. When deciding what quantity to purchase (Cournot) or what price to offer (Bertrand) for a crop, exporters form expectations about how farmers respond. In other words, they internalize the upward sloping crop supply curve in Equation 1: each additional unit they purchase increases the price of every other unit. Because Cournot competition yields intuitive expressions for farmer shares at the crop level (see Equation 8), I present the equilibrium under Cournot competition here and show the equilibrium under Bertrand competition in the appendix.

The domestic price of other inputs, p_j^m , and the international price of output, p_j^x , are exogenous. Each exporter maximizes profits

$$\max_{q_{ij}, m_{ij}} \{ p_j^x x_{ij} - p_{ij} q_{ij} - p_j^m m_{ij} \}$$

subject to the supply curve in Equation 1. The first order condition for crops, q_{ij} , can be written:

$$(2) \quad \text{farmer share}_{ij} = \frac{p_{ij} q_{ij}}{p_j^x x_{ij}} = \alpha \times \underbrace{\left(1 + \frac{1}{\epsilon_{ij}} \right)^{-1}}_{\text{markdown}}$$

where $\frac{1}{\epsilon_{ij}} \equiv \frac{\partial \log p_{ij}}{\partial \log q_{ij}}$ is the (inverse) price elasticity of crop supply.

Equation 2 says that the farmer share defined in Section I.C.2 depends on two things: value added (captured by α) and market power (captured by $\frac{1}{\epsilon_{ij}}$).

Under perfect competition, $\frac{1}{\epsilon_{ij}} = 0$, so that the farmer share of exporter revenue equals the output elasticity of crops, α . When the exporter has market power, he internalizes the upward sloping supply of crops, $\frac{1}{\epsilon_{ij}} > 0$, and the farmer share is “marked down” from the perfectly competitive level. The steeper the supply curve faced by the exporter (higher $\frac{1}{\epsilon_{ij}}$), the more market power he has, the wider the markdown, and the lower the farmer share. Alternatively, the more value the exporter adds to the crop (lower α), the lower the farmer share. These are exactly the two explanations for low farmer shares discussed in Section I.C.2.

Given Cournot competition between exporters trying to procure crop j and the supply curve in Equation 1, the supply elasticity has the following closed form expression:

$$(3) \quad \frac{1}{\epsilon_{ij}} = \frac{1}{\eta}(1 - s_{ij}) + \frac{1}{\theta}s_{ij}$$

where $s_{ij} = \frac{p_{ij}q_{ij}}{\sum_{i(j)} p_{ij}q_{ij}}$ is the relative size of exporter i in crop j as defined in Section I.C.3. In other words, the supply elasticity, ϵ_{ij} , is the weighted harmonic mean of the elasticity of substitution across crops, θ , and across exporters, η , where the relative sizes of exporters form the weights.¹⁰ Substituting into Equation 2 yields the equilibrium farmer share:

$$(4) \quad \text{farmer share}_{ij} = \alpha \times \left[1 + \frac{1}{\eta}(1 - s_{ij}) + \frac{1}{\theta}s_{ij} \right]^{-1}$$

If $\eta > \theta$, Equation 4 implies a negative relationship between the farmer share and the relative size of the exporter, precisely the relationship documented in Section I.C.3. The elasticity of substitution across crops, θ , and across exporters, η , determine the strength of this relationship. Equation 4 therefore forges a connection between my stylized facts about agricultural value chains and my theory of crop choice and exporter market power.

¹⁰This is analogous to [Atkeson and Burstein \(2008\)](#), where the exporter-specific *demand* elasticity is a weighted harmonic mean of the elasticities of substitution across and within nests from the point of view of *consumers* and the weights are determined by exporter market shares of the *output* market.

II.D Monopsony power in equilibrium

An equilibrium of the model is defined as follows:

Definition: Given a set of international prices for output $\{p_j^x\}_j$, domestic prices for other inputs $\{p_j^m\}_j$, and parameters $\{\alpha, \eta, \theta\}$, an *equilibrium* is a vector of relative exporter sizes $\{s_{ij}\}_{i,j}$ consistent with farmer optimization (Equation 1) and exporter optimization (Equation 4).

To provide intuition on how monopsony power works in equilibrium, fix the elasticities of substitution η and θ . As s_{ij} increases toward 1, the substitutability across crops, θ , receives more weight in exporter decisions. In contrast, as s_{ij} decreases toward 0, the substitutability across exporters within a crop, η , receives more weight. Since $\eta > \theta$, the supply elasticity ϵ_{ij} decreases as s_{ij} increases. Larger exporters face steeper supply curves and pay farmers a lower share of export revenue. Intuitively, when a single exporter dominates the market for a given crop, farmers can only switch to other crops. Since it is harder for farmers to plant a new crop than to find a new exporter in the same crop than to plant a new crop ($\eta > \theta$), farmers will be less sensitive to prices than if the exporter had a smaller market share. The less price-sensitive are farmers, the more market power the exporter can exert.

Now, fix the size of the exporter. As substitutability across crops, θ , decreases, so does the supply elasticity, ϵ_{ij} . All exporters face steeper supply curves and pay farmers a lower share of export revenue. Intuitively, it has become harder for the farmer to switch to other crops. As a result, prices will play a smaller role in farmer decisions, so that supply will be less elastic and exporters will have more market power. A similar argument holds for substitutability across exporters within a crop, η .

These predictions are summarized in the following proposition:

Proposition: Crop supply becomes less elastic, exporter market power increases, and the crop-level farmer share falls as s_{ij} increases, θ decreases, or η decreases.

III Estimation

In the model, two key elasticities govern market power: the elasticity of substitution across crops, θ , and the elasticity of substitution across exporters within a crop, η . In this section, I estimate these elasticities using exporter responses to

international demand shocks. I also conduct validation exercises.

III.A Identification using pass-through of demand shocks

To make the connection between theory and data more explicit, take logs on both sides of Equation 4. In addition, let the log output elasticity vary by crop. Finally, take a linear approximation of the log markdown. This yields the regression equation in Column 3 of Table 3:

$$(5) \quad \log(\text{farmer share}_{ij}) = \underbrace{\log \alpha_j + \log \frac{\eta}{1 + \eta}}_{\delta_j} - \underbrace{\frac{\eta}{1 + \eta} \left(\frac{1}{\theta} - \frac{1}{\eta} \right)}_{-\beta} s_{ij} + u_{ij}$$

where u_{ij} captures classical measurement error. The size of the coefficient, β , is informative of the difference between η and θ . However, I cannot disentangle them with this regression alone, as the fixed effect, δ_j , contains both η and α_j . This is a well-known issue in the markup literature (De Loecker and Warzynski, 2012), typically addressed by estimating the production function and backing out market power. Instead, I use the structure of the model to estimate η and θ directly.

Consider what happens when there is a sudden increase in the international price for exporter i of crop j . In order to expand exports and meet the growing demand, he must first purchase more crops from farmers by offering a higher price. However, because he has market power and internalizes the upward sloping supply curve for crops, he knows that each additional unit raises the price of every other unit. As a result, he expands crop purchases by less than if his supply curve were flat. The more market power he has, the steeper his supply curve, and the lower the pass-through of the demand shock to farmer income.¹¹

In the appendix, I show that the pass-through of a shock to the international price of crop j , $\Delta \log p_j^x$, to the farmer price offered by exporter i , $\Delta \log p_{ij}$, takes

¹¹This is analogous to a monopolist who faces a sudden decrease in marginal cost but does not pass it through to consumers.

the following form (holding fixed the behavior of other exporters):

$$(6) \quad \rho(s_{ij}) \equiv \frac{\Delta \log p_{ij}}{\Delta \log p_j^x} = \left[1 + \frac{(\frac{1}{\theta} - \frac{1}{\eta})s_{ij}(1 - s_{ij})(1 + \eta)}{1 + \frac{1}{\eta}(1 - s_{ij}) + \frac{1}{\theta}s_{ij}} \right]^{-1}$$

Clearly, $\eta > \theta$ implies that $\rho < 1$, so that pass-through is incomplete under market power. In the appendix, I show that ρ is also decreasing in s_{ij} under this condition. Equation 6 implies that for a given change in international prices the corresponding change in farmer price will be smaller for relatively large exporters. This reflects the intuition that pass-through declines with relative exporter size and forms the basis of my estimation procedure.

In practice, strategic interaction among exporters implies that I *cannot* hold fixed the behavior of other exporters. To illustrate, suppose a relatively large exporter purchases more crops from farmers in response to an idiosyncratic demand shock. This acts as a negative supply shock to the remaining exporters, so that they purchase fewer crops from farmers. This, in turn, acts as a *positive* supply shock to the large exporter. The large exporter's desired increase in crop quantity therefore requires a smaller price increase than suggested by his supply curve prior to the shock. Strategic interaction thus implies that pass-through declines *more* steeply with exporter size, so that I cannot estimate η and θ directly from Equation 6.

III.B Estimation in the presence of strategic interaction

Because of strategic interaction, I recover the parameters of the model through indirect inference. I estimate all parameters jointly, but outline the estimation procedure separately for each group of parameters. Appendix C.2 provides further details.

III.B.1 Estimating η and θ

In order to take Equation 6 to the data, I estimate the following pass-through regression:

$$(7) \quad \Delta \log p_{ij,t} q_{ij,t} - \Delta \log x_{ij,t} = \delta_{jt} + \beta s_{ij,t-1} + \gamma \Delta \log p_{ij,t}^x + \zeta s_{ij,t-1} \times \Delta \log p_{ij,t}^x + u_{ij,t}$$

where $\varepsilon_{ij,t}$ is an error term. The coefficient γ measures the average pass-through

of the demand shock, while the coefficient ζ measures how pass-through varies with exporter size. As discussed above, these coefficients are informative of the elasticities η and θ . However, because of strategic interaction among exporters, I use the full structure of the model to back out the elasticities from pass-through coefficients.

I proceed in several steps: (1) estimate Equation 7 in the actual data, (2) simulate Equation 7 in the model, (3) pick η and θ so that the coefficients γ and ζ from the model match their counterparts in the data. In addition to being tractable, this procedure mitigates the concern with under-reporting of purchases from farmers, as only differential *changes* in under-reporting among exporters of different sizes would threaten the estimates.

In order to estimate Equation 7 in the data, I first construct the demand shocks. I follow a standard shift-share specification combining exporter trade shares from my microdata with shifts in international prices from COMTRADE (Gaulier and Zignago, 2010):

$$\Delta \log p_{ijt}^x = \sum_d \lambda_{ijd,t-1} \Delta \log p_{jdt}^x$$

where d indicates a destination country, $\lambda_{ijd,t-1}$ is the share of exporter i 's sales to that country, and $\Delta \log p_{jdt}^x$ is the log change in price for imports of product j in the destination country (excluding imports from Ecuador). Figure A2 in the appendix plots the distribution of the shocks.

Table 4 displays the results of pass-through regressions using these shocks. Column 1 shows the baseline specification from Equation 7. Column 2 includes product and year fixed effects to control for systematic differences across products and years. Column 3 controls for time-varying exporter characteristics, as in Table 3. The coefficients, denoted $\hat{\gamma}$ and $\hat{\zeta}$, are consistent with the predictions in Section III.A. Pass-through is incomplete ($\hat{\gamma} < 1$), and it decreases with relative exporter size ($\hat{\zeta} < 0$). The magnitudes in Column 3 imply that the largest exporters increase farmer prices by only $\frac{.355-.239}{.355} = 32.7\%$ as much as the smallest exporters following an international price shock.

To estimate Equation 7 in the model, I proceed in several steps (see Appendix C.1 for further details). First, I draw the productivity of each exporter from an exogenous distribution. For each guess of η and θ , I solve the model. Next, I shock

each exporter by drawing from the distribution of international price changes. I solve the model again to create a simulated panel. Finally, I estimate Equation 7 using the simulated panel. The resulting pass-through coefficients, denoted $\gamma(\eta, \theta)$ and $\zeta(\eta, \theta)$, are functions of η and θ .

I pick η and θ so that the pass-through coefficients estimated from the simulated data match the coefficients estimated from the actual data and reported in Table 4:

$$(\hat{\eta}, \hat{\theta}) = \arg \min_{\eta, \theta} \left\{ \|\hat{\gamma} - \gamma(\eta, \theta)\| + \|\hat{\zeta} - \zeta(\eta, \theta)\| \right\}$$

Table 4: Exporter responses to price shocks

	$\Delta \log pq - \Delta \log x$	$\Delta \log pq - \Delta \log x$	$\Delta \log pq - \Delta \log x$
	(1)	(2)	(3)
s	0.061 (0.054)	0.073 (0.068)	0.073 (0.073)
$\Delta \log p^x$	0.228 (0.118)	0.354 (0.124)	0.355 (0.124)
$s \times \Delta \log p^x$	-0.093 (0.256)	-0.226 (0.268)	-0.239 (0.269)
FE	No	Yes	Yes
Controls	No	No	Yes
Observations	767	767	767
R ²	0.008	0.049	0.052

Notes: Table summarizes price pass-through regressions. Dependent variable is the change in log farmer price, defined as the change in log payments to farmers minus the change in log quantity exported. Independent variables are the change in the log international price, calculated from COMTRADE data using a shift-share approach described in the text, the lagged exporter size, and their interaction. Column 3 controls include changes in the wage bill and payments to non-farm suppliers, and an indicator for exporters with lagged market share less than 1%. Clustered standard errors are shown in parentheses.

III.B.2 Estimating α

Aggregating 4 across exporters yields an intuitive expression for the crop-level farmer share:

$$(8) \quad \text{farmer share}_j = \alpha \times \left[1 + \frac{1}{\eta} \left(1 - HHI_j \right) + \frac{1}{\theta} HHI_j \right]^{-1}$$

where $HHI_j \equiv \sum_{i(j)} s_{ij}^2$ is the sum of squared exporter sizes, also known as the Herfindahl-Hirschman Index of market concentration. Equation 8 implies that the lower the effective number of exporters for a given crop (higher HHI), the lower the overall farmer share.

I pick α so that the overall farmer share generated by the model matches the farmer share observed in the data. For each guess of α and the other parameters, I solve the model and calculate the crop-level farmer share from Equation 8, taking HHI_j is taken from the simulated data. Let $\phi(\alpha)$ denote the average farmer share. I pick α so that $\phi(\alpha)$ matches its counterpart in the data, denoted $\hat{\phi}$ and reported in Figure 1b:

$$\hat{\alpha} = \arg \min_{\alpha} \|\hat{\phi} - \phi(\alpha)\|$$

III.B.3 Other parameters

I assume that (log) exporter productivity, $\log z$, and price shocks, $\Delta \log p^x$, follow normal distributions:

$$\log z \sim N(\mu_z, \sigma_z^2) \text{ and } \Delta \log p^x \sim N(\mu_d, \sigma_d^2)$$

For exporter productivity, I choose (μ_z, σ_z^2) to match the distribution of log exporter revenue in the data.¹² For demand shocks, I choose (μ_d, σ_d^2) to match the distribution of log changes in international prices in the data. Finally, the number of crops, M , and the number of exporters for each crop, $\{N_j\}_j$, are chosen to match the histograms in Figure 1a.

III.B.4 Parameter estimates

Table 5 summarizes estimates of the model parameters under Cournot competition.¹³ The elasticities of substitution across exporters, η , and across crops, θ , are small, indicating that exporters face steep supply curves and exercise market power over farmers. The output elasticity of crops, α , is large relative to the farmer share, further indicating a high degree of market power. I explore the economic meaning of these estimates in detail below.

¹²In the appendix, I show how one could estimate the productivity distribution non-parametrically.

¹³See the appendix for estimates under Bertrand competition.

Table 5: Parameter estimates

Parameter	Estimate	Moment	Value
(a) Key parameters			
η	1.93	Baseline pass-through, $\hat{\gamma}$	0.35
θ	0.40	Decline in pass-through with size, $\hat{\zeta}$	-0.23
α	0.45	Average farmer share, $\hat{\phi}$	0.24
(b) Other parameters			
μ_z	13.98	Quantiles of log exporter revenue	
σ_z	2.27		
μ_d	0.02	Quantiles of log price changes	
σ_d	0.11		
M	157	Number of crops	
N_j	1-10	Number of exporters per crop	

III.C Model validation

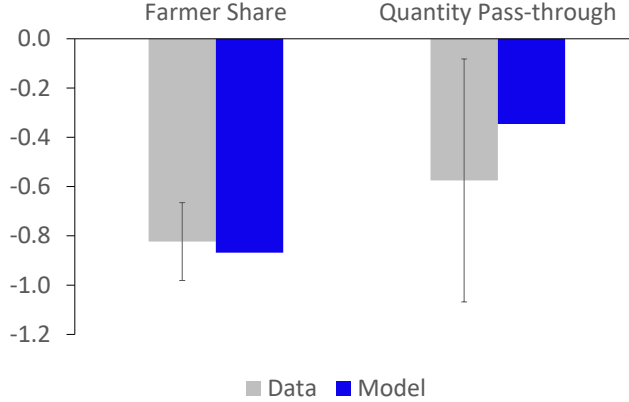
I validate the model by (a) comparing moments not targeted in the estimation procedure between the model and the data and (b) comparing the heterogeneity in production and transport costs implied by the model with estimates from the agricultural trade literature.

III.C.1 Internal validation

I validate the model internally using two non-targeted moments, which are functions of exporter size. The first moment is the negative relationship between farmer shares and exporter size, which I documented in Table 3. Although the average farmer share was targeted in estimation, the relationship between farmer shares and exporter size was not. I regress log farmer shares on exporter size in the simulated data and compare the coefficient on exporter size to Column 1 of Table 3. Column 1 of Figure 3 indicates that the relationship in the model is somewhat steeper than in the data, but the two coefficients are not statistically distinguishable. In the appendix, I estimate an overidentified version of the model which matches this coefficient in addition to the coefficients from the pass-through regression, and obtain similar estimates.¹⁴

¹⁴The precision of this coefficient helps reduce the standard errors of the parameters, which I also calculate in the appendix.

Figure 3: Non-targeted functions of exporter size



Notes: Figure plots non-targeted moments as a function of exporter size. First column shows the coefficient on exporter size from a farmer share regression (see Table 3). Second column shows the coefficient on the interaction between log international price changes and exporter size from a quantity pass-through regression (see Table A5). Grey bars indicate point estimates in the actual data. Black lines indicate standard error bars. Blue bars indicate point estimates in the simulated data.

The second moment is the negative relationship between quantity pass-through and exporter size. In the appendix, I prove that the pass-through of international price changes to farmer *quantities* in the model also declines with exporter size. I confirm this prediction in Table A5. The decline in price pass-through was targeted in estimation, but the decline in quantity pass-through was not. I regress log changes in quantity on log changes in international price, lagged exporter size, and their interaction and compare the coefficient on the interaction term to Column 1 of Table A5. Column 2 of Figure 3 indicates that quantity pass-through declines less steeply with exporter size in the model, but the two coefficients are not statistically distinguishable. However, the relationship is not precisely estimated in the data.

III.C.2 External validation

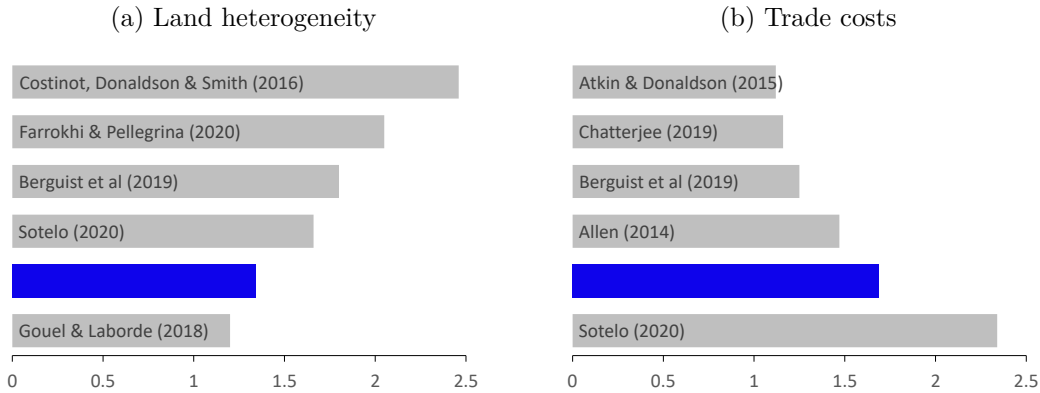
I validate the model externally by comparing my estimates of θ and η to those implied by the literature on agricultural production and trade in developing countries. Recall the interpretation of θ in Section II as a measure of land heterogeneity: the higher is θ , the less heterogeneous is the land, and the more suitable it is for producing a variety different crops. Several studies estimate this

heterogeneity directly using data on land use and yields across crops. In the appendix, I show how to calculate the land heterogeneity implied by my estimate of θ . Figure 4a compares this value to those from the literature. They are generally larger than my estimate of 1.35, indicating a smaller degree of heterogeneity than in my setting. Importantly, I include the largest number of distinct products, which may explain why I find more heterogeneity. Consistent with this explanation, Gouel and Laborde 2021 is both the only other study to include animal products and the only study to find lower heterogeneity. Sotelo 2020 finds a value similar to mine in Peru, the most agroclimatically similar country to Ecuador among those studied.

Now, recall the interpretation of η in Section II as a measure of heterogeneity in costs of reaching different exporters. To the best of my knowledge, no study estimates this heterogeneity directly in an agricultural setting. However, a large literature estimates iceberg trade costs across space. I show in the appendix that under some assumptions, my estimate of η implies an average iceberg trade cost of 1.69. Figure 4b shows the average estimated trade cost for several studies that focus on agriculture in developing countries. They are generally smaller than my estimate, indicating lower trade costs on average. The most comparable study is Chatterjee 2020, where trade costs allow local intermediaries in India to exercise market power over farmers. Lacking the kind of spatial data he uses to define each geographic market, I define a single market for each crop, which may explain why my estimates are larger. On the other hand, my estimates are *smaller* than in Sotelo 2020, which uses spatial data from Peru, the country most geographically similar to Ecuador among those studied.¹⁵

¹⁵The countries represented are Ethiopia, Nigeria, India, Ghana, Philippines, and Peru.

Figure 4: External validation of θ and η



Notes: Panel A plots estimates of land heterogeneity parameters from selected papers in grey, and the corresponding value implied by $\hat{\theta}$ in blue. See text of Appendix C.9 for conversion details. See Table A6 for source details. Panel B plots estimates of iceberg trade costs from selected papers in grey, and the corresponding value implied by $\hat{\eta}$ in blue. See text of Appendix C.10 for conversion details. See Table A7 for source details.

IV Measuring exporter monopsony power

Equipped with estimates of η and θ , I can now quantify monopsony power across the agricultural sector.¹⁶ First, I use the actual data to calculate the implied markdowns faced by farmers in Ecuador. Second, I conduct simulations to compare the level of farmer income between the estimated model and a counterfactual in which exporters behave competitively, rather than strategically. Third, I decompose the aggregate effect of market power into different channels. Finally, I discuss alternate explanations for the results.

IV.A Crop markdowns in Ecuador

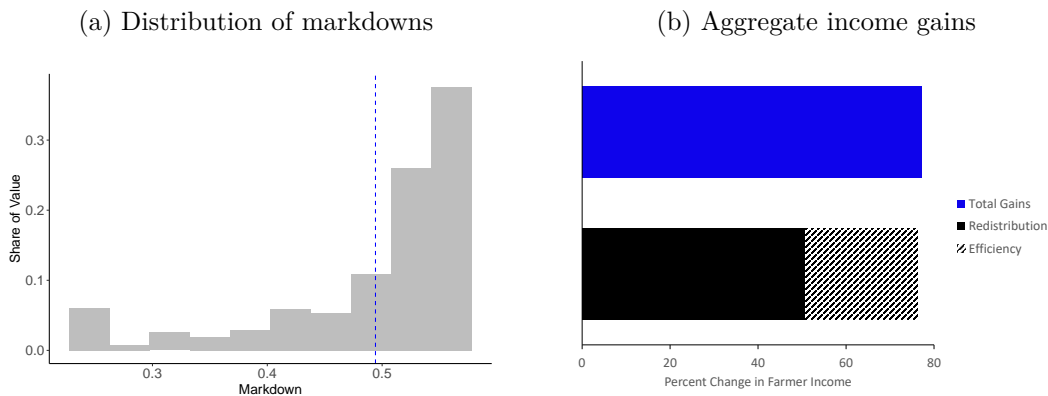
To explore the microeconomic impacts of market power, I combine parameter estimates with value chain data in order to measure how much farmer prices are marked down from their marginal revenue products. Recall from Section II that the equilibrium markdown can be written as follows:

$$(9) \quad \text{markdown}_{ij} = \left[1 + \frac{1}{\eta}(1 - s_{ij}) + \frac{1}{\theta}s_{ij} \right]^{-1}$$

¹⁶I focus on Cournot competition and present results under Bertrand competition in the appendix.

Figure 5a plots the distribution of markdowns, obtained by plugging the estimated η and θ and observed s_{ij} into Equation 9. The weighted average is 0.49, implying that farmers receive around half of their marginal revenue product. While the majority of exporters pay farmers 50-60% of their marginal product, some exporters, including of important crops like coffee and palm, pay less than 30%.

Figure 5: Measurements of monopsony power



Notes: Panel A plots the distribution of markdowns across exporters. Markdowns were calculated using s_{ij} from the data and the estimated η and θ . Dashed blue line indicates average markdown of 49%. Panel B shows percent increase in farmer income between the estimated model with monopsony power and a counterfactual model with perfect competition. Blue area indicates total gains. Black area indicates gains from redistributing exporter profits to farmers. Shaded area indicates gains from greater efficiency.

IV.B Farmer income under perfect competition

To explore the aggregate implications of market power, I consider a counterfactual economy in which exporters act competitively, rather than strategically. Under perfect competition, exporters still face upward sloping crop supply curves, whose shapes are determined by the parameters η and θ . However, they do not internalize their influence over the price, but rather perceive a perfectly elastic supply curve, $\frac{1}{\epsilon_{ij}} = 0$. Crop prices are no longer marked down from their marginal revenue product, so that farmers receive the competitive farmer share, α .

This has two effects. First, farmers earn higher income for supplying the same crop to the same exporter, since markdowns are eliminated across the entire sector. This is a pure redistribution from exporters to farmers. However, there

are also efficiency gains. In my theory of crop choice, farmers trade off the price of a given exporter and a given crop with their idiosyncratic shock for producing that crop and supplying that exporter. This implies that some farmers do not produce the crop in which they are most productive, simply because its price index is too low. Conditional on a crop, some farmers do not supply the exporter that is closest to them, simply because his price is too low. Removing market power lessens this tradeoff and allows some farmers to produce their best crop and supply their closest exporter. These are efficiency gains.

To quantify these channels, I first simulate the model with and without market power. The total impact of market power is the log difference in farmer income between the two scenarios. To measure the gains from redistribution, I calculate farmer income using quantities from the market power baseline and prices from the perfect competition counterfactual. To measure efficiency gains, I do the opposite, using market power prices and perfect competition quantities:

$$\underbrace{\log \sum_i p_{ij}^{PC} q_{ij}^{PC} - \log \sum_i p_{ij}^{MP} q_{ij}^{MP}}_{\text{total gains}} = \underbrace{\log \sum_i p_{ij}^{PC} q_{ij}^{MP} - \log \sum_i p_{ij}^{MP} q_{ij}^{MP}}_{\text{redistribution}} + \underbrace{\log \sum_i p_{ij}^{MP} q_{ij}^{PC} - \log \sum_i p_{ij}^{MP} q_{ij}^{MP}}_{\text{efficiency}} + \text{interactions}$$

where the superscript MP denotes the baseline with market power and PC denotes the counterfactual with perfect competition.

Figure 5b displays the results of the decomposition. I find that farmer income would be 77.1% higher in the absence of market power. Redistribution from exporters to farmers increases income by 50.7%, accounting for almost two thirds of the gains.¹⁷ Greater efficiency accounts for the remaining third, a 25.6% increase in farmer income.

Although all farmers gain from perfect competition, the gains are not equally shared. In the appendix, I show how increases in farmer income vary with the baseline level of crop market concentration, HHI_j . Gains range from around 54%

¹⁷In terms of welfare, redistribution represents a gain for farmers and a loss for exporters. If exporter profits are rebated to farmers, the overall welfare gain may be small or even negative. However, this assumption in unreasonable in this context.

in relatively competitive crops, such as bananas, to 122% in the least competitive crops, including cocoa.

IV.C Alternative explanations

Market power reduces both the variance and the mean of farmer income. Section IV.B suggest that farmer income is lower when exporters have more market power, while Section III.A implies that farmer income is less responsive to international demand shocks. A competing explanation for this mean-variance trade-off is that exporters insure farmers against shocks. In the appendix, I conduct two exercises, which suggest that the results are not fully explained by insurance.

First, I test the hypothesis that exporters insure farmers against both positive and negative shocks. The decline in farmer shares and pass-through with exporter size implies that larger exporters offer more insurance. In this case, exporters should help smooth the effects of both positive and negative shocks. Instead, I find that pass-through declines with exporter size only for *positive* shocks. In contrast, negative shocks have similar pass-through among small and large exporters. If insurance were the sole mechanism at play, the pass-through of international price shocks would not depend on sign of the shock.

Next, I calculate the level of farmer risk aversion that rationalizes the estimated mean-variance trade-off. In the data, farmer income increases 42% on average following a 100% increase in the international price shock. Under perfect competition, this pass-through would equal 100%. Assuming Cournot competition at baseline, farmer income would be 71% higher under perfect competition. A back-of-the-envelope calculation shows that if farmers are indifferent between perfect competition and market power, their coefficient of relative risk aversion must be greater than 5. This corresponds to the highest category of risk aversion estimated among farmers in Ethiopia (Yesuf and Bluffstone, 2009). If insurance were the only mechanism at play, farmers would have to be unreasonably risk averse to prefer the baseline equilibrium.

Market power generates massive profits for exporters. Section IV.B suggests that farmer income would be substantially higher simply from reallocating these profits. Another possible explanation is that these profits represent fixed costs of sourcing and exporting crops. This does not affect my estimation strategy, since the pass-through regressions effectively difference out exporter-specific time-

invariant unobservables. However, it may affect my estimates of the aggregate effect of monopsony power, as exporters exit the market under perfect competition. In the appendix, I conduct two exercises, which suggest that the results are not fully explained by fixed costs.

First, I plot the distribution of profit rates implied by the data and parameter estimates. Profits range from 5% of exporter revenue to more than 50%. Larger exporters have significantly higher profit rates, even within a crop. If these profit rates purely reflected exporter-specific fixed costs, investments in sourcing and distribution would exhibit decreasing returns to scale. This is inconsistent with evidence on intermediation costs (Ganapati, 2021).

Second, I compute an upper bound for the fixed costs of exporting bananas – the most competitive crop. This back-of-the-envelope calculation suggests that the fixed cost is no greater than 52% of the farmer income paid by the smallest exporter. Even with this upper bound, fixed costs do not fully offset the 54% increase in farmer income across the banana sector under perfect competition.

V Pro-farmer trade policies

Perfectly competitive markets are a useful benchmark, but they are a far cry from the policies currently in place to curtail market power around the world. In this section, I use the model to examine two of the most common such policies: Fair Trade certifications and price floors. I model Fair Trade as a perfectly competitive exporter in each crop and show that this raises farmer income both directly and indirectly, by reducing the market power of other exporters. In contrast, a price floor in each crop raises farmer income, but *increases* the market power of some exporters, partially offsetting the direct effect. As a result, Fair Trade is more effective in raising farmer incomes.

V.A Fair Trade

Fair Trade refers a series of product certifications designed to foster the sustainable production of commodities. Certified commodities include flowers, bananas, sugar, coffee, cocoa, and other fruits and vegetables. Similar certifications exist for fish and meat. In order for a product to be certified, both exporters and producers must meet certain criteria, which are typically enforced by an NGO. Exporters agree to pay a minimum price that covers the cost of sustainable farm-

ing, as well as a Fair Trade premium typically earmarked for further investment in farming communities. In return, farmers guarantee safe working conditions and sound environmental practices. Because these guarantees are costly, only a subset of producers are Fair Trade certified. For coffee – the largest product in the Fair Trade market – less than 40% of available quantity is certified.

Fair Trade exporters are able to pay farmers higher prices because consumers are willing to pay a premium for Fair Trade branded products (Hainmueller, Hiscox and Sequeira 2015). Some Fair Trade exporters are even owned by farmers via cooperatives, who overcome the fixed cost of organizing with assistance from an NGO. These cooperatives have an incentive to internalize markdowns and pay higher prices (Bacon, Mendez and Stuart 2008).

I focus on the competitive effect of Fair Trade and abstract from selection, non-monetary benefits, costs of certification, downstream demand, and fixed costs of exporting. Dragusanu, Giovannucci and Nunn (2014) offer a comprehensive survey of Fair Trade, including these aspects.¹⁸ Following Podhorsky 2015, I model the potential impact of Fair Trade in Ecuador by introducing a perfectly competitive exporter in each crop market. This tractably and flexibly captures many of the ways Fair Trade works in practice. In particular, I show in the appendix that the model is isomorphic to (a) buyers specifying a minimum Fair Trade price and (b) farmers exporting directly through a cooperative.

The Fair Trade exporter faces the same supply curve as other exporters of a given product, but pays farmers their marginal revenue product, which is higher than the oligopsony price. A new exporter would increase competition and force other exporters to raise prices, even if he behaved strategically. That he instead behaves competitively, and therefore pays a higher price conditional on his productivity, further raises prices. Fair Trade therefore has a positive direct and indirect effect on prices. These effects reflect the primary goals of Fair Trade: increasing prices and improving bargaining power among farmers. Furthermore, their importance has been documented both theoretically (Podhorsky 2015) and empirically (Dragusanu and Nunn 2018).

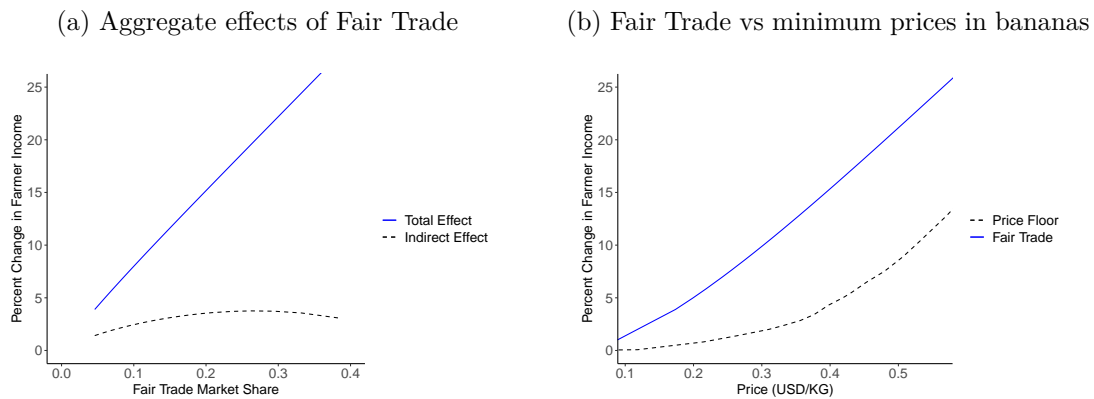
The overall effect of Fair Trade depends on the productivity of the new exporter. The more productive he is, the higher the price he can offer to farmers,

¹⁸I provide a brief discussion in the appendix.

and the more market share he pulls away from exporters with market power. Figure 6 summarizes how the increase in farmer income varies with the market share of the Fair Trade.¹⁹ The blue solid line shows that a Fair Trade exporter who captures 15% of the market (corresponding to the median productivity level in the data) increases farmer income by 12%. As the new exporter becomes among the most productive in the economy, he captures 40% of the market, and the gains increase to 25%, or about one third of the gains from perfect competition in Figure 5b. These gains are quantitatively similar to causal estimates from the coffee sector (Dragusanu and Nunn 2018), but apply to a much broader range of products.

To get a sense of the indirect and direct effects of the Fair Trade exporter, I estimate how farmer income would change if the new exporter behaved strategically. The dashed black line indicates that the gains from Fair Trade are driven by the direct effect on farmers supplying the Fair Trade exporter.

Figure 6: Effect of Fair Trade on farmer income



Notes: Panel A plots market share of a counterfactual Fair Trade entrant in each crop on the x-axis and resulting percent change in farmer income relative to the baseline model on the y-axis. Solid blue line indicates counterfactual in which the exporter implements a perfectly competitive Fair Trade policy. Dashed black line indicates counterfactual in which the exporter has monopsony power. Panel B plots counterfactual Fair Trade price of bananas in USD/KG on the x-axis and resulting percent change in farmer income relative to the baseline model on the y-axis. Solid blue line indicates counterfactual in which price applies only to the Fair Trade exporter. Dashed black line indicates the counterfactual in which it applies to all exporters.

¹⁹Figure A6 in the appendix shows how the change in farmer income varies directly with the productivity of the Fair Trade exporter.

V.B Minimum prices

A common alternative to Fair Trade is for governments to set a price floor across *all* exporters of a given product.²⁰ Minimum price support is growing, especially for exported commodities in developing countries. Compared to conditional subsidies, these policies are more straightforward to implement, but can create more distortions.

To illustrate the effect of a price floor, consider exporters for whom the minimum price is binding. These exporters move along their supply curves. If they are productive enough that they can still earn profits, they will pay the minimum price and purchase more crops at a lower markdown. If they are not productive enough to earn positive profits moving along their supply curves, they will pay the minimum price and purchase *fewer* crops until the marginal revenue product equals the minimum price. This increases the market power of more productive firms and undoes some of the positive price effects. The strength of these effects depends crucially on the level of the minimum price. If the minimum price is low, most exporters will be able to pay, and the net effect will be positive.²¹ As the minimum price becomes too high, no exporters can afford to pay, and demand contracts so much that farmers may be worse off.

Figure 7b summarizes how farmer income in the banana sector would change under a uniform price floor.²² In 2011, the average farmgate price of exported bananas in Ecuador was around 30 cents/kg. The black dashed line indicates that imposing a price floor at this level across all banana exporters would have almost no effect on farmer income. Instead, imposing a similar price floor only for Fair Trade exporters would increase farmer income by 9%. Fair Trade bananas have experienced substantial growth in Ecuador over the past 10 years, with prices reaching 50 cents in 2021. Although the impact of Fair Trade has not been studied in Ecuador, estimates from Peru suggest that farmer income grew by 20% when the Fair Trade price increased from 30 to 50 cents in the early 2000s (Ruben and Fort, 2012). These results are consistent with Figure 7b.

²⁰In Ecuador, bananas and palm are the only exported products with price floors, but these were not binding during the sample period (Cunha et al. 2019).

²¹This is analogous to a minimum wage increasing employment in the presence of labor market power (Berger et al. forthcoming).

²²Figure A7 in the appendix shows how farmer income would change across all products.

Compared to a universal price floor, Fair Trade does not distort the behavior of smaller exporters (Podhorsky 2015), making it a more effective policy for raising farmer income. However, the effect of each policy depends on the ability of farmers to substitute across crops. In the appendix, I show how the predicted effects of Fair Trade and universal price floors vary with the elasticity of substitution across crops, θ . As it becomes easier for farmers to switch crops, farmer income gains become more similar across the two policies. This is because exporters are more competitive, which reduces both the additional competitive effect of Fair Trade (positive) and the distortionary effect of universal price floors (negative).

VI Conclusion

Recent decades have seen the rise of both market concentration and globalization. Understanding the consequences of concentration is especially important in the agricultural sector in emerging economies, where globalization offers millions of farmers a path out of poverty. In Ecuador, these consequences are large.

To overcome the challenge of measuring inequality in value chains, I link firm-level customs data with domestic firm-firm transactions across the universe of exported cash crops. I document that farmers earn significantly less if they sell to an exporter who dominates the market for a crop. To quantify the importance of exporter monopsony power, I develop a model in which farmers choose a crop to produce and an exporter to supply. The more costly it is for farmers to switch crops or switch exporters within a crop, the more that farmer shares fall with exporter size. The elasticities of substitution across crops and across exporters within a crop are therefore crucial to measuring market power. Together, the theory and data allow me to estimate these elasticities using exporter responses to international price shocks. The estimated model implies that farmers in products as diverse as fruit and fish are paid a fraction of their marginal revenue products.

Despite the prevalence of market power, globalization can still provide farmers a path out of poverty. Fair Trade increases farmer income substantially while avoiding some of the distortions created by universal price floors. However, the effects of these policies depend on the substitutability of crops and exporters. Further work is needed to examine their effects in settings with more substitutability, as well as the effects of other pro-farmer policies.

References

- Adao, Rodrigo, Paul Carrillo, Arnaud Costinot, Dave Donaldson, and Dina Pomeranz**, “Imports, Exports, and Earnings Inequality: Measures of Exposure and Estimates of Incidence,” *Quarterly Journal of Economics*, forthcoming.
- Aksoy, M Ataman and John C Beghin**, *Global agricultural trade and developing countries*, The World Bank, 2004.
- Allen, Treb**, “Information frictions in trade,” *Econometrica*, 2014, 82 (6), 2041–2083.
- Alviarez, Vanessa, Keith Head, and Thierry Mayer**, “Global giants and local stars: How changes in brand ownership affect competition,” 2020.
- Arnoud, Antoine, Fatih Guvenen, and Tatjana Kleineberg**, “Benchmarking Global Optimizers,” Technical Report, National Bureau of Economic Research 2019.
- Atkeson, Andrew and Ariel Burstein**, “Pricing-to-market, trade costs, and international relative prices,” *American Economic Review*, 2008, 98 (5), 1998–2031.
- Atkin, David and Dave Donaldson**, “Who’s getting globalized? The size and implications of intra-national trade costs,” Working paper, NBER 2015.
- Bacon, Christopher M., V. Ernesto Mendez, and Eugenia Flores Gomez Douglas Stuart**, “Are Sustainable Coffee Certifications Enough to Secure Farmer Livelihoods? The Millennium Development Goals and Nicaragua’s Fair Trade Cooperatives,” *Globalizations*, 2008, pp. 259–274.
- Bartkus, Viva Ona, Wyatt Brooks, Joseph P Kaboski, and Carolyn E Pelnik**, “Big Fish in Thin Markets: Competing with the Middlemen to Increase Market Access in the Amazon,” Working paper, NBER 2021.
- Berger, David W, Kyle F Herkenhoff, and Simon Mongey**, “Labor market power,” *American Economic Review*, forthcoming.
- Bergquist, Lauren, Benjamin Faber, Thibault Fally, Matthias Hoelzlein, Edward Miguel, and Andres Rodriguez-Clare**, “Scaling Agricultural Policy Interventions: Theory and Evidence from Uganda,” Working paper 2019.
- Bergquist, Lauren Falcao and Michael Dinerstein**, “Competition and entry in agricultural markets: Experimental evidence from Kenya,” *American Economic Review*, 2020, 110 (12), 3705–47.
- Cardell, N Scott**, “Variance components structures for the extreme-value and logistic distributions with application to models of heterogeneity,” *Econometric Theory*, 1997, pp. 185–213.

- Carrillo, Paul, Dina Pomeranz, and Monica Singhal**, “Dodging the taxman: Firm misreporting and limits to tax enforcement,” *American Economic Journal: Applied Economics*, 2017, 9 (2), 144–64.
- Casaburi, Lorenzo and Tristan F Reed**, “Interlinked Transactions and Competition: Experimental Evidence from Cocoa Markets,” *American Economic Journal: Applied Economics*, forthcoming.
- Chatterjee, Shoumitro**, “Market Power and Spatial Competition in Rural India,” Working paper 2020.
- Cheong, David, Marion Jansen, and Ralf Peters**, *Shared Harvests: Agriculture, Trade and Employment*, UNCTAD, 2013.
- Costinot, Arnaud, Dave Donaldson, and Cory Smith**, “Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world,” *Journal of Political Economy*, 2016, 124 (1), 205–248.
- Cunha, Barbara, Jose Daniel Reyes, and Samuel Jaime Pienknagura**, “Ecuador Trade and Investment Competitiveness Report,” Technical Report, The World Bank 2019.
- Dhingra, Swati and Silvana Tenreyro**, “The Rise of Agribusiness and the Distributional Consequences of Policies on Intermediated Trade,” Working paper, CEPR 2020.
- Dragusanu, Raluca and Nathan Nunn**, “The effects of Fair Trade certification: evidence from coffee producers in Costa Rica,” Working paper, NBER 2018.
- , **Daniele Giovannucci, and Nathan Nunn**, “The economics of fair trade,” *Journal of economic perspectives*, 2014, 28 (3), 217–36.
- FAO**, “The Changing Role of Multinational Companies in the Global Banana Trade,” Technical Report 2014.
- Farrokhi, Farid and Heitor S Pellegrina**, “Global Trade and Margins of Productivity in Agriculture,” Working paper, NBER 2021.
- Ganapati, Sharat**, “The Modern Wholesaler: Global Sourcing, Domestic Distribution, and Scale Economies,” Working paper 2021.
- Gaulier, Guillaume and Soledad Zignago**, “BACI: International Trade Database at the Product-Level. The 1994-2007 Version,” Working Paper 2010-23, CEPII 2010.
- Gouel, Christophe and David Laborde**, “The crucial role of domestic and international market-mediated adaptation to climate change,” *Journal of Environmental Economics and Management*, 2021, 106, 102408.
- Gourieroux, Christian, Alain Monfort, and Eric Renault**, “Indirect inference,” *Journal of applied econometrics*, 1993, 8 (S1), S85–S118.

- Grant, Matthew and Meredith Startz**, “Cutting out the middleman: The structure of chains of intermediation,” Working paper 2021.
- Hainmueller, Jens, Michael J Hiscox, and Sandra Sequeira**, “Consumer demand for fair trade: Evidence from a multistore field experiment,” *Review of Economics and Statistics*, 2015, *97* (2), 242–256.
- Janvry, Alain De, Craig McIntosh, and Elisabeth Sadoulet**, “Fair trade and free entry: can a disequilibrium market serve as a development tool?,” *Review of Economics and Statistics*, 2015, *97* (3), 567–573.
- Loecker, Jan De and Frederic Warzynski**, “Markups and firm-level export status,” *American economic review*, 2012, *102* (6), 2437–71.
- , **Pinelopi K Goldberg, Amit K Khandelwal, and Nina Pavcnik**, “Prices, markups, and trade reform,” *Econometrica*, 2016, *84* (2), 445–510.
- Lowder, Sarah K, Jakob Scoet, and Terri Raney**, “The number, size, and distribution of farms, smallholder farms, and family farms worldwide,” *World Development*, 2016, *87*, 16–29.
- McFadden, Daniel**, “Modeling the choice of residential location,” *Transportation Research Record*, 1978, (673).
- Morlacco, Monica**, “Market Power in Input Markets: Theory and Evidence from French Manufacturing,” Working paper 2020.
- Nakamura, Emi and Dawit Zerom**, “Accounting for incomplete pass-through,” *The review of economic studies*, 2010, *77* (3), 1192–1230.
- Patten, Diana Van and Esteban Mendez-Chacon**, “Multinationals, Monopsony and Local Development: Evidence from the United Fruit Company,” Working paper 2021.
- Podhorsky, Andrea**, “A positive analysis of Fairtrade certification,” *Journal of Development Economics*, 2015, *116*, 169–185.
- Pomeranz, Dina**, “No taxation without information: Deterrence and self-enforcement in the value added tax,” *American Economic Review*, 2015, *105* (8), 2539–69.
- Ruben, Ruerd and Ricardo Fort**, “The Impact of Fair Trade Certification for Coffee Farmers in Peru,” *World Development*, 2012, *40* (3), 570 – 582.
- Rubens, Michael**, “Market structure, oligopsony power and productivity,” Working paper 2021.
- Sotelo, Sebastian**, “Domestic trade frictions and agriculture,” *Journal of Political Economy*, 2020, *128* (7), 2690–2738.
- Wong, Sara A**, *The effects of SPS and TBT measures on banana and pineapple trade in Ecuador*, Espae, 2008.
- Yesuf, Mahmud and Randall A Bluffstone**, “Poverty, risk aversion, and path dependence in low-income countries: Experimental evidence from Ethiopia,” *American Journal of Agricultural Economics*, 2009, *91* (4), 1022–1037.

Appendices for online publication

A Data appendix

A.1 Data construction details

In this section, I provide additional details on the construction of my value chain dataset and discuss some robustness checks.

Section I discusses the three main sources of data: customs records, VAT receipts, and a business registry. A fourth dataset includes matched employee-employer information from 2008-2011. The data are derived from Social Security Tax declarations and record the earnings and employers for each worker and year. I use these data to calculate employment and wage bill for each firm. A fifth dataset includes annual income tax forms filed annual by each firm. There are two types of forms – the F101, which is filed by large firms, and the simplified F102, which is filed by small firms and self-employed individuals. I use these to calculate total sales for each firm, as well as important variables such as the self-employment indicator.

The main text outlines the matching procedure used to construct the supply chain. I begin with exporters in the customs data and merge all of their suppliers from the VAT data. Suppliers in the farming and fishing sectors are considered producers, so they are terminal nodes in the value chain. Suppliers in sectors such as transportation and storage provide other inputs to exporting, so they are also terminal nodes for my purposes. Suppliers in the wholesale sector are potentially domestic intermediaries, so I repeat the matching procedure for them. After performing the procedure 3 times, I find that 90% of crop purchases are made directly by exporters (chain length 1), 9% by a single domestic intermediaries (chain length 2), and less than 1% by chains of domestic intermediaries (chain length 3). As a result, I restrict my analysis to chains of length no greater than 2. This is consistent with the average length of chains for imported goods in Nigeria ([Grant and Startz, 2021](#)).

I classify crops based on the primary 6-digit HS code of the exporter as described in the main text. Alternatively, I can classify crops based on the 5-digit ISIC code of the producer. Because ISIC codes contain multiple HS codes, this yields more aggregated crop categories and hence less concentrated crop markets.

However, these markets still exhibit a high concentration of exporters, low farmer shares, and a declining relationship between farmer shares and exporter size.

I define farmer income as sales of agricultural firms in the main analysis. An alternative approach is to define farmer income as the sum of (a) the wage bill reported by larger farms and (b) the sales reported by smaller farms.²³ This leads to lower estimates of farmer income for a given firm, but allows me to incorporate exporting farms in the cross-sectional facts. In practice, the average farmer share increases, but the negative relationship with exporter size remains.

A.2 Additional statistics

Table A1 summarizes the network across 2-digit products.

Table A1: Value chain statistics by product

2-digit Product	No. Exporters	No. Farmers
Live animals	3	3
Fish and crustaceans	180	8,650
Dairy produce	6	1,406
Other animal products	4	23
Live plants	476	1,153
Vegetables	44	2,162
Fruit and nuts	301	11,301
Coffee, tea, spices	33	2,486
Cereals	22	6,446
Mill products	7	50
Oil seeds	20	159
Vegetable extracts	2	2
Other vegetable products	8	36
Animal or vegetable fats and oils	25	17,909
Meat and fish preparations	43	2,533
Sugars and sugar confectionery	11	3,724
Cocoa and cocoa preparations	77	25,336
Cereal preparations	12	1,299
Vegetable and fruit preparations	47	7,988
Other preparations	14	2,827
Beverages	16	1,157
Waste from the food industries	31	4,159
Tobacco products	16	999

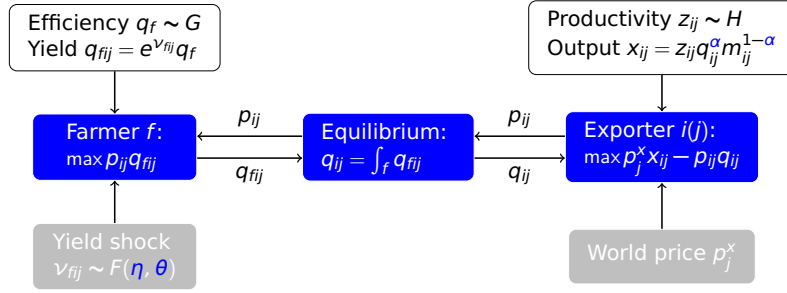
²³Adao, Carrillo, Costinot, Donaldson and Pomeranz (forthcoming) follow this approach for the manufacturing sector in Ecuador.

B Theory appendix

B.1 Model diagram

Figure A1 in summarizes the structure of the model. White boxes depict endowments and technologies, grey boxes depict model shocks, and blue boxes depict optimization and equilibrium conditions. Black arrows signify the flow of goods and payments. Blue text denotes key model parameters. See Section II for details.

Figure A1: Model structure



B.2 Derivation of CES supply curve

The farmer maximizes $y_{ij} = \log p_{ij} + \log q_f + \frac{\nu_{fj}^c}{1+\theta} + \frac{\nu_{fij}^e}{1+\eta}$ across i and j . The maximum satisfies $y_{ij} > y_{kl}$ for all k and l . For any k and l , the terms $\log q_f$ on both sides of the inequality cancel, so that the maximum is independent of farmer capacity.

The expected quantity supplied by farmer f to exporter i of crop j is $q_{fij} = q_f \times \Pr(fij)$. Integrating over farmers yields the total quantity of crop j supplied to exporter i :

$$q_{ij} = \int_0^1 \Pr(fij) q_f dG = \frac{p_{ij}^\eta}{\sum_{i(j)} p_{ij}^{1+\eta}} \frac{(\sum_{i(j)} p_{ij}^{1+\eta})^{\frac{1+\theta}{1+\eta}}}{\sum_j (\sum_{i(j)} p_{ij}^{1+\eta})^{\frac{1+\theta}{1+\eta}}} \underbrace{\int_0^1 p_{ij} q_f dG}_Y$$

Multiplying both sides by p_{ij} and summing across crops and exporters, we have $Y = \sum_{i,j} p_{ij} q_{ij}$, so that Y is total spending by exporters on crops.

Define the crop-level price and quantity indexes

$$p_j = \left(\sum_{i(j)} p_{ij}^{1+\eta} \right)^{\frac{1}{1+\eta}}, \quad q_j = \left(\sum_{i(j)} q_{ij}^{\frac{1+\eta}{\eta}} \right)^{\frac{\eta}{1+\eta}}$$

Substituting above yields the CES supply system for crops

$$q_{ij} = p_{ij}^{\eta} p_j^{\theta-\eta} \underbrace{\left(\sum_j p_j^{1+\theta} \right)^{-1}}_X Y$$

Note that $q_j = p_j^{\theta} X$, which implies that I can write the inverse supply curve

$$p_{ij} = q_{ij}^{\frac{1}{\eta}} q_j^{\frac{1}{\theta} - \frac{1}{\eta}} X^{\frac{1}{\theta}}$$

Finally, define the aggregate price and quantity indexes

$$P = \left(\sum_j p_j^{1+\theta} \right)^{\frac{1}{1+\theta}}, \quad Q = \left(\sum_j q_j^{\frac{1+\theta}{\theta}} \right)^{\frac{\theta}{1+\theta}}$$

Using these definitions and the fact that $q_j = p_j^{\theta} X = p_j^{\theta} \left(\sum_j p_j^{1+\theta} \right)^{-1} Y$, it is straightforward to show that $PQ = Y$. This implies that $X = \frac{Y}{P^{1+\theta}}$. Substituting into the supply curves yields the following expressions:

$$(A1) \quad q_{ij} = \left(\frac{p_{ij}}{p_j} \right)^{\frac{1}{\eta}} \left(\frac{p_j}{P} \right)^{\frac{1}{\theta}} \frac{Y}{P}$$

$$(A2) \quad p_{ij} = \left(\frac{q_{ij}}{q_j} \right)^{\frac{1}{\eta}} \left(\frac{q_j}{Q} \right)^{\frac{1}{\theta}} \frac{Y}{Q}$$

B.3 Bertrand competition

Given Bertrand competition between exporters trying to procure crop j and the supply curve in Equation 1, the supply elasticity has the following closed form:

$$(A3) \quad \epsilon_{ij} = \eta(1 - s_{ij}) + \theta s_{ij}$$

where s_{ij} is the relative size of exporter i in crop j . In other words, the supply elasticity, ϵ_{ij} , is the weighted mean of the elasticity of substitution across crops, θ , and across exporters, η , where the relative sizes of exporters form the weights. Substituting into Equation 2, the equilibrium farmer share is:

$$(A4) \quad \text{farmer share}_{ij} = \alpha \times \left[1 + \frac{1}{\eta(1 - s_{ij}) + \theta s_{ij}} \right]^{-1}$$

Since $\eta > \theta$, Equation A4 implies a negative relationship between the farmer share and the relative size of the exporter, just like Equation 4. Aggregating across exporters yields the crop-level farmer share:

$$(A5) \quad \text{farmer share}_j = \alpha \times \left[1 + \sum_{i(j)} \frac{s_{ij}}{\eta(1 - s_{ij}) + \theta s_{ij}} \right]^{-1}$$

This equation is analogous to 8, but difficult to interpret without an analog to the HHI.

One can show that for any $\eta \neq \theta$, the markdown under Bertrand competition:

$$(A6) \quad \left[1 + \frac{1}{\eta(1 - s_{ij}) + \theta s_{ij}} \right]^{-1}$$

is greater than the markdown under Cournot competition:

$$\left[1 + \frac{1}{\eta}(1 - s_{ij}) + \frac{1}{\theta} s_{ij} \right]^{-1}$$

One can further show that for $\eta > \theta$, the pass-through of an international price change is lower under Cournot. For a given η , θ , and s_{ij} , Bertrand competition clearly implies less market power among exporters.

The implications of Bertrand competition for *estimating* market power are less clear. Given the relationship between pass-through and exporter size in the data, Bertrand competition will yield smaller estimates of η and θ than Cournot competition, indicating steeper supply curves and hence more market power. However, given η , θ , and the distribution of farmer shares in the data, Bertrand competition will also yield smaller estimates of α than Cournot competition, in-

dicating narrower markdowns and hence less market power. These counteracting forces can simultaneously yield lower estimates of the market power parameters η and θ and smaller gains from removing market power.

B.4 Pass-through of international price shocks

Log-linearize around the equilibrium in Equation 4:

$$\Delta \log p_{ij} q_{ij} = \Delta \log p_j^x + \Delta \log x_{ij} - \frac{(\frac{1}{\theta} - \frac{1}{\eta}) s_{ij}}{1 + \frac{1}{\eta} + (\frac{1}{\theta} - \frac{1}{\eta}) s_{ij}} \Delta \log s_{ij}$$

Constant returns to scale imply that log changes in crop exports are the sum of log changes in crop quantities and log changes in exporter productivity: $\Delta \log x_{ij} = \Delta \log z_{ij} + \Delta \log q_{ij}$. Holding fixed the behavior of other exporters, the nested CES supply curve further implies that log changes in exporter size can be expressed in terms of log changes in crop prices: $\Delta \log s_{ij} = (1 + \eta)(1 - s_{ij}) \Delta \log p_{ij}$. Substituting above and simplifying, we have:

$$\Delta \log p_{ij} = \left[1 + \frac{(\frac{1}{\theta} - \frac{1}{\eta})(1 + \eta) s_{ij} (1 - s_{ij})}{1 + \frac{1}{\eta} + (\frac{1}{\theta} - \frac{1}{\eta}) s_{ij}} \right]^{-1} \times (\Delta \log p_j^x + \Delta \log z_{ij})$$

Assuming that international price shocks are orthogonal to productivity shocks and rearranging yields an expression for the *partial equilibrium* pass-through:

$$\rho(s_{ij}) \equiv \frac{\Delta \log p_{ij}}{\Delta \log p_j^x} = \left[1 + \frac{(\frac{1}{\theta} - \frac{1}{\eta}) s_{ij} (1 - s_{ij}) (1 + \eta)}{1 + \frac{1}{\eta} (1 - s_{ij}) + \frac{1}{\theta} s_{ij}} \right]^{-1}$$

Clearly, pass-through is incomplete as long as $\eta > \theta$. In addition, one can show that pass-through is lower on average for larger exporters.

First, note that the derivative of the pass-through as a function of exporter market size can be written as follows:

$$\frac{\partial \rho}{\partial s_{ij}} = \frac{(1 + \eta)(\frac{1}{\theta} - \frac{1}{\eta}) [(\frac{1}{\theta} - \frac{1}{\eta}) s_{ij} (1 - s_{ij}) - (1 - 2s_{ij})]}{\{1 + \frac{1}{\eta} + (\frac{1}{\theta} - \frac{1}{\eta}) s_{ij} [(1 - s_{ij})(1 + \eta) + 1]\}^2}$$

For exporter size near 0, this expression is negative and large in absolute value. For exporter size near 1, this expression is positive but small in absolute value. Pass-through declines rapidly as size increases near 0, but only increases slowly as size increases near 1. This suggests that pass-through is lower on average among larger exporters.

Next, recall from Section III.A that because of strategic interaction among exporters, the data do not reveal the partial equilibrium pass-through. Strategic interaction makes small exporters more responsive to price shocks and large exporters *less* responsive in general equilibrium. In other words, the partial equilibrium pass-through underestimates the general equilibrium pass-through for small exporters and *overestimates* it for large exporters. This magnifies the decline in pass-through in the previous paragraph.

The model also yields predictions for the pass-through of international price changes to *quantities*:

$$\frac{\Delta \log q_{ij}}{\Delta \log p_j^x} = \frac{\Delta \log p_{ij}}{\Delta \log p_j^x} \left(\frac{\Delta \log p_{ij}}{\Delta \log q_{ij}} \right)^{-1} = \rho(s_{ij}) \times \left(\frac{1}{\eta}(1 - s_{ij}) + \frac{1}{\theta}s_{ij} \right)^{-1}$$

The first term is the price pass-through, which is less than 1 and declines with exporter size. The term in parentheses can be greater or less than 1, so there is no clear prediction for average quantity pass-through. However, since $\eta > \theta$, this term increases with exporter size, so that quantity pass-through unambiguously declines with size. Finally, note that if $\eta > 1$ and $\theta < 1$, quantity pass-through is higher than price pass-through as $s_{ij} \rightarrow 0$ and lower than price pass-through as $s_{ij} \rightarrow 1$. This implies that quantity pass-through must be declining *faster* with exporter size.

C Estimation appendix

C.1 Solving the model

To solve the model, I first guess crop market shares. Then, I solve for scaled crop supply elasticities and prices and use the prices to update market shares, iterating until the shares converge. Finally, I rescale to obtain crop prices and quantities. For a vector of parameters (η, θ, α) and a draw of productivities $\{z_{ij}\}$, the algorithm is as follows:

- Guess equal market shares $s_{ij} = \frac{1}{N_j}$
- Scaled equilibrium
 - Calculate supply elasticity $\epsilon_{ij} = \left(\frac{1}{\eta}(1 - s_{ij}) + \frac{1}{\theta}s_{ij} \right)^{-1}$
 - Calculate scaled prices $\hat{p}_{ij} = \left(\alpha \frac{\epsilon_{ij}}{1 + \epsilon_{ij}} z_{ij} s_{ij}^{-\frac{\eta - \theta}{1 + \eta}} \right)^{\frac{1}{1 + \theta}}$

- Update market shares $s_{ij} = \frac{\hat{p}_{ij}^{1+\eta}}{\sum_{i \in j} \hat{p}_{ij}^{1+\eta}}$
- Iterate until market shares converge
- Unscaled equilibrium
 - Calculate scaled price indexes $\hat{p}_j = (\sum_{i \in j} \hat{p}_{ij}^{1+\eta})^{\frac{1}{1+\eta}}$, $\hat{p} = (\sum_j \hat{p}_j^{1+\theta})^{\frac{1}{1+\theta}}$
 - Re-scale prices $p_{ij} = \hat{p}_{ij} \times \hat{p}^\theta$
 - Re-scale price indexes $p_j = (\sum_{i \in j} p_{ij}^{1+\eta})^{\frac{1}{1+\eta}}$, $p = (\sum_j p_j^{1+\theta})^{\frac{1}{1+\theta}}$
 - Calculate quantities $q_{ij} = (\frac{p_{ij}}{p_j})^\eta (\frac{p_i}{p})^\theta$

C.2 Simulated Method of Moments

I estimate (η, θ, α) via Simulated Method of Moments. The details are as follows:

- Guess (η, θ, α) . Draw productivities $\log z_{ij} \sim N(\mu_z, \sigma_z^2)$. Solve model and treat as data with $t = 1$.
- Draw shocks $\Delta \log p_{ijt}^x \sim N(\mu_p, \sigma_p^2)$. Solve model again and treat as data with $t = 2$.
- Estimate regressions in the simulated data

$$\Delta \log p_{ijt} = \delta_{jt} + \beta s_{ij,t-1} + \gamma \Delta \log p_{ijt}^x + \zeta s_{ij,t-1} \times \Delta \log p_{ijt}^x + u_{ijt}$$

- Estimate regressions in the real data

$$\Delta \log p_{ijt} q_{ijt} - \Delta \log x_{ijt} = \hat{\delta}_{jt} + \hat{\beta} s_{ij,t-1} + \hat{\gamma} \Delta \log p_{ijt}^x + \hat{\zeta} s_{ij,t-1} \times \Delta \log p_{ijt}^x + \hat{u}_{ijt}$$

- Calculate farmer shares in the simulated data

$$\phi = \sum_j \frac{p_j q_j}{\sum_k p_k q_k} \alpha \times \left[1 + \frac{1}{\eta} \left(1 - HHI_j \right) + \frac{1}{\theta} HHI_j \right]^{-1}$$

$$\hat{\phi} = \frac{\sum_{i(j),j} p_{ij} q_{ij}}{\sum_{k(l),l} p_{kl}^x q_{kl}}$$

- Pick (η, θ, α) to minimize $[\hat{m} - m(\eta, \theta, \alpha)]' W [\hat{m} - m(\eta, \theta, \alpha)]$.


where $\hat{m} = (\hat{\gamma}, \hat{\zeta}, \hat{\phi})'$ is the vector of data moments, $m(\eta, \theta, \alpha) = (\gamma(\eta, \theta, \alpha), \zeta(\eta, \theta, \alpha), \phi(\eta, \theta, \alpha))'$ is the vector of model moments, and W is a weighting matrix.

I perform the optimization using a Multi Level Single Linkage (MLSL) global algorithm with a Nelder-Mead local minimizer, as implemented by the NLOPTR package in R. This algorithm has been shown to perform well for Simulated Method of Moments (Arnaud, Guvenen and Kleineberg 2019).

C.3 Specifying demand shocks

Figure A2 plots the distributions of demand shocks under two different specifications of the shift-share design described in Section III.B. Both specifications use shares of export revenue by destination. The first, shown in blue, uses shifts in import prices at the destination (excluding imports from Ecuador) obtained from CEPII’s *World Trade Flows Characterization* database. It is well-approximated by a normal distribution with mean 0.02 and standard deviation 0.11. The second, shown in black, uses shifts in import *expenditures* at the destination (again excluding imports from Ecuador) obtained from CEPII’s *BACI* database. This generates substantially more dispersion in demand shocks, and is well-approximated by a normal distribution with mean 0.05 and standard deviation 0.15. When solving the model, I can draw price shocks directly from the distributions in the data. For the sake of reproducibility, I draw from the fitted normal distributions instead.

Figure A2: Percent change in international prices



Notes: Solid blue line plots density of percent change in international prices. Dashed black line plots density of percent change in international expenditures.

C.4 Recovering exporter productivities

When estimating the model, I pick the mean and standard deviation of log exporter productivity to match the distribution of log exporter revenue in the data. However, it is possible to recover exporter productivities non-parametrically following the procedure in Berger et al. (forthcoming). First, note that for exporters i and i' of crop j , dividing scaled crop prices from above yields:

$$\frac{\hat{p}_i}{\hat{p}_{i'}} = \left(\frac{\psi(s_i)}{\psi(s_{i'})}\right)^{\frac{1}{1+\theta}} \left(\frac{z_i}{z_{i'}}\right)^{\frac{1}{1+\theta}} \left(\frac{s_i}{s_{i'}}\right)^{-\frac{\eta-\theta}{(1+\eta)(1+\theta)}}$$

where I have suppressed the j subscript and $\psi(s_i) = (1 + \frac{1}{\epsilon_i})^{-1}$ is the optimal markdown as a function of exporter size. Note also that the equilibrium exporter size $s_{ij} = (\frac{\hat{p}_{ij}}{\hat{p}_j})^{1+\eta}$, which implies that $\frac{\hat{p}_i}{\hat{p}_{i'}} = (\frac{s_i}{s_{i'}})^{\frac{1}{1+\eta}}$. Substituting above and rearranging yields a simple expression for the relative productivities of i and i' :

$$\frac{z_i}{z_{i'}} = \frac{\psi(s_{i'})/s_{i'}}{\psi(s_i)/s_i}$$

This equation says that a more productive exporter (higher z_i) pays farmers a lower markdown relative to his size (lower $\psi(s_i)/s_i$). Intuitively, more productive exporters in the model are both larger and pay lower markdowns, so it is reasonable to infer relative productivity from relative markdowns and relative sizes.

C.5 Bertrand competition

Table A2 shows estimates of the key parameters under Bertrand competition. Quantitatively, Bertrand competition indeed implies both a lower elasticity of substitution across crops and lower levels of market power. However, the results are qualitatively similar to the case with Cournot competition.

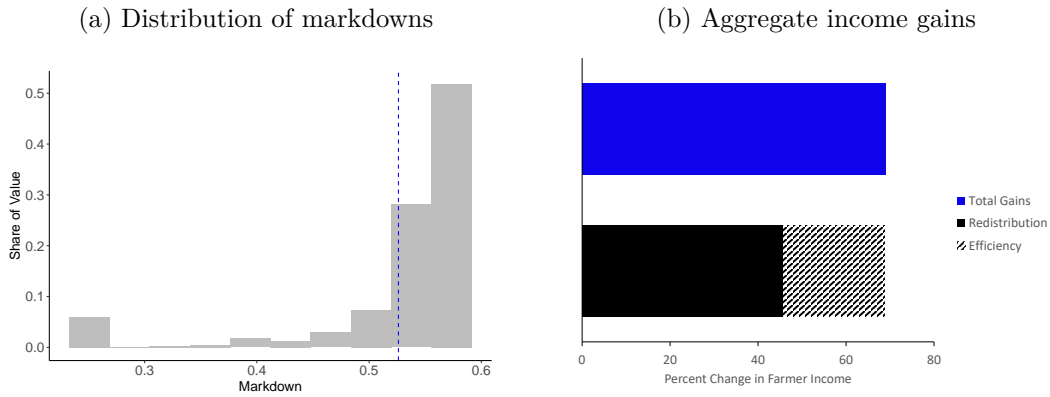
Table A2: Key parameters, Bertrand competition

Parameter	Cournot	Bertrand	Moment	Value
η	1.93	2.00	$\hat{\gamma}$	0.35
θ	0.40	0.21	$\hat{\zeta}$	-0.23
α	0.45	0.38	$\hat{\phi}$	0.24

Figure A3 plots micro and macro measurements of monopsony power under Bertrand competition. Panel A plots the distribution of markdowns. As expected,

the distribution shifts to the right, indicating that exporters pay farmers a larger share of their marginal revenue product and hence are more competitive. The weighted average is only 0.53, so market power is still substantial. In Panel B plots the change in farmer income under perfect competition. As expected, the overall gains (66.1%) are lower when the baseline is Bertrand than when the baseline is Cournot. However, the breakdown between redistribution (43.4%) and efficiency (21.9%) is similar to the Cournot case.

Figure A3: Measurements of monopsony power (Bertrand competition)



Notes: Panel A plots the distribution of markdowns across exporters. Markdowns were calculated using s_{ij} from the data and the estimated η and θ . Dashed blue line indicates average markdown of 49%. Panel B shows percent increase in farmer income between the estimated model with monopsony power and a counterfactual model with perfect competition. Blue area indicates total gains. Black area indicates gains from redistributing exporter profits to farmers. Shaded area indicates gains from greater efficiency.

C.6 Overidentified model

In this section, I estimate an overidentified version of the model under both Cournot and Bertrand competition. I proceed as in Section III.B, with one important modification. In addition to matching the baseline pass-through (γ in Equation 7), the decline in pass-through with exporter size (ζ in Equation 7), and the average farmer share, I match the decline in farmer share with exporter size (β in Equation 5). The theory implies that this coefficient is a function of η and θ , as discussed in Section II.D. Furthermore, it is precisely estimated in Table 3, unlike the coefficient on the interaction term in Table 4. This will be particularly helpful for estimating θ .

To estimate the model under Bertrand competition, I make two modifications to the estimation procedure in Section III.B. First, I compute the optimal farm price using the Bertrand supply elasticity (Equation A3) rather than the Cournot supply elasticity (Equation 3). Second, I choose the output elasticity α to match the Bertrand farmer share (Equation A5) rather than the Cournot farmer share (Equation 8).

Table A3 presents estimates of the key parameters. The overidentified model features stronger potential market power than the baseline model in the form of lower elasticities of substitution η and θ . However, the actual market power implied by the output elasticity α is similar to that of the baseline model. Note that the Cournot model matches all moments well, despite being overidentified. However, the Bertrand model struggles to generate both the steep decline in pass-through and the steep decline in farmer shares as a function of exporter size.

Table A3: Key parameters, overidentified model

Parameter	Cournot	Bertrand	Moment	Value (Data)	Value (Cournot)	Value (Bertrand)
η	1.90	1.94	$\hat{\gamma}$	0.35	0.36	0.44
θ	0.41	0.37	$(\hat{\zeta}, \hat{\beta})$	(-0.23,-0.82)	(-0.22,-0.83)	(-0.16,-0.89)
α	0.44	0.38	$\hat{\phi}$	0.24	0.24	0.24

C.7 Standard errors

The variance-covariance matrix of the key parameters, V , is given by:

$$V = (1 + \frac{1}{S})(G'WG)^{-1}G'W\Omega WG(G'WG)^{-1}$$

where G is the matrix of partial derivatives of the model moments, $m(\eta, \theta, \alpha)$, with respect to (η, θ, α) , Ω is the variance-covariance matrix of the moments, and S is the number of simulations (Gourieroux, Monfort and Renault, 1993). The standard error of each parameter is then given by the square root of the corresponding element along the diagonal of V . Replacing W with the optimal weighting matrix, Ω^{-1} , we have:

$$V = (1 + \frac{1}{S})(G'\Omega^{-1}G)^{-1}$$

For the exactly-identified model, $m = (\gamma, \zeta, \phi)'$ as above, and G and Ω are 3x3 matrices. For the over-identified model, $m = (\gamma, \zeta, \phi, \beta)$, G is a 4x3 matrix, and Ω is a 4x4 matrix. I compute G using a 1% deviation in each of the estimated parameters. For Ω , I use the variance-covariance matrix of data moments. Table A4 below summarizes the standard errors for the two models under Cournot competition. Although η is precisely estimated in both models, θ is precisely estimated only in the overidentified model. This is because the coefficient on exporter size in the farmer share regression, β , is highly significant and informative of the difference between η and θ . Finally, note that α is imprecisely estimated in both specifications. In the model, α is calculated using estimates of η and θ and therefore inherits their uncertainty. In the data, there is considerable variation in farmer shares, which adds further uncertainty.

Table A4: Standard errors

Parameter	Estimate (Exactly Identified)	SE (Exactly Identified)	Estimate (Overidentified)	SE (Overidentified)
η	1.93	0.99	1.90	0.94
θ	0.40	0.36	0.41	0.14
α	0.45	1.70	0.44	0.47

C.8 Pass-through to farmer quantities

Above, I showed that pass-through to farmer *quantities* declines with exporter size, and declines more steeply than pass-through to prices. I test these predictions by estimating the following regression:

$$(A7) \quad \Delta \log x_{ijt} = \delta_{jt} + \beta s_{ij,t-1} + \gamma \Delta \log p_{ijt}^x + \zeta s_{ij,t-1} \times \Delta \log p_{ijt}^x + u_{ijt}$$

where the terms are defined as in Equation 7. Table A5 displays the results of different specifications analogous to those of Table 4. As predicted by the theory, quantity pass-through decreases significantly with size. The point estimate on the interaction term is more negative than in Table 4, indicating that quantity pass-

through declines more steeply.²⁴ The positive correlation between price responses in Table 4 and quantity responses in Table A5 support the interpretation of international price shocks as demand shocks for exporters. By shifting the demand curve for exporters, these shocks trace out their supply curves and identify buyer market power.

Table A5: Quantity responses to price shocks

	$\Delta \log x$	$\Delta \log x$	$\Delta \log x$
	(1)	(2)	(3)
s	-0.138 (0.103)	0.001 (0.131)	0.130 (0.139)
$\Delta \log p^x$	0.055 (0.226)	0.014 (0.238)	0.063 (0.237)
$s \times \Delta \log p^x$	-0.575 (0.493)	-0.685 (0.516)	-0.735 (0.514)
FE	No	Yes	Yes
Controls	No	No	Yes
Observations	767	767	767
R ²	0.005	0.047	0.062

Notes: Table summarizes quantity pass-through regressions. Dependent variable is the change in log quantity exported. Independent variables are the change in the log international price, the lagged exporter size, and their interaction. Column 3 controls include changes in the wage bill and payments to non-farm suppliers, and an indicator for exporters with lagged market share less than 1%. Clustered standard errors are shown in parentheses.

C.9 External validation of θ

To compare crop-specific productivity shocks in my model to those in the agricultural trade literature, assume there is a single exporter for each crop, so that the only relevant shock is $\frac{\nu_{fj}}{1+\theta}$. A farmer with efficiency q_f now produces $e^{\frac{\nu_{fj}}{1+\theta}} q_f = e^x q_f$ units of crop j , where x follows a Gumbel distribution with scale parameter $\frac{1}{1+\theta}$. In the literature, land heterogeneity typically follows a Frechet distribution with shape parameter $\tilde{\theta}$. It remains to convert the cost shock to a productivity shock, and the Gumbel parameter to the associated Frechet parameter.

²⁴The theory makes no clear prediction for average quantity pass-through, but the data suggest it is substantially lower than average price pass-through.

Rewrite the cost shock $z = e^x$. The CDF of z is $G(z) = P(e^x \leq z) = P(x \leq \log z) = F(\log z)$, where F is the CDF of x . Substituting $\log z$ into the CDF for the Gumbel distribution, we obtain the CDF of the Frechet distribution with shape parameter $1 + \theta$. Therefore, my estimate of $\hat{\theta} = 0.40$ corresponds to a shape parameter of 1.40 for the distribution of land heterogeneity. The following table reports this estimate, along with those from a selection of papers.

Table A6: Sources for Figure 4a

Reference	Land heterogeneity	Source
Costinot et al. 2016	2.46	Table 2
Farrokhi and Pellegrina 2021	2.05	Table 2
Bergquist et al. 2019	1.80	Section 4
Sotelo 2020	1.66	Section 5
This paper	1.40	Section C.9
Gouel and Laborde 2021	1.2	Section 6.2

C.10 External validation of η

To compare exporter-specific cost shocks in my model to those in the agricultural trade literature, assume there is a single crop, so that the only relevant shock is $\frac{\nu_{fi}}{1+\eta}$. A farmer with efficiency q_f delivers $e^{\frac{\nu_{fi}}{1+\eta}} q_f = e^x q_f$ units to exporter i , where x follows a Gumbel distribution with scale parameter $\frac{1}{1+\eta}$. In addition, assume that trade costs are the only source of heterogeneity in exporter-specific costs. In the literature, trade costs are typically deterministic and takes an iceberg form. As a result, I compare the mean trade cost estimates from the literature to the mean implied by my estimates, expressed in iceberg form.

Following the derivation above, the Gumbel distribution with scale parameter $\frac{1}{1+\eta}$ is equivalent to the Frechet distribution with scale parameter $1 + \eta$. The mean of a Frechet distribution with scale parameter $1 + \eta$ is $\Gamma(1 - \frac{1}{1+\eta})$, where $\Gamma(\cdot)$ is the gamma function. Substituting my estimate of $\eta = 1.93$ yields a mean of 1.42. To convert this to iceberg form, I divide the 90th percentile of the Frechet distribution by the average, yielding an average trade cost of 1.69. The following table reports this estimate, along with those from a selection of papers.

Table A7: Sources for Figure 4b

Reference	Iceberg trade cost	Source
Atkin and Donaldson 2015	1.12	Section 4.3
Chatterjee 2020	1.16	Section 6.1.1
Bergquist et al. 2019	1.25	Section 4
Allen 2014	1.47	Table 7
This paper	1.69	Section C.10
Sotelo 2020	2.34	Reported in Table 4

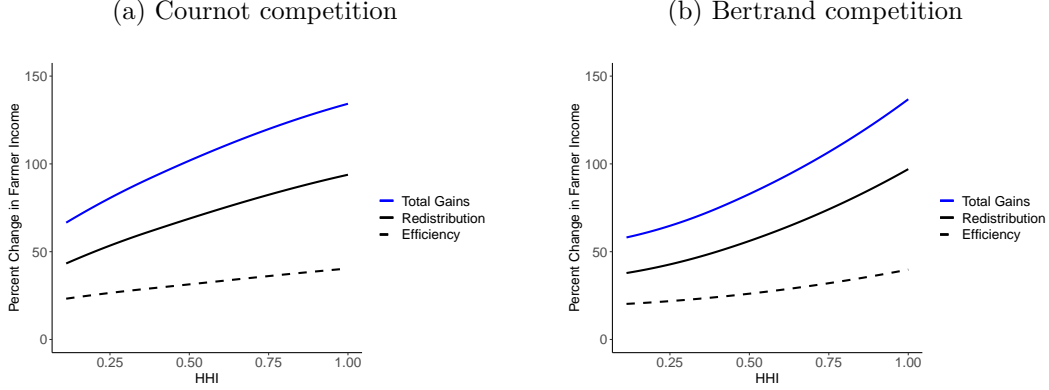
D Measurement appendix

D.1 Heterogeneous gains from perfect competition

Although all farmers gain from perfect competition, the gains are not equally shared. Panel A of Figure [A4](#) shows how increases in farmer income vary with the baseline level of crop market concentration, HHI_j , under Cournot competition. Gains range from around 54% in relatively competitive crops, such as bananas, to 122% in the least competitive crops, including cocoa. Both redistribution and efficiency gains increase with crop market concentration, but redistribution increases disproportionately more.

Panel B shows a similar pattern for Bertrand competition. Note that the gains are smaller than under Cournot competition for the least concentrated markets, but larger for the most concentrated markets. This is related to the result that the Lerner Index is linear in market shares under Cournot competition, but convex under Bertrand competition ([Alvarez, Head and Mayer 2020](#)).

Figure A4: Farmer income gains and crop market concentration



Notes: Figure plots the HHI from the estimated model with monopsony on the x-axis and the percent change in farmer income between the estimated model and a counterfactual with perfect competition on the y-axis. HHI is defined as the sum of squared market shares within each crop. Solid blue line indicates total gains by crop. Solid black line indicates gains from redistribution. Dashed black line indicates efficiency gains. Panel A assumes Cournot competition, and Panel B assumes Bertrand competition.

D.2 Asymmetric pass-through regressions

Farmer income is lower under market power, but also less sensitive to trade shocks. If larger exporters offer more insurance, this mean-variance trade-off will be correlated with exporter size, as in Tables 3 and 4. In this case, exporters should help smooth the effects of both positive and negative shocks. To test this, I estimate the following regression:

$$\begin{aligned}
 \Delta \log p_{ijt} q_{ijt} - \Delta \log x_{ijt} &= \delta_{jt} + \beta s_{ij,t-1} + \kappa \mathbf{1}(\Delta \log p_{jt}^x > 0) + \gamma \Delta \log p_{jt}^x \\
 \text{(A8)} \quad &+ \lambda \mathbf{1}(\Delta \log p_{jt}^x > 0) \times s_{ij,t-1} + \mu \mathbf{1}(\Delta \log p_{jt}^x > 0) \times \Delta \log p_{jt}^x \\
 &+ \zeta s_{ij,t-1} \times \Delta \log p_{jt}^x + \xi \mathbf{1}(\Delta \log p_{jt}^x > 0) \times s_{ij,t-1} \times \Delta \log p_{jt}^x + u_{ijt}
 \end{aligned}$$

where $\mathbf{1}(\Delta \log p_{jt}^x > 0)$ is an indicator for whether the international price shock is positive and the other terms are defined as in Equation 7. The coefficient on the triple interaction term, ξ , indicates whether large exporters respond differentially to positive shocks compared to small exporters. Table A8 displays the results. The negative and marginally significant coefficients in the third row suggest that the decline in pass-through with exporter size is driven by *positive* shocks. In contrast, negative shocks have similar pass-through among small and

large exporters. If insurance were the sole mechanism at play, the pass-through of international price shocks would not depend on sign of the shock.

Table A8: Asymmetric responses to price shocks

	$\Delta \log pq - \Delta \log x$	$\Delta \log pq - \Delta \log x$	$\Delta \log pq - \Delta \log x$
	(1)	(2)	(3)
$\Delta \log p^x$	0.290 (0.487)	0.394 (0.509)	0.382 (0.509)
$s \times \Delta \log p^x$	0.148 (0.656)	0.085 (0.687)	0.099 (0.688)
$\mathbf{1}(\Delta \log p^x > 0) \times s \times \Delta \log p^x$	-1.395 (0.924)	-1.300 (0.965)	-1.359 (0.968)
FE	No	Yes	Yes
Controls	No	No	Yes
Observations	947	947	947
R ²	0.010	0.050	0.050

Notes: Table summarizes asymmetric pass-through regressions. Dependent variable is the change in log farmer price. Independent variables are the change in the log international price, the lagged exporter size, an indicator for positive changes in the log international price, and all interactions. Column 3 controls include changes in the wage bill and payments to non-farm suppliers, and an indicator for exporters with lagged market share less than 1%. Clustered standard errors are shown in parentheses.

D.3 Calculation of farmer risk aversion

Recall from above that log changes in international prices are approximately normally distributed with mean $\mu_d = 0.02$ and variance $\sigma_d^2 = 0.11$. Starting from an equilibrium, farmer income Y_t follows a Geometric Brownian Motion:

$$dY_t = \rho\mu_d Y_t dt + \rho\sigma_d Y_t dW_t$$

where ρ is the pass-through rate and W_t is a Wiener process. In a one-period model with initial income Y_0 , farmer income follows a log-normal distribution with mean $\mu_y = \log Y_0 + \rho\mu_d - \frac{\rho^2\sigma_d^2}{2}$ and variance $\sigma_y^2 = \rho^2\sigma_d^2$.

Suppose that farmers have Constant Relative Risk Aversion (CRRA) preferences with coefficient γ . Given a log-normal income process with mean μ_y and variance σ_y^2 , the certainty equivalent, x , is:

$$x = e^{\mu_y + \frac{\sigma_y^2(1-\gamma)}{2}}$$

The farmer is indifferent between receiving x with certainty and receiving income according to the risky log-normal process. Therefore, the farmer is indifferent

between two income processes if they have the same certainty equivalent. Substituting the expressions for μ_Y and σ_Y^2 into the certainty equivalent formula, we have:

$$Y_0^{PC} e^{\rho_{PC}\mu_d - \frac{\rho_{PC}^2\sigma_d^2\gamma}{2}} = Y_0^{MP} e^{\rho_{MP}\mu_d - \frac{\rho_{MP}^2\sigma_d^2\gamma}{2}}$$

where PC denotes the equilibrium with perfect competition and MP denotes the equilibrium with market power. Solving for γ yields:

$$\gamma = \frac{2 \log(Y_0^{PC}/Y_0^{MP})}{2(\rho_{MP} - \rho_{PC})\mu_d + (\rho_{PC}^2 - \rho_{MP}^2)\sigma_d^2}$$

Plugging in $\mu_d = 0.02$, $\sigma_d^2 = 0.11$, $Y_0^{PC} = 1.77Y_0^{MP}$, $\rho_{PC} = 1$, and $\rho_{MP} = 0.42$ yields $\gamma = 7.36$. This corresponds to the highest category of risk aversion estimated in [Yesuf and Bluffstone \(2009\)](#).

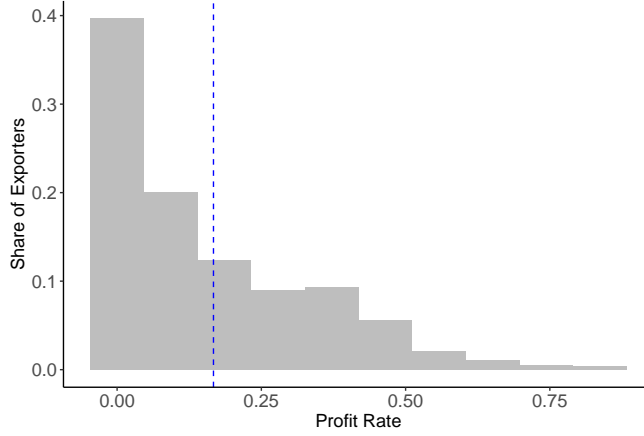
D.4 Distribution of implied profit rates

Substituting the first order conditions for q_{ij} and m_{ij} into the profit function and dividing yields the following expression for the profit rate:

$$\frac{\pi_{ij}}{p_j^x x_{ij}} = \text{farmer share}_{ij} \times \left(\frac{1}{\eta}(1 - s_{ij}) + \frac{1}{\theta}s_{ij} \right)$$

Using data on farmer shares and s_{ij} together with estimates of η and θ , I compute the implied profit rate for each exporter. [Figure A5](#) plots the distribution across exporters. Profits are large, averaging 16% of exporter revenue. The correlation between profit rates and exporter size is not obvious: the farmer shares decreases with exporter size, while the term in parentheses increases. In regressions not shown, I find that the latter effect dominates, so that larger exporters have higher profit rates, even controlling for crop-year fixed effects. This is inconsistent with the idea that these profits represent investments in sourcing and exporting technology, which typically exhibit increasing returns to scale.

Figure A5: Distribution of profit rates



Notes: Figure plots the distribution of the profit rate across exporters. Profit rates were calculated using farmer shares and s_{ij} from the data and the estimated η and θ . Dashed blue lines indicates that the average profit rate is 0.16.

D.5 Calculation of fixed costs to exporting

If banana exporters face a common fixed cost f to exporting, they will only export if profits are sufficiently large to cover the fixed cost. Substituting the first order conditions for q_{ij} and m_{ij} into the profit function, we have:

$$\pi_{ij} = p_{ij}q_{ij} \times \left(\frac{1}{\eta}(1 - s_{ij}) + \frac{1}{\theta}s_{ij} \right) \geq f$$

The smallest banana exporter has negligible market share. Plugging in $\eta = 1.93$, $\theta = 0.41$, and $s_{ij} = 0$, this implies that f is no larger than 52% of the farmer income paid by the smallest exporter. Under perfect competition, farmer income is 54% higher on average across the entire sector (Figure A4). The farmer income of the smallest exporter is by definition smaller than that of the average exporter. This suggests that farmer income is at least 2% higher under perfect competition, even net of fixed costs.

E Policy appendix

E.1 Additional features of Fair Trade

In my analysis of Fair Trade, I abstract from selection, non-monetary benefits, costs of certification, downstream demand, and fixed costs of exporting. Here, I discuss how these issues might affect my results.

The net effect of selection is unclear. Higher quality farmers may face lower costs of certification, so that there is positive selection (Dragusanu and Nunn 2018). In this case, my model will underestimate the gains. On the other hand, lower quality farmers may perceive higher benefits from certification, so that there is negative selection (Ruben and Fort 2012). In that case, my model will overestimate the gains. For a theoretical model that incorporates selection, see Podhorsky (2015).

Certification costs reduce the net benefits of Fair Trade for farmers, so that my model will overestimate the gains. However, existing estimates of such costs correspond to a 25% reduction in the gains (De Janvry et al., 2015), so that farmer income still increases by 9% in my preferred specification. On the other hand, non-monetary aspects of Fair Trade increase the net benefits to farmers, so that my model will underestimate the gains.

Experimental evidence suggests that *sales* of Fair Trade products in supermarkets are 10% higher than those of identical conventional products (Hainmueller et al., 2015). Since farmers earn such a small share of sales to begin with, this more than covers the 12% increase in farmer payments in my preferred specification.

Finally, if the fixed cost of exporting directly is such that the marginal strategic exporter is indifferent between entering and not entering, then the indirect effect of Fair Trade implies a cost of no more than 4% of farmer income. This is modest relative to the increase in farmer income caused by the policy.

E.2 Alternative formulations of Fair Trade

In the main text, I model Fair Trade by introducing a perfectly competitive exporter with productivity drawn from the estimated distribution. Here, I show how this maps to different ways of implementing Fair Trade.

One way of implementing Fair Trade is by helping farmers form cooperatives to export directly. In Section II.C, I show that a profit-maximizing exporter with market power offers farmers a price that is marked down from their marginal revenue product. In contrast, a perfectly competitive exporter offers farmers their marginal revenue product. Suppose an exporter with market power instead maximized farmer income subject to non-negative profits. He solves:

$$\max_{q_{ij}} \{p_{ij}q_{ij} - \mu[\alpha p_j^x x_{ij} - p_{ij}q_{ij}]\}$$

where μ is the Lagrange multiplier on the constraint I have used the fact that $p_j^m m_{ij} = (1 - \alpha)p_j x_{ij}$. Clearly $\mu > 0$ for an interior solution. The first order conditions for q and μ , respectively, are:

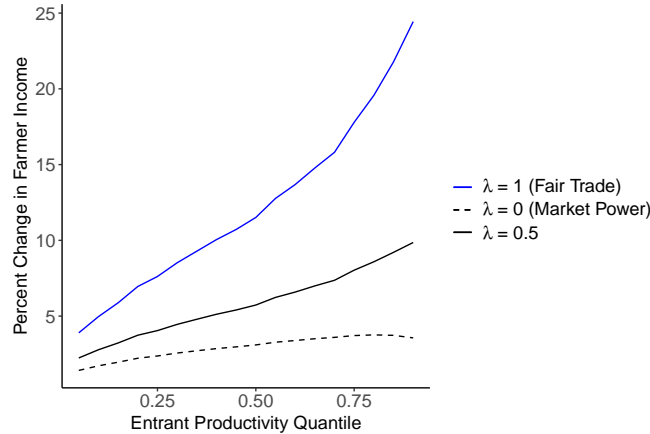
$$p_{ij}q_{ij} = \frac{\mu}{\mu-1} \times \alpha^2 \times \left(1 + \frac{1}{\epsilon_{ij}}\right)^{-1} p_j^x x_{ij}$$

$$p_{ij}q_{ij} = \alpha p_j^x x_{ij}$$

Both equations are satisfied when $\frac{\mu}{\mu-1} = \alpha^{-1} \times \left(1 + \frac{1}{\epsilon_{ij}}\right)$, which implies that exporters pay farmers their marginal revenue product. In other words, a strategic exporter who maximizes farmer income behaves competitively.

A Fair Trade exporter may maximize a weighted sum of profits and farmer income (Podhorsky, 2015), offering a price between the marginal revenue product. Figure A6 shows the effect on farmer income when the Fair Trade exporter places weight λ on farmer income and $1 - \lambda$ on profits. The solid black line corresponds to the case where $\lambda = 0.5$. The blue line is the limiting case where $\lambda = 1$ (perfect competition), while the dashed line is the limiting case where $\lambda = 0$ (market power). Note that a relatively large weight on farmer income is required to approach the Fair Trade gains.

Figure A6: Effect of Fair Trade on farmer income



Notes: Figure plots productivity quantile of counterfactual exporter on the x-axis and resulting percent change in farmer income on the y-axis. Dashed black line indicates counterfactuals in which exporter maximizes profits. Solid blue line indicates Fair Trade counterfactual in which exporter maximizes farmer income. Solid black line indicates counterfactual in which exporter maximizes a weighted sum of profits and farmer income with equal weights.


Another way of implementing Fair Trade is to specify a minimum price for certified products (plus a premium). This is distinct from a broad price floor because only *certain* buyers must adhere to the minimum price, i.e. those participating in the Fair Trade program. As a result, the Fair Trade minimum price does not have the distortionary effect of a broad price floor. Fair Trade products can always be sold to conventional buyers. Therefore, the Fair Trade price must be higher than the conventional price to guarantee enough takeup from farmers. To prevent too much takeup, Fair Trade buyers typically specify a maximum certified quantity (Podhorsky, 2015). In the model, the Fair Trade exporter sets the optimal price given his productivity level and buys the market-clearing quantity. Price and quantity are strictly increasing in productivity, so that varying the productivity traces out the curve of Fair Trade prices and certified quantities.

Figure A7 compares the effect of a Fair Trade price to that of a universal price floor. The dashed black line indicates the percent change in farmer income as a function of the price quantile for a policy that specifies a minimum price for all exporters, as in Figure 7b. The solid blue line does the same for a policy that specifies a minimum price only for Fair Trade exporters. For a given price, the Fair Trade policy always yields larger gains under the baseline value of θ . As θ increases, the difference shrinks, to the point that a sufficiently high universal price floor yields larger gains.

Figure A7: Fair Trade vs minimum prices (all products)

(a) $\theta = 0.40, \eta = 1.93$

(b) $\theta = 0.80, \eta = 1.93$



Notes: Figure plots price quantile of counterfactual price floor on the x-axis and resulting percent change in farmer income on the y-axis. Dashed black line indicates counterfactuals with a universal price floor. Solid blue line indicates Fair Trade counterfactual with the same price. Elasticity of substitution across crops, θ , varies as indicated in each title.

