Scaling Agricultural Policy Interventions[*]

Lauren Falcao Bergquist[†], Benjamin Faber[‡], Thibault Fally[§],
Matthias Hoelzlein[¶], Edward Miguel[‖], Andres Rodriguez-Clare[**]

August 2023

Policies aimed at raising agricultural productivity have been a centerpiece in the fight against global poverty. We propose a new approach for quantifying large-scale agricultural policy counterfactuals that can both complement and be informed by evidence from field and quasi-experiments. We develop a quantitative model of farm-level agricultural trade that captures important, but typically neglected features of this setting, including homogeneous goods and additive trade costs. We propose a new solution method in this environment that relies on rich but widely available microdata. We harness field and quasi-experiments for parameter estimation, and showcase our approach in the context of subsidies for modern inputs in Uganda. We find that the average welfare gain from treatment, for the same sample of households, falls by 20% when implemented at scale. At the same time, the effect for the poorest households increases at scale as the gains shift from land onto labor, reducing the regressivity of the local intervention by more than half. We further document how these forces are shaped by the granular economic geography often missing in existing quantitative models and by the geographical scale of implementation, with new implications for randomized saturation designs. Finally, we discuss practical considerations for combining our toolkit with evidence from field and quasi-experiments.

JEL classification: F15, F63, O13

[*]This research was supported by funding from the International Growth Centre (IGC). We thank audiences at multiple seminars for comments and Rodrigo Adao, Wyatt Brooks, Pablo Fajgelbaum and Sebastian Sotelo for helpful discussions. All errors are our own.

[†]Department of Economics and Jackson School of Global Affairs, Yale and NBER
[‡]Department of Economics, UC Berkeley and NBER
[§]Department of Agricultural and Resource Economics, UC Berkeley and NBER
[¶]Department of Economics, Notre Dame
[‖]Department of Economics, UC Berkeley and NBER
[**]Department of Economics, UC Berkeley and NBER

1 Introduction

Roughly two thirds of the world's population living below the poverty line work in agriculture (Castaneda *et al.* , 2016). In this context, policies aimed at raising agricultural productivity, such as programs providing access, training and subsidies for modern inputs and production techniques, have been a centerpiece in the fight against global poverty.[1] To inform these policies using rigorous evidence, much of the recent literature studies local interventions with variation in policy exposure across households or local markets generated by randomized control trials (RCTs) or natural experiments. While rightly credited for revolutionizing the field of development economics, field and quasi-experiments often face the well-known limitation that their estimates may not speak to the broader general equilibrium (GE) effects that emerge once policies are scaled up. At the same time, an earlier literature in agriculture and development, employing computable general equilibrium (CGE) analysis to quantify GE implications, often rely on less well-identified moments for parameter estimation and largely abstracts from modeling the granular economic geography of farm production, consumption and trade costs that underlies the propagation of shocks and their incidence in GE.[2]

To make progress on these challenges, we propose a new methodology for quantifying large-scale policy counterfactuals at the level of households in agricultural settings. We aim to quantify how the average treatment effect and distributional implications of a local intervention differ – for the same group of households – if the treatments are scaled up to a larger segment of the population. Our approach can both complement evidence from field and quasi-experiments and be informed by it.[3] To be able do so, we introduce in our theory several well-known features of agricultural trade across local markets that have been outside the scope of quantitative models and their solution methods in international trade and economic geography.[4] The first is that individual crops are best described as homogeneous goods, counter to the common assumption of differentiated varieties that may be more suitable for cross-country trade in manufacturing. As a result, policy shocks at scale may change which markets are connected through trade in different crops, an important extensive margin of adjustment that is ruled out in most quantitative models of trade.[5] Second, we allow for both additive (per-unit) and multiplicative (ad-valorem) components of trade costs, across households and across agricultural markets. Additive trade

[1]See e.g. Caldwell *et al.* (2019) for a review of recent impact evaluations in this space.

[2]See e.g. de Janvry & Sadoulet (1995) for a review of this literature.

[3]Similar to recent work by e.g. Brooks & Donovan (2020), Gollin *et al.* (2021), Lagakos *et al.* (2021) and Porzio *et al.* (2022), our analysis combines a structural model with rich microdata and evidence from RCTs or quasi-experiments to quantify GE counterfactuals that are frequently outside the scope of reduced-form estimation.

[4]In addition to the two features we focus on here, we allow for non-homothetic preferences – so that food price changes can have distributional implications beyond affecting household incomes – and technology choice in crop production – such that the adoption of modern inputs can more flexibly affect the production function with respect to other inputs. Our approach to modeling technology choice is similar to that in Farrokhi & Pellegrina (2022).

[5]Notable exceptions are Costinot & Donaldson (2016) and Sotelo (2020) who require rich additional data (on either production possibility frontiers or farm-gate prices) to solve the model. Since such data requirements can be hard to satisfy, especially at the level of households, we propose a new solution method that unlocks the scope for counterfactual analysis in such an environment.

costs can give rise to incomplete and heterogeneous pass-through of local price changes for crops and inputs across trading pairs, in contrast to the conventional ad-valorem (”iceberg”) assumption with complete pass-through.[6]

Both of these features are fundamental for modeling a granular and realistic economic geography that underlies the propagation of shocks across markets and households within them in this setting. But as we show, they break the convenient properties of “structural gravity” (Head & Mayer, 2014), including the use of “exact hat algebra” as the conventional solution method in the literature (see e.g. Costinot & Rodríguez-Clare (2014)). After laying out the model, we propose a new solution method for counterfactual analysis in this environment that relies on rich but widely available microdata on household location, production and consumption. We first show that we can use information on trade costs between and within markets in combination with data on household-level expenditure shares and agricultural production quantities to set up a price discovery problem. This entails solving for equilibrium farm-gate prices and trade flows that rationalize the observed consumption and production decisions given a graph of trade costs connecting households and markets. In turn, with knowledge of farm-gate prices and trade costs, we can then proceed to solve for the counterfactual equilibrium, following an approach that then becomes similar to exact hat algebra.

This approach has several advantages. First, we are able to solve the model without imposing structural gravity and without introducing stark new data requirements – such as requiring data on the full set of initial farm-gate prices. Second, our solution method ensures that the economy is in equilibrium before solving for counterfactuals: the household prices we obtain from the price discovery are by construction consistent with the calibrated trade costs and the consumption and production decisions we observe in the data.[7] From a computational perspective, our solution method is capable of handling high-dimensional GE counterfactuals at the level of individual households who populate the macroeconomy. This allows us to match the unit of observation often used in experiments (individual households), as well as to speak to distributional effects at this granular level.

We showcase our approach by evaluating the local versus at-scale implications of a subsidy for modern inputs (chemical fertilizers and hybrid seed varieties) in Uganda.[8] Drawing on the strength of experiments for identification, we estimate the model’s key demand and supply elasticities using exogenous variation in consumer and producer prices from existing RCTs (Bergquist & Dinerstein, 2020, Carter *et al.* , 2020). On the supply side, we also make use of a natural experiment that exploits changes in crops’ world market prices that propagate differently to local markets as a function of (additive) trade costs to the nearest border crossing. To calibrate cross-market trade costs, we make use of estimates from Bergquist *et al.* (2022), using Ugandan

[6] Price changes at an origin pass through differently both across destinations with higher (-) or lower (+) additive trade costs, and across crops or inputs with higher (+) or lower (-) unit values within a given bilateral route.

[7] For example, Sotelo (2020) uses province-level crop prices from agricultural surveys to calibrate and solve the model, but these price data are not, in general, model-consistent given calibrated trade costs.

[8] One of the most widespread and costly agricultural policies in low-income countries. See e.g. Jayne & Rashid (2013) and discussion below.

market and trader survey microdata to provide information on market-to-market trade flows and crop prices at origin and destination across crops. To calibrate within-market trade costs between farmers and their local markets, we use observed gaps in the Ugandan National Panel Survey (UNPS) between farm-gate prices and local markets in combination with knowledge of farmer-level trade flows to and from the markets. Finally, we use Ugandan administrative data on household location, production and consumption to calibrate the model to the roughly 4.5 million households who populate the country.

We then use the calibrated model to conduct counterfactual analyses. We first study how the average and distributional effects of this policy differ between a local intervention and one implemented at scale. We run two types of counterfactuals for each of the roughly 4,500 rural parishes in Uganda. In each parish, we randomly select 2.5 percent of the local population (a sample of roughly 100,000 households nationwide, or 25 per parish). We first solve for counterfactual changes in household welfare due to an intervention that targets a 75% cost subsidy for modern inputs only at each of these local treatment groups, keeping the rest of Uganda unexposed – akin to implementing roughly 4,500 separate RCTs. We then compare these local effects to the welfare changes experienced by the same sample of households under an intervention that scales the subsidy policy to all rural households in Uganda.

Pooling all local randomized interventions, we find that the average effect of the subsidy at small scale is a 4.4 percent increase in household real income. This is driven almost entirely by farmers saving on costs for the subsidized inputs and using more of them, while output and other input prices remain mostly unaffected. However, at scale we find that the welfare effect – for the same sample of farmers receiving the same intervention as in the local experiment – changes by as much as + or -5 percentage points across households. This is large relative to the local treatment effect: over a third of households experience a change greater than 50 percent of their local effect. On average, the at-scale intervention produces a smaller welfare effect by about 20 percent (only a 3.6 percentage point gain). However, not all households are worse off at scale: about 20 percent experience at-scale effects that exceed their gains from the local intervention.

The distributional implications underlying these differences turn out to be key. The local intervention is highly regressive: land-rich farmers experience an 8 percent real income gain, while land-poor farmers experience only a 2.5 percent gain. In contrast, we find that the at-scale intervention is significantly less regressive, as land-poor farmers do better at scale (their gains increase from 2.5 to 4 percent) while the land-rich fare worse (their gains drop from 8 to 6 percent). This is driven mostly by income effects rather than differential price index changes. Because land-rich households use modern inputs more intensively before the intervention, the income gains from the local subsidy (with output and factor prices mostly unaffected) are concentrated among this group. At scale, however, GE effects on average decrease the local market prices of modern input-intensive crops and increase the price of local labor. The resulting reduction in agricultural revenues and increase in labor compensation benefit households with higher initial reliance on wage labor relative to land-rich households. We document that these differences at scale are most pronounced among more remote regions, where local market prices are less constrained by large nearby centers or border crossings, and among crops and farmers with higher initial usage of modern inputs, where the asymmetry between partial equilibrium gains and GE effects on local prices is larger on average.

We then use our methodology to provide new insights relevant for experimental approaches to estimating GE effects. A growing literature employs "randomized saturation designs", which randomize not only treatment across individuals, but also the treatment saturation rate across geographic areas ("clusters"), to elicit GE effects with experimental variation (e.g. Baird *et al.* (2011), Burke *et al.* (2019), Egger *et al.* (2022)). Due to constraints on statistical power and feasibility of implementation, such designs often limit the comparison to two discrete levels of saturation, implemented within clusters that are typically villages or groups of villages. In order to identify the impact of policies at scale (e.g. at 100% national saturation), one must thus typically extrapolate from these two points of saturation, subject to two important assumptions: i) that GE forces are monotonic and linear with respect to changes in the saturation rate; and ii) that the GE forces experienced at the level of local clusters are representative of the effects of saturation at a broader geographical scale (e.g. nationwide).

We can assess the plausibility of these two assumptions by exploring how welfare implications evolve as a function of saturation rates at different geographical scales. On the reassuring side, our findings point to GE effects that are roughly linear as a function of the nationwide saturation rate: starting with the local intervention that treats 2.5% of farmers in each parish, we estimate how that original sample of farmers fare when the program is sequentially scaled up in steps of 10% of the remaining rural Ugandan population, up to 100% saturation. We find that the average gains to the initially treated farmers decline close to linearly as a function of scale-up to the rest of the country. This provides some reassurance about the lessons that can be drawn from designs relying on just two discrete saturation rates.

However, our results also suggest some caution about these designs. Because it is nearly impossible to randomize nationwide saturation rates, experiments typically randomize saturation rates at some lower, sub-national level. We find that the geographical scale of saturation meaningfully changes conclusions about the policy's impact. In our setting, we find that increases in saturation at the national level decrease the average rural welfare gains and flatten their regressivity; however, when we instead implement the same counterfactuals in steps of 10% of the population within *subcounties* (a large but feasible unit for randomization saturation),[9] we find almost no change in average welfare gains even at 100% saturation within the subcounty.[10] Our findings suggest caution when extrapolating from GE effects observed in designs that randomize saturation within smaller geographic units to the effects that would be observed at a broader

[9]Uganda is made up of roughly 800 subcounties. These are on the larger side of common definitions of "clusters," which we do here to be conservative.

[10]This is not due to the absence of GE forces under full sublocation saturation, but rather due to their different nature compared to at national scale; also affecting distributional conclusions. Sublocation saturation leads to a weaker reduction in the gains among rich households and a stronger increase among poorer households, which cancel out on average in a way that they do not under national saturation.

scale of rollout.

We conduct two additional counterfactuals to investigate the role of the granular economic geography that our model embraces, comparing our results to those using existing approaches in the literature. In a first comparison, we evaluate the welfare impact at scale in a setting without trading frictions – as if all households were selling into one integrated domestic market. Modeling a single market has been standard in an earlier literature using CGE models, as well as in a more recent literature in macroeconomics on quantifying the aggregation of local shocks if they were to occur to all agents in the economy (e.g. Buera *et al.* (2017), Sraer & Thesmar (2018), Fujimoto *et al.* (2019)). In a second comparison, we allow for trade costs, but instead consider the common workhorse structure of quantitative trade models (e.g., Costinot & Rodríguez-Clare (2014) and Baqaee & Farhi (2019)), featuring ad valorem trade costs and structural gravity with differentiated varieties – implying that all origin-destination market pairs engage in bilateral trade unless the costs of doing so are prohibitive (and thus remain unaffected by policy changes).[11] In both cases, we find meaningful differences in the average and distributional implications of the subsidy at scale compared to our framework, and discuss the mechanisms that are missed when imposing coarser assumptions about the economic geography.

We also explore model validation tests and the sensitivity of our findings across different modeling assumptions. One important innovation of our theory is to use the model-based price discovery algorithm to be able to solve the model with the new economic features we allow for in this setting. This involves solving for farm-gate prices at the level of households and trade flows that rationalize the observed consumption and production decisions given a graph of trade costs. For model validation, we are able to use data on market prices and trade flows for 260 Ugandan markets in the trader surveys collected by Bergquist *et al.* (2022). This allows us to assess to what extent the model-based estimates of local crop prices and predicted trading relationships capture variation in prices and trade flows of those same markets in the survey data. We also document our findings across parameter ranges that deviate from our preferred estimates on both the supply and demand sides of the model. While results do not vary strongly across alternative demand-side parameters in our application, the magnitude of the GE adjustments are sensitive to the estimated supply elasticities. This highlights the important role that RCTs and well-identified natural experiments can play in identifying key model parameters in a given policy environment. We conclude with a brief discussion of practical considerations when combining our toolkit with evidence based on fieldwork or quasi-experiments.

[11] Another important difference is in the question studied: whereas standard quantitative trade models aim to measure the aggregate welfare effect of a shock, we are interested in linking the quantitative analysis to the outcomes typically studied in impact evaluations: the average and distributional effects across individual farmers or households.

2 Model and Solution Method

We develop a rich but tractable quantitative model that is able to capture the granular economic geography of household location, production, consumption and trading that one can observe in the data, as well as a number of well-known features of agricultural trade that we also document in the microdata (see Appendix 1 and Appendix 2). These features deviate from the workhorse structure of quantitative trade and economic geography models: i) the vast majority of local markets do not trade with one another in a given crop, pointing to a limited degree of product differentiation within crops; ii) trade costs appear to be additive (charged per unit of weight) rather than multiplicative (ad valorem); iii) preferences are non-homothetic, with falling expenditure shares on food consumption as incomes rise; and iv) the adoption of modern inputs, such as chemical fertilizer or hybrid seeds, changes the relative cost shares of traditional inputs (land and labor).

In line with these features, our model features heterogeneous producers and consumers who interact across a realistic geography. The economy is populated by farmers who are endowed with land of heterogeneous suitability for different crops, which are modeled as homogeneous goods. Farmers trade both labor and crops in their nearest local market. These local markets are connected with all other markets and the rest of the world by a graph based on existing transport infrastructure. Our model allows trade costs between farmers and markets and between markets to have both an additive and an ad valorem (iceberg) component. Farmers are also allowed to choose between different production techniques, where the adoption of modern inputs may affect the production function with respect to traditional inputs. Preferences are non-homothetic, such that GE price changes in agriculture can affect initially richer or poorer households asymmetrically through the price index.

Environment

There are two kinds of agents: farmers indexed by $i \in \mathcal{I}$ and urban households indexed by $h \in \mathcal{H}$. There is also an agent that we call Foreign, which is indexed by F and stands for the rest of the world. In general, each of these agents in the economy is indexed by o (origin) or d (destination) when dealing with the trade network, and with $j \in \mathcal{J} \equiv \mathcal{I} \cup \mathcal{H} \cup \{F\}$ when dealing with agent behavior. To save on notation, we dispense for now with the notion of markets and just think of agents interacting directly with each other. We bring back such local markets when we impose particular restrictions on trade costs and labor migration in the final model section below.

Final goods are indexed by k and can be agricultural goods, $k \in \mathcal{K}_A$, or manufacturing goods, $k \in \mathcal{K}_M$. In turn, inputs (besides land) are indexed by n and can be intermediate goods used in agriculture, $n \in \mathcal{N}_I$, or labor used both in agriculture and manufacturing, $n = L$. We use g as a generic index that encompasses both final goods and inputs, $g \in \mathcal{G} \equiv \mathcal{K}_A \cup \mathcal{K}_M \cup \mathcal{N}_I \cup \{L\}$, and let $p_{j,g}$ denote the price at which agent j can buy or sell good g. We will refer to the collection of agents excluding Foreign as "Home", which will correspond to Uganda in our quantitative analysis.

Farmers own land and labor in quantities Z_i and L_i, and they produce agricultural goods (crops) using their own land (i.e., land is not tradable) as well as labor and intermediate goods (such as fertilizer and seeds).[12] Urban households own labor in quantity L_h and produce a manufacturing good using labor.

Trade in good g from o to d is subject to iceberg and additive trade costs. Iceberg trade costs are $\tau_{od,g} \geq 1$ and additive trade costs are $t_{od,g} \geq 0$ in units of a "transportation good". We assume that this good is produced by Foreign and that there are no trade costs for this good, so that all agents can access it at the same price.[13] Setting this price equal to one by choice of numeraire, $t_{od,g}$ becomes the actual additive transportation cost from o to d for good g. Thus, for example, if agent j buys good g from farmer i then her price is $p_{j,g} = \tau_{ij,g} p_{i,g} + t_{ij,g}$. We assume that these trade costs satisfy the triangular inequality: $\tau_{od} \leq \tau_{oo'} \cdot \tau_{o'd}$ and $t_{od} \leq t_{oo'} + t_{o'd}$ for any o, o', d.

For manufacturing goods we follow the convention in the trade literature and assume that they only face iceberg transportation costs, hence $t_{od,g} = 0$ for all $g \in \mathcal{K}_M$. Similarly, as in the Armington model of trade, we assume that each urban household as well as Foreign produce a differentiated manufacturing good, and use $g(h)$ to refer to the manufacturing good produced by urban household h and $g(F)$ to refer to the manufacturing good produced by Foreign.

We assume that Home is "small" in the sense that the prices of goods produced in Foreign (i.e., crops, intermediate goods and Foreign's manufacturing good) are exogenous and given by $p^*_{F,g}$, while Foreign's demand for the manufacturing goods associated with any of our economy's urban centers is not affected by any variables in Home other than its price. In the case of intermediate goods we go one step further and assume that they are imported from Foreign at exogenous prices $p^*_{i,n}$ for all $i \in \mathcal{I}$ and $n \in \mathcal{N}_I$ – this provides the needed flexibility to consider counterfactuals in which arbitrary subsets of farmers experience declines in fertilizer prices through the implementation of a government program or RCT.[14]

Finally, regarding notation, we use $\{x_{ij}\}$ to denote the vector of some variable x_{ij} for all combinations of indices i and j, and $\{x_{ij}\}_i$ to denote the vector of x_{ij} for the given i and for all j.

Next, we turn to preferences, technology and equilibrium. To simplify the exposition, we present the model imposing the specific functional forms on preferences and technology that we will use in our quantitative analysis. However, we emphasize that the quantitative approach developed below can be used with other functional forms as long as long as they satisfy a set of common properties that we lay out in Appendix 4.B.

[12] We model land as not tradable in line with empirical evidence showing that land markets in sub-Saharan Africa and other low-income regions are generally thin, with sparse rental markets and in some cases "almost non-existent" transactions (Acampora *et al.* , 2022) (see also e.g. Deininger *et al.* (2008); Holden *et al.* (2010)).

[13] This implies that the policies we study do not lead to additional GE effects through changing (endogenous) transportation costs in the country.

[14] We thus focus on the impact of input subsidies on farmers, and ignore potential knock-on effects on domestic production of those inputs.

Preferences

We assume here non-homothetic preferences in the form of Stone-Geary demand for consumption of agricultural and manufacturing goods, so that households need to consume a minimum amount of a composite agricultural good, $\bar{C}_A$. This composite is a CES-aggregate of the consumption of individual agricultural goods with elasticity of substitution σ, while individual manufacturing goods are similarly aggregated with elasticity of substitution η. Letting $\xi_{j,k}$ denote the expenditure share of agent j on good k and $\xi_k\left(\{b_{j,k}p_{j,k}\}_j, I_j\right)$ be the corresponding expenditure share function (assumed common across all agents in Home), we then have

$$\xi_{j,k} = \xi_k\left(\left\{b_{j,k}p_{j,k}\right\}_j, I_j\right) = \begin{cases} \frac{\left(b_{j,k}p_{j,k}\right)^{1-\sigma}}{P_{j,A}^{1-\sigma}}\left(\zeta + (1-\zeta)\frac{P_{j,A}\bar{C}_A}{I_j}\right) & \text{for } k \in \mathcal{K}_A \\ \frac{\left(b_{j,k}p_{j,k}\right)^{1-\eta}}{P_{j,M}^{1-\eta}}(1-\zeta)\left(1-\frac{P_{j,A}\bar{C}_A}{I_j}\right) & \text{for } k \in \mathcal{K}_M. \end{cases}$$

Here $\{b_{j,k}\}$ are demand shifters and I_j is income of agent j. Price indices $P_{j,A}$ and $P_{j,M}$ correspond to the CES-aggregates for agriculture and manufacturing of agent j, respectively.

Turning to Foreign, our small-open economy assumption for Home implies that Foreign's demand (in value) for manufacturing good $g(h)$ can be specified directly as a function of this goods's individual price, $X_{F,g(h)}\left(p_{F,g(h)}\right)$. We assume that this is given by

$$X_{F,g(h)}\left(p_{F,g(h)}\right) = D_{F,g(h)}p_{F,g(h)}^{1-\eta},$$

where $D_{F,g(h)}$ is some constant.

Technology

Farmers produce agricultural goods $k \in \mathcal{K}_A$ using land, labor and intermediate goods with techniques $\omega \in \Omega$. The production function for a farmer i producing good k with technique ω is assumed Cobb-Douglas with cost share $\alpha_{i,n,k,\omega}$ for input $n \in \mathcal{N}_I \cup \{L\}$. We assume that $\sum_n \alpha_{i,n,k,\omega} < 1$ and let $\alpha_{i,z,k,\omega} \equiv 1 - \sum_n \alpha_{i,n,k,\omega}$ be the corresponding cost share of land.

Letting $r_{i,k,\omega}$ denote the return to an effective unit of land allocated by farmer i to produce agricultural good k with technique ω, then at an interior solution to the farmer's optimization problem we must have

$$a_{i,k,\omega}p_{i,k} = c_{i,k,\omega}(\{p_{i,n}\}_i, r_{i,k,\omega}) \equiv \left(\prod_n p_{i,n}^{\alpha_{i,n,k,\omega}}\right) r_{i,k,\omega}^{\alpha_{i,z,k,\omega}}, \tag{1}$$

where $a_{i,k,\omega}$ is a Hicks-neutral productivity shifter. Equation (1) determines $r_{i,k,\omega}$ as an implicit function of prices, $p_{i,k}$ and $\{p_{i,n}\}_i$, and productivity $a_{i,k,\omega}$.

Farmer i allocates land endowment Z_i across different agricultural goods (or simply "crops") and techniques to maximize their total land returns, $\sum_{k,\omega} r_{i,k,\omega}Z_{i,k,\omega}$, where $Z_{i,k,\omega}$ measures the effective units of land allocated by farmer i to produce crop k with technique ω. We allow for decreasing marginal productivity in how physical units of land Z_i can be converted into efficiency units of land for different crops and techniques, with possibly a different elasticity of substitution between crops and techniques. Specifically, we assume that the feasible set for the allocation of efficiency units of land across crops and techniques satisfies

$$\left(\sum_k \left(\sum_\omega Z_{i,k,\omega}^{\kappa/(\kappa-1)}\right)^{\frac{\mu}{\mu-1}\frac{\kappa-1}{\kappa}}\right)^{\frac{\mu-1}{\mu}} \leq Z_i,$$

where κ is the elasticity governing the allocation of land across techniques within a given crop in the lower nest, while μ is the elasticity governing the allocation of land across crops in the upper nest.[15] Letting $\pi_{i,k,\omega}$ denote the share of land returns of farmer i coming from production of crop k with technique ω, one can show that[16]

$$\pi_{i,k,\omega} = \pi_{k,\omega}\left(\{r_{i,k,\omega}\}_i\right) \equiv \frac{r_{i,k,\omega}^\kappa}{\sum_\omega r_{i,k,\omega}^\kappa} \frac{\left(\sum_\omega r_{i,k,\omega}^\kappa\right)^{\mu/\kappa}}{\sum_k \left(\sum_\omega r_{i,k,\omega}^\kappa\right)^{\mu/\kappa}},$$

with total returns to land given by

$$Y\left(\{r_{i,k,\omega}\}_i\right) = \left(\sum_k \left(\sum_\omega r_{i,k,\omega}^\kappa\right)^{\mu/\kappa}\right)^{1/\mu}.$$

Finally, letting $q_{i,k,\omega}$ denote output of crop k for farmer i with technique ω, then

$$q_{i,k,\omega} = q_{i,k,\omega}\left(\{p_{i,g}\}_i, \{r_{i,k,\omega}\}_i\right) = \frac{\pi_{k,\omega}\left(\{r_{i,k,\omega}\}_i\right) Y\left(\{r_{i,k,\omega}\}_i\right) Z_i}{\alpha_{i,z,k,\omega} p_{i,k}}.$$

Turning to urban households, we assume that each urban area is associated with a single representative urban household who produces a differentiated manufacturing good. We keep the technology simple by assuming that manufacturing production is linear in labor, so that the quantity of manufacturing good $g(h)$ produced by urban household h is $a_h L_h$. Given that labor supply is perfectly inelastic, we can then treat $q_h \equiv a_h L_h$ as the urban households' endowment

[15]This is a nested constant elasticity of transformation production function as in Powell & Gruen (1968), De Melo (1988), Burfisher (2021). One can also verify that this can be obtained from an extension of the Roy-Frechet microfoundations in Costinot & Donaldson (2016) and Sotelo (2020), but now allowing for a nested Frechet structure, as in Farrokhi & Pellegrina (2022). In particular, assuming that farmer i has a continuum of plots of land with measure Z_i, and that each plot of land has productivities $X_{i,k,\omega}$ independently drawn from the joint distribution $H(\boldsymbol{x}_i) = \exp\left(-\gamma^{-1}\sum_k\left(\sum_\omega x_{i,k,\omega}^{-\kappa}\right)^{\mu/\kappa}\right)$ with $\gamma = \Gamma(1-1/\mu)$, then this would lead to the production function above. The Roy-Frechet microfoundations would imply the restriction $1 < \mu \leq \kappa$, so that the density is always positive and the mean is well defined, but this is not necessary for the more general case of a nested CES PPF that we work with here.

[16]With a slight abuse of notation, for all vectors associated with farmers' production, we write $\{X_{i,k,\omega}\}_i$ for the vector $\{X_{i,k,\omega}\}_{i,k\in\mathcal{K}_A}$ for any variable X.

of manufacturing good $g(h)$.

Equilibrium

In equilibrium, rural and urban households maximize utility taking prices as given, prices respect no-arbitrage conditions given trade costs, and all markets clear. We assume that markets are competitive, but potentially subject to a rich and granular set of frictions in the transactions between agents that we capture by allowing for (additive and ad valorem) agent- and good-specific trading costs in all input and output markets.[17]

To formalize the definition of equilibrium, let $\chi_{j,g}\left(\{b_{j,k}p_{j,k}\}_j, \{r_{j,k,\omega}\}_j, \{p_{j,g}\}_j, I_j\right)$ be the excess demand function (in value) of agent j for good g given demand shifters, prices, returns, and income, and let $\chi_{F,g}(\bullet)$ be the corresponding excess demand function for Foreign. The excess demand functions $\chi_{j,g}(\bullet)$ for farmers, urban households and Foreign are determined by the results in the previous subsections, and can be found in Appendix 4.A. The equilibrium is a set of prices, $\{p_{j,g}\}$ and trade flows $\{x_{od,g}\}$ (measured in quantity at the destination), such that $p_{j,g} = p^*_{j,g}$ for all $j \in \mathcal{I}$ and all $g \in \mathcal{N}_I$, $p_{F,g(F)} = p^*_{F,g(F)}$, excess demand is equal to the difference between purchases and sales for each agent j and good $g \in \mathcal{K}_A \cup \mathcal{K}_M \cup \{L\} \setminus \{g(F)\}$,

$$\chi_{j,g}\left(\{b_{j,k}p_{j,k}\}_j, \{r_{j,k,\omega}\}_j, \{p_{j,g}\}_j, I_j\right) = p_{j,g}\left(\sum_o x_{oj,g} - \sum_d \tau_{jd,g}x_{jd,g}\right), \tag{2}$$

and no-arbitrage conditions hold for all $g \notin \mathcal{N}_I$,

$$\tau_{od,g}p_{o,g} + t_{od,g} \geq p_{d,g} \perp x_{od,g}, \quad \forall o, d, \tag{3}$$

with farmer i's income equal to the sum of land returns and wage income $p_{i,L}L_i$,

$$I_i = Y_i\left(\{r_{i,k,\omega}\}_i\right) Z_i + p_{i,L}L_i, \quad \forall i \in \mathcal{I}, \tag{4}$$

urban household h's income given by

$$I_h = p_{h,g(h)}q_h, \quad \forall h \in \mathcal{H}, \tag{5}$$

and $r_{i,k,\omega}$ satisfying (1) $\forall i \in \mathcal{I}, k \in \mathcal{K}_A, \omega \in \Omega$.[18] Here the symbol $\perp$ between a weak inequality and a variable indicates that the weak inequality holds as equality if the variable is strictly positive. For example, if farmer i sells crop k to agent j then $x_{ij,k} > 0$ equation (3) implies that

[17] Note that trading frictions are also present in local labor markets when farmers are hiring or selling labor. The presence of additive trade costs also implies that pass-through is not log linear. This leads to richer comparative statics than in models with only iceberg trade costs and perfect competition or even monopolistic competition with fixed markups.

[18] We can exclude the manufacturing good produced in Foreign from the set of equilibrium conditions since for this good we know that $p_{j,g(F)} = \tau_{Fj,g(F)}p^*_{F,g(F)}$ for all j. Also, in parallel to our treatment of land for farmers, we assume that there is no market for household labor in urban areas, and hence the equilibrium system does not have to determine the price of this good.

$\tau_{ij,k}p_{i,k} + t_{ij,k} = p_{j,k}$, while if $\tau_{ij,k}p_{i,k} + t_{ij,k} > p_{j,k}$ then equation (3) implies that $x_{ij,k} = 0$.

The equilibrium conditions across all crops, manufacturing goods and labor imply that there is trade balance, which is given by the condition that Foreign runs a deficit in goods that is paid for by the economy's total expenditure on the transportation good (which is an income to Foreign).

Solution of Counterfactuals

We are interested in computing the effect of shocks to technology (e.g. due to climate change or weather shocks), intermediate good prices (e.g. due to government subsidies or extension programs), or changes in trade costs (e.g. due to rural road building), all for the agricultural sector. Using hat notation (i.e., $\hat{x} = x'/x$), these shocks are given by $\{\hat{a}_{j,k,\omega}\}$,$\{\hat{p}^*_{i,n}\}_{n\in\mathcal{N}_I}$ and $\{\hat{\tau}_{od,k}, \hat{t}_{od,k}\}$. In the counterfactual equilibrium, equations (1)-(5) can be written as

$$\hat{p}_{i,k}p_{i,k} = c_{i,k,\omega}(\{\hat{p}_{i,n}p_{i,n}\}_i, \hat{r}_{i,k,\omega}r_{i,k,\omega})/\hat{a}_{i,k,\omega}a_{i,k,\omega},\ \forall i \in \mathcal{I}, k \in \mathcal{K}_A, \omega \in \Omega, \tag{6}$$

$$\begin{aligned}\chi_{j,g}\left(\{b_{j,k}\hat{p}_{j,k}p_{j,k}\}_j, \{\hat{r}_{j,k,\omega}r_{j,k,\omega}\}_j, \{\hat{p}_{j,g}p_{j,g}\}_j, \hat{I}_jI_j\right) = \\ = \hat{p}_{j,g}p_{j,g}\left(\sum_o x'_{oj,g} - \sum_d \hat{\tau}_{jd,g}\tau_{jd,g}x'_{jd,g}\right),\end{aligned} \tag{7}$$

$$\hat{\tau}_{od,g}\tau_{od,g}\hat{p}_{o,g}p_{o,g} + \hat{t}_{od,g}t_{od,g} \geq \hat{p}_{d,g}p_{d,g} \perp x'_{od,g},\ \forall o, d, \tag{8}$$

$$\hat{I}_i = \left[Y\left(\{\hat{r}_{i,k,\omega}r_{i,k,\omega}\}_i\right)/Y\left(\{r_{i,k,\omega}\}_i\right)\right](1-\lambda_{i,L}) + \hat{p}_{i,L}\lambda_{i,L},\ \forall i \in \mathcal{I}, \tag{9}$$

$$\hat{I}_h = \hat{p}_{g(h)},\ \forall h \in \mathcal{H}, \tag{10}$$

with $\hat{p}_{i,n} = \hat{p}^*_{i,n}$ for all $i \in \mathcal{I}$ and $n \in \mathcal{N}_I$, and where $\lambda_{i,L} = p_{i,L}L_i/I_i$ is the share of farmer's total income coming from wage income.

Our assumptions imply that we can use exact-hat algebra (Dekle *et al.* , 2007) to solve for the endogeneous changes in the prices of manufacturing goods produced in Home, $\{\hat{p}_{j,g}\}$ for all $g \in \mathcal{K}_M \setminus \{g(F)\}$ (recall that $p_{F,g(F)}$ is fixed at $p^*_{F,g(F)}$). To see this, start by noting that since there are no additive trade costs in manufacturing then equation (3) implies that if $x_{od,g} > 0$ then $\tau_{od,g}p_{o,g} = p_{d,g}$. Adding up equation (2) over all j implies that our equilibrium system entails

$$\sum_j \chi_{j,g}\left(\{b_{j,k}p_{j,k}\}_j, \{p_{j,g}\}_j, I_j\right) = 0, \quad \forall g \in \mathcal{K}_M \setminus \{g(F)\},$$

where we have dropped $\{r_{j,k,\omega}\}_j$ from the argument of $\chi_{j,g}$ since land returns do not affect excess demand for manufacturing goods (conditional on income). The counterfactual version of this equation is

$$\sum_j \chi_{j,g}\left(\{b_{j,k}\hat{p}_{j,k}p_{j,k}\}_j, \{\hat{p}_{j,g}p_{j,g}\}_j, \hat{I}_jI_j\right) = 0, \quad \forall g \in \mathcal{K}_M \setminus \{g(F)\}. \tag{11}$$

We can use exact-hat algebra to compute the left-hand side of this equation as a function of hat changes in prices and income levels given data on income levels, expenditure shares, and Home exports of manufacturing goods.[19] Given that $\hat{p}_{j,g(h)} = \hat{\tau}_{hj,g(h)}\hat{p}_{h,g(h)}$ and $\hat{I}_h = \hat{p}_{h,g(h)}$ for all j and h, and taking as given counterfactual prices of agricultural goods and counterfactual income levels of farmers (which are solved separately as explained below), then equation (11) constitutes a system of equations – one for each $g(h), h \in \mathcal{H}$ – that we can use to solve for the hat changes in the prices of the manufacturing goods produced in Home, $\{\hat{p}_{h,g(h)}\}$.

For agricultural goods and labor in rural production we have additive trade costs and so the first step that we followed above for manufacturing goods does not give us an equation like (11). Moreover, since these are homogeneous goods then prices are not directly pinned down by the price at their origin, which is no longer predetermined. Formally, we need to deal with the fact that the right-hand side of equation (7) as well as equation (8) are in terms of counterfactual *levels*, and so we cannot use exact-hat algebra – we need information on the full vector of prices in the initial equilibrium $\{p_{j,g}\}$ for all $j \in \mathcal{I} \cup \mathcal{H}$ and $g \in \mathcal{K}_A \cup \{L\}$, and corresponding trade costs to solve the system. We next explain how we can recover these prices in a manner that is consistent with the model and the microdata.

As we discuss in Section 4, from the microdata we can either observe or directly infer the following set of variables: expenditure shares on agricultural goods for farmers and urban households, $\{\xi_{i,k}, \xi_{h,k}\}_{k\in\mathcal{K}_A}$, Foreign crop prices, $\left\{p^*_{F,k}\right\}_{k\in\mathcal{K}_A}$, physical crop output and cost shares for farmers, $\{q_{i,k,\omega}\}_{k\in\mathcal{K}_A}$ and $\{\alpha_{i,n,k,\omega}\}_{k\in\mathcal{K}_A}$, labor endowments of farmers, $\{L_i\}$, income of urban households $\{I_h\}$, and trade costs $\{t_{od,g}\}, \{\tau_{od,k}\}$.[20] We denote this set of observable variables used for price discovery in agriculture by

$$\mathbb{D}_A = \left\{\{\xi_{i,k}, \xi_{h,k}, p^*_{F,k}, q_{i,k,\omega}, \alpha_{i,n,k,\omega}\}_{k\in\mathcal{K}_A}, L_i, I_h, t_{od,g}, \tau_{od,k}\right\}.$$

Assuming that all these variables come from the initial equilibrium in our model, we can now rewrite excess demand functions for agricultural goods and labor (i.e., $\chi_{j,g}(\bullet)$ for $g \in \mathcal{K}_A \cup \{L\}$) for farmers, urban households and Foreign as functions of prices $\{p_{j,g}\}_{g\in\mathcal{K}_A\cup\{L\}}$ and data $\mathbb{D}_A$. We can then "discover" agricultural goods' prices and wages $\{p_{j,g}\}_{g\in\mathcal{K}_A\cup\{L\}}$ in the initial equilibrium

[19] Under our assumption on preferences and technology, $\chi_{j,g}\left(\{\hat{p}_{j,k}b_{j,k}p_{j,k}\}_j, \{\hat{p}_{j,g}p_{j,g}\}_j, \hat{I}_j I_j\right)$ can be evaluated as a function of $\{\hat{p}_{j,g}\}_j$ and $\hat{I}_j$ given data on expenditure shares of agent j on all goods k, Home exports of good g, and income I_j. We have data for income levels of urban households and expenditure shares on agriculture goods, while data on exports and expenditure shares for each manufacturing good produced in Home are inferred from trade costs and aggregate manufacturing exports, revenues and expenditure. As explained further in Appendix 4.E, this is a standard procedure in the trade literature when dealing with intra-national trade flows (see for example Donaldson & Hornbeck, 2016 and Faber & Gaubert, 2019). Income levels of farmers are obtained in the price discovery step for agriculture described below.

[20] Additive trade costs are relevant only for agricultural goods and rural labor, while iceberg trade costs are relevant for agricultural and manufacturing goods and labor, but we do not make this explicit to avoid notation clutter.

as a solution to the following system of equations for $g \in \mathcal{K}_A \cup \{L\}$:

$$\chi_{j,g}\left(\{p_{j,g}\}_{g\in\mathcal{K}_A\cup\{L\}};\mathbb{D}_A\right) = p_{j,g}\left(\sum_o x_{oj,g} - \sum_d \tau_{jd,g}x_{jd,g}\right)\forall j, \tag{12}$$

$$\tau_{od,g}p_{o,g} + t_{od,g} \geq p_{d,g} \perp x_{od,g},\ \forall o, d. \tag{13}$$

where these excess-demand functions are again formally presented in Appendix 4.A. Given crop prices and data $\mathbb{D}_A$ we can in turn compute farmer income levels, $\{I_i\}$, and shares of land returns by crop and technique, $\{\pi_{i,k,\omega}\}$.[21]

In Appendix 4.D, we describe how to transform this price discovery step into an equivalent problem of finding the equilibrium of an exchange economy that is integrated as a small open economy with the rest of the world. We then show that, if there are no additive trade costs, the goods in such an economy satisfy the connected substitutes condition in Berry *et al.* (2013) and hence there is a unique equilibrium in which all agents are directly or indirectly connected through trade. This implies that there is a unique (connected) solution to the price-discovery step. Although we can no longer establish uniqueness analytically if there are additive trade costs, our numerical analysis suggests that this is indeed the case in our context.[22]

Finally, we can obtain counterfactual trade flows $\left\{x'_{od,g}\right\}$ and prices changes $\{\hat{p}_{j,g}\}$ as a solution to the system of equations (6)-(10) given shocks $\{\hat{a}_{j,k,\omega}\}$,$\{\hat{p}^*_{i,n}\}_{n\in\mathcal{N}_I}$ and $\left\{\hat{\tau}_{od,k}, \hat{t}_{od,k}\right\}$, data $\mathbb{D}_A$ (used for the price discovery step in agriculture), and data on manufacturing exports by Home as well as expenditure shares in manufacturing, $\{\xi_{i,k}, \xi_{h,k}\}_{k\in\mathcal{K}_M\setminus\{g(F)\}}$ (used for the counterfactual analysis in manufacturing).[23]

Additional Assumptions Used in Application

Consistent with the data described in Appendix 1 and Appendix 2, we allow input shares to vary not only across crops and Ugandan regions but also across techniques within crops. We introduce two techniques for each crop: traditional, $\omega = 0$, and modern, $\omega = 1$. We will map these two techniques to data in terms of observed use of modern intermediates (chemical fertilizer and hybrid seeds in our setting) in production: the traditional technique makes use of land and labor whereas the modern technique also makes use of the intermediate goods. Formally, $\alpha_{i,n,k,1} > 0 = \alpha_{i,n,k,0},\ \forall i, n \in \mathcal{N}_I, k$. Thus, the choice of a modern technique will increase the im-

[21]We obtain farmer income levels as $I_i = \sum_{k\in\mathcal{K}_A,\omega}\left(1 - \sum_n \alpha_{i,n,k,\omega}\right)p_{i,k}q_{i,k,\omega} + p_{i,L}L_i$, and shares of land returns by crop and technique as $\pi_{i,k,\omega} = \left(1 - \sum_n \alpha_{i,n,k,\omega}\right)p_{i,k}q_{i,k,\omega}/\left(I_i - p_{i,L}L_i\right)$.

[22]The sufficient conditions in our proof of uniqueness no longer hold in the presence of additive trade costs because the demand for foreign goods is no longer necessarily increasing with the price of domestic goods. In lieu of an analytical proof of uniqueness, we explore it numerically by considering 100 different initial guesses for prices drawn randomly along the range of possible prices given the exogenous international prices and trade costs. Reassuringly, we find the same equilibrium in all cases.

[23]The system of equations (6)-(10) also includes values for $\{r_{i,k,\omega}\}$ and $\{p_{i,n}\}$, which we do not observe. However, we can again use exact-hat algebra to evaluate the excess demand function on the LHS of (7) by using land-rent shares $\{\pi_{j,k,\omega}\}_j$ (which we obtain from the price discovery step in agriculture), and similarly evaluate $c_{i,k,\omega}(\{\hat{p}_{i,n}p_{i,n}\}_i, \hat{r}_{i,k,\omega}r_{i,k,\omega})$ using exact-hat algebra by using cost shares $\{\alpha_{i,n,k,\omega}\}_{i,k,\omega}$ in $\mathbb{D}_A$.

portance of intermediates and decrease the importance of land or labor (as well as the relative shares of the traditional inputs as observed in the data).

Our application will include millions of farmers, so we make a series of assumptions on trade costs to ensure that the counterfactual analysis described above is computationally feasible. First, consistent with the data, we assume that trade in agricultural goods is subject to additive trade costs, $t_{od,g} \geq 0$ and $\tau_{od,g} = 1$ for $g \in \mathcal{K}_A$. Second, we assume that trade costs have a hub-and-spoke structure, with each individual agent being directly connected to only one local market (hub). Formally, we denote markets by m and let $\mathcal{J}(m)$ denote the set of agents connected with market m. Trade costs between any two agents $j \in \mathcal{J}(m)$ and $j' \in \mathcal{J}(m')$ satisfy

$$\tau_{jj',g} = \tau_{jm,g} \cdot \tau_{mm',g} \cdot \tau_{m'j',g} \tag{14}$$

and

$$t_{jj',g} = t_{jm,g} + t_{mm',g} + t_{m'j',g}. \tag{15}$$

This assumption on trade costs allows us to define market-level prices from the prices of agents belonging to that market. In particular, if $j \in \mathcal{J}(m)$ is a net seller of good g then the market m price of good g is given by $p_{m,g} = \tau_{jm,g} p_{j,g} + t_{jm,g}$, while if $j \in \mathcal{J}(m)$ is a net buyer of good g then $p_{m,g}$ is such that $p_{j,g} = \tau_{mj,g} p_{m,g} + t_{mj,g}$. In Appendix 4.F, we show that these market-level prices are well defined in the sense that each of these equations yields the same price. In our application we will refer to the markets where farmers live as parishes in Uganda and to the markets where urban households live as cities.

Third, we assume that markets trade on a fully connected graph based on Uganda's road network as well as the location of border crossings with Foreign. This means that the trade cost between any two markets can be computed as the product (for iceberg trade costs) or sum (for additive trade costs) of trade costs along a sequence of markets that are directly connected by a road or by a border crossing in the case of Foreign. Finally, we assume that labor markets are local with prohibitive costs of selling or hiring labor across markets. While we thus abstract from migration in our application below,[24] the model we develop above can readily incorporate it (see Appendix 4.G).

3 Combining the Model with Local Experiments

The complex forces governing how shocks propagate across markets described in the model above are difficult to capture in local experiments, which are typically either too small to give rise to GE forces that emerge at scale or use variation in relative exposures for identification (with parts of GE forces absorbed by intercepts or fixed effects). However, local experiments can play a critical role in informing the estimation of policy impacts at scale. At-scale counterfactu-

[24] Meaningful migration responses have not been found empirically in the context of the typical agricultural policies we consider here (e.g. Huntington & Shenoy (2021)), or in the context of broader shocks to agricultural productivity due to extreme weather events (e.g. Emerick & Burke (2016) and Emerick (2018)).

als in part depend on the locations, trade costs, and output and input markets that are directly exposed by the policy (see next section). But they are also determined by a number of key elasticities that govern responses to direct policy exposure on both the supply and demand sides of the economy. Local experiments and quasi-experiments can offer a critical improvement in the identification of these responses, relative to a calibration exercise that is based on observational variation.

In this section we describe the estimation of the elasticities governing demand (ζ, σ, η) and supply (κ, μ), and show how to exploit the advantage of local experiments in combination with the model. Though we will use local experiments from a variety of East African countries – Uganda, Kenya, and Mozambique – we believe using well-identified estimates from experiments conducted in the region offer a large advance over calibration using existing estimates in the literature, which are mostly drawn from outside the region and/or outside of agriculture.[25] As discussed as part of the counterfactual analysis, we also explore the findings across a range of alternative parameter values and model assumptions. Appendix 1 provides additional details about the data used in the estimation.

Demand Estimation

To estimate our key demand parameter, σ, the elasticity of substitution between crops in consumption, we bring to bear a demand-side experiment conducted in Bergquist & Dinerstein (2020). This experiment was conducted in open-air maize markets in rural western Kenya, 30km from the Ugandan border. In their experiment, individual consumers who approached maize traders to make a purchase were offered a price discount, the size of which was randomized across ten possible amounts. The value of the discount ranged from from roughly 0-15% of the baseline price. Using the subsidy as exogenous variation in consumer prices, the experiment measured resulting quantities purchased.

To estimate σ in the model, we run the following specification:

$$\log x_{i,m,sd} = \alpha + \beta \log p_{i,m,sd} + \theta_{m,sd} + \epsilon_{i,m,sd},$$

regressing log quantity purchased by individual i from seller s in market m on date d on log price, instrumenting for price with the randomized subsidy amount. Because the subsidy was randomized across consumers buying from the same seller in the same market-day, we run specifications including either market-by-date fixed effects ($\theta_{m,d}$) or seller-by-market-by-date fixed effects ($\theta_{m,sd}$), presented in Columns 2 and 4 of Table 1, respectively. Both specifications yield estimates close to -1. We therefore calibrate our model with $\sigma = 1$.

One possible limitation of the above experiment is that the subsidies were fairly short-run in their duration. One may worry that the short-term demand elasticities estimated here do not map well to the long-term demand elasticities that are presumably more relevant for shaping the

[25] Rural areas across East Africa share many features, including crops grown, farming methods (mostly rain-fed agriculture), and overall levels of development.

impact of policies at scale. Bergquist & Dinerstein (2020) tackle this issue explicitly, exploiting the randomized order of their treatment periods to test for evidence of inter-temporal dynamics in demand (see Appendix C of Bergquist & Dinerstein (2020)). They find limited evidence of stockpiling, which they attribute to the relative infrequency of storage in their setting (Burke *et al.* , 2019). We are therefore less concerned about this issue in our setting (though we do explore the sensitivity of the counterfactual analysis across a range of higher or lower values for σ in Section 5). However, more generally, RCTs do tend to be relatively short-run in their duration. There may therefore be a trade-off between improved identification vs. temporal mismatch when using RCTs to estimate these key elasticities, rather than observational variation. The recent push within the experimental literature to run more long-run experiments is promising for the future availability of longer-run elasticity estimates for scale-up projections (Bouguen *et al.* , 2019).

To calibrate the demand parameter ζ, that governs non-homotheticity in food consumption, we use the following relationship that holds subject to utility maximization under Stone-Geary:

$$\frac{P_{i,A}\bar{C}_A}{I_i} = \frac{\xi_{i,A} - \zeta}{(1-\zeta)},$$

where the left-hand side is the share of household income spent on subsistence, and $\xi_{i,A}$ is the observed share spent on total food consumption, $\xi_{i,A} = \sum_{g \in \mathcal{K}_A} \xi_{i,g}$. We use the typical feature of these preferences that the share of income spent on subsistence approaches zero for the richest households, setting the left-hand side equal to zero, and calibrating ξ_A with the average share of expenditure spent on total food consumption among the richest 5 percent of Ugandan households (which is close to 0.1 in the UNPS data). This yields an estimate of $\zeta = 0.1$, implying that the share spent on subsistence is on average 38 percent across Ugandan households. For the elasticity of substitution across manufacturing varieties we choose $\eta = 5$, in line with the literature in international trade.

Supply Estimation

Applying the model described above to the common example of input subsidy programs, we show in this section how one can use a small-scale RCT of a fertilizer subsidy program to estimate the first key supply elasticity κ, which governs farmers' choice of land allocation across modern or traditional planting technologies within crops. To do so, we advantage of the experiment run in Carter *et al.* (2020), in which randomly selected farmers in Mozambique were offered fertilizer and improved seeds at a subsidized price. Data collected on farmers' use of modern technologies and output by plot allows the estimation of the impact of changing input prices (instrumented by treatment) on land allocations across technologies. To estimate κ, we derive the following estimation equation from Section 2:

$$\log\left(\frac{\pi_{i,k,1}}{\pi_{i,k,0}}\right) = -\left(\kappa\frac{\alpha^{input}_{i,k,1}}{\alpha^{land}_{i,k,1}}\right)\log p^{input}_{i,k} + \epsilon_{i,k},$$

where we have the relative land allocations of modern vs traditional production techniques within maize production on the left-hand side, and the log price of intermediates ($p^{input}_{i,k}$) on the right-hand side. The extent to which a price shock for modern inputs affects land allocations across production techniques within crops will be a function of the supply elasticity in the lower nest, κ, as well as the relative cost shares of intermediates and land in modern production, $\alpha^{input}_{i,k,1}$ and $\alpha^{land}_{i,k,1}$ respectively.

Using the data in Carter *et al.* (2020), we construct a price index for intermediates as the weighted average of prices of chemical fertilizer and hybrid seeds, with weights proportional to their relative cost shares. We then instrument this price with the randomized subsidy treatment.[26] Table 2 presents the estimation results of the first stage, reduced form and the IV point estimates. For each, we report results both from a single post-treatment cross-section or using baseline and post-treatment panel data with round and community fixed effects.[27] The IV point estimate in columns 5 and 6 is 0.83 and 0.85. Using the ratio of cost shares of land over fertilizer and hybrid seeds, this implies that $\kappa = 2.5$. We use this estimate of the lower-nest (within-crop) elasticity as our baseline, and explore the sensitivity of the counterfactual analysis across a range of higher or lower values for κ.

We complement this RCT with a natural experiment in Uganda which allows us to estimate the upper-tier supply elasticity in our model for substitution of land allocations across crops, μ.[28] The estimation equation derived from the parametrization in Section 2 above is as follows:

$$\log\left(\sum_{\omega'}\pi^{-1}_{i,\omega',t|k}\left(\frac{q_{i,k,\omega',t}}{\prod_n l^{\alpha_{i,n,k,\omega'}}_{i,n,k,\omega',t}}\right)^{\frac{1}{\alpha^{land}_{i,k,\omega'}}\frac{\kappa}{\kappa-1}}\right)^{\frac{\kappa-1}{\kappa}} = \left(\frac{\mu-1}{\mu}\right)\log\pi_{i,k,t} + \log Z_{i,t} + \log\tilde{b}_{i,k,t} \quad (16)$$

The left-hand side of (16) is farmer i's harvest quantities for crop k aggregated across both techniques in survey year t (summed across both seasons) adjusted for the reported quantities of labor, modern intermediates ($l_{i,n,k,\omega,t}$) and the share of land allocated to technique ω conditional on producing crop k ($\pi_{i,\omega,t|k}$). This represents an observable measure of land productivity for a crop k and farmer i as the harvest amounts we observe under either production technique are deflated by the inputs used across all plots of land allocated to crop k. The first term on the

[26]Given these data record just one snapshot of production, where some farmers were allocating 100% of production to either modern or traditional techniques, we aggregate both left and right-hand sides to the level of local villages broken up by treatment status, summing land allocations on the left and taking average prices on the right. This is to avoid the assumption that those farmers could never make use of the other technology.

[27]Carter *et al.* (2020) also explore the spillover effects of the subsidy on non-treated farmers along the personal networks of treated farmers. They report that such dynamic effects were not present in the first post-treatment round that we use for estimation here.

[28]The experiment in Carter *et al.* (2020) did not induce changes in the allocation of land across crops that one could use for estimating μ.

right-hand side, $\log \pi_{i,k,t}$, is the land share for crop k (summed over both techniques) used in producing the harvests on the left-hand. The final two terms capture farmer-specific production shocks over time and across crops and farmer i's land endowment, which we capture by including crop-by-year fixed effects ($\theta_{k,t}$), farmer-by-crop fixed effects ($\phi_{i,k}$) and an error term $\epsilon_{i,k,t}$. Alternatively, to allow for region-specific shocks across crops over time, we also replace $\theta_{k,t}$ with region-by-crop-by-year fixed effects ($\theta_{r,k,t}$). The regression coefficient of interest, $\frac{\mu-1}{\mu}$, is thus estimated using changes in land allocations within farmer-by-crop cells controlling for average changes by crop across farmers over time.

To estimate μ convincingly, we require plausibly exogenous variation in land allocations ($\log \pi_{i,k,t}$) across crops over time by farmers that are not confounded with unobserved local productivity shocks. To this end, we make use of the fact that additive trade costs (charged per unit) imply that shocks to world market prices across crops k should lead to a larger reallocation of land shares for farmers closer to the border, as the percentage change in local producer prices is $\frac{\Delta p_{world}}{p_{world,t0}+bordercost_i}$. We use shocks to world prices for coffee, as world coffee prices are both highly relevant (more than 90% of Ugandan coffee production is exported)[29] and likely exogenous to domestic production (Uganda accounts for less than 2% of world coffee sales). We thus construct the instrument as the interaction of the log distance to the nearest border crossing for farmer i, a dummy for whether crop k is coffee, and the log of the relative world price of coffee relative to the average world price of the other eight crops. Note that the fixed effects ϕ_{ik} and θ_{kt} absorb all but the triple interaction term. The identifying assumption is that farmers' productivity shocks in coffee production relative to other crops over time are not related to the the interaction of the timing of coffee's relative world prices with distance to the border.

As documented in appendix Figure A.3, the relative world price of coffee dropped significantly over our sample period 2005-2013. All else equal, land shares used for coffee production should have thus fallen more strongly closer to the border. Panel A of Table 3, which presents the first-stage regression, documents that this is indeed the case: the negative point estimate on our instrument implies that negative relative world price changes for coffee decrease land allocation to coffee more for farmers closer to the border. This relationship holds both before and after including region-by-crop-by-technology-by-time fixed effects, and when using all years of data (2005, 2009, 2010, 2011 and 2013) or just using long changes 2005-2013. In Panel B, we report estimation results before adjusting farmer harvests ($q_{i,k,t}$) by inputs used in production in the denominator of the left-hand side.[30] Panel C presents the second-stage estimation of equation (16). We find statistically significant point estimates in the range of 0.45-0.75. Recall that

[29] Among the 9 main crops we study in Uganda, only coffee falls into this category: the share of exports to production for coffee exceeds 90 percent in all years of our sample, whereas the sum of exports plus imports over domestic production is close to zero (below 4 percent) for the other crops.

[30] Judging from Panel B, it does not seem to be the case that OLS estimates are biased upward compared to IV estimation. If anything, the IV point estimates of harvest on land shares are somewhat larger than in OLS. This could suggest that unobserved idiosyncratic productivity shocks pose less of an omitted variable concern in this setting compared to potentially significant measurement error in the reported land shares allocated to different crops and across different technology regimes on individual farmer plots in the survey data.

this point estimate captures $\beta = \frac{\mu-1}{\mu}$; this therefore implies estimates of μ in the range of 1.8-4. Reassuringly, these are close to existing estimates of this parameter reported in Sotelo (2020) ($\mu = 1.7$). To be conservative, we pick the low estimate of $\mu = 1.8$ as our baseline calibration.[31]

4 Using Microdata for Calibration of a Granular Economic Geography

In addition to estimating key elasticities with local experiments, our approach also takes seriously the value of using microdata to reflect the granular economic geography of an economy when estimating effects at scale. In this section, we show how we populate the vector of observable data used in the price discovery and counterfactual solution we laid out in Section 2:

$$\mathbb{D}_A = \left\{ \{\xi_{i,k}, \xi_{h,k}, p^*_{F,k}, q_{i,k,\omega}, \alpha_{i,n,k,\omega}\}_{k\in\mathcal{K}_A}, L_i, I_h, t_{od,g}, \tau_{od,k} \right\}$$

Appendix 1 provides a summary and additional discussion of the administrative microdata that we use below in the calibration. The rich administrative microdata we use here are increasingly available in many low and middle-income countries.

We restrict the set of crops $\mathcal{K}_A$ to the 9 most commonly grown crops in Uganda: matooke (banana), beans, cassava, coffee, groundnuts, maize, millet, sorghum and sweet potatoes. As documented in Appendix 2, they account for 99 percent of the land allocation for the median farmer and for 86 percent of the aggregate land allocation. Further, we allow for a single intermediate input ($n \in \mathcal{N}_I$) that encompasses chemical fertilizer and hybrid seed varieties.

To estimate the cost shares of intermediates, labor and land in the production function of each crop x technology x location, $\alpha_{i,n,k,\omega}$, we take the median of the cost shares that we observe across households in the UNPS microdata for each of the 4 regions of the country, and appropriately weighted using sampling weights. Appendix Table A.9 presents the cost shares observed in production across the 9 major crops and the two technology regimes (averaged across the 4 regional sets of parameters we use in the calibration).

To calibrate the model to the full set of local markets and households populating Uganda, we need household-level information on pre-existing production quantities ($q_{i,k,\omega}$) and expenditure shares across crops and sectors ($\xi_{i,g}$, $\xi_{h,g}$) for the full population of households we observe in the census microdata, which is generally not available as part of census data.[32] Instead, we use the UNPS, which includes such detailed household-level information for a nationally representative sample of Ugandan households, to project these outcomes on a number of household and location characteristics that are *also* observed in the 100 percent sample microdata from the 2002 population census. Outcomes of interest are total harvest by production technique in each crop, expenditure share on food, expenditures by crop within food and trade costs to the local market (that we estimate among UNPS households as discussed above). For each of these

[31] This is conservative in terms of welfare impacts, and in terms of the difference between local-vs-at-scale effects.

[32] Household labor endowments (L_i) are observed in the census data directly and equal to the number of working-age household members in our calibration. Urban income (I_h) is computed by multiplying UNPS average urban incomes with a city's population. Foreign prices for crops and inputs ($\{p^*_{F,g}\}_g$) are from the FAO database.

outcomes from the UNPS on the left-hand side, we project them (using survey weights in the UNPS) on household and location characteristics observed in both datasets and use the predictions for extrapolation to the 100% census population. These characteristics are (in levels): age and education of the head of the household, number of dependents, number of household members, an asset ownership index (computed using the same assets), potential yield of a given farmer's location from the FAO/GAEZ database, dummies for subsistence farming and urban households, district dummies and survey year fixed effects.[33] For this estimation, we employ Poisson pseudo-maximum likelihood (PPML), which has the nice property of preserving aggregates in the predicted population data.

Trading Frictions

To calibrate trade frictions across local markets, we use survey microdata collected by Bergquist *et al.* (2022) on bilateral trade flows between Ugandan markets, in addition to origin and destination prices. They collect trade flow data in a survey of maize and beans traders located in 260 markets across Uganda (while not nationally representative, these markets are spread throughout the country). Traders are asked to list the markets in which they purchased and sold each crop over the previous 12 months. They complement this data with a panel survey, collected in each of the 260 markets every two weeks for three years (2015-2018), in which prices are measured for maize and beans.

This information can be used to limit the calibration of cross-market trade costs to trading market pairs only within a given period. Consistent with the stylized facts in Appendix 2, we estimate additive trade costs as a function of road distances between markets. Using only bilateral price gaps from market pairs during months in which they observe positive trade flows between the pair (following spatial arbitrage in the model), with information on the road distance between the markets from the transportation network database, we estimate the following specification:

$$t_{od,g,t} = (p_{d,g,t} - p_{o,g,t}) = \alpha + \beta\,(RoadDistance_{od}) + \epsilon_{od,g,t},$$

where t indexes survey rounds and the error term $\epsilon_{od,g,t}$ is clustered at the level of bilateral pairs (od). $RoadDistance_{od}$ is measured in road kilometers traveled along the transportation network. We estimate a single function of trade costs with respect to road distances across all goods, so that $t_{od,g} = t_{od}$.[34] The estimated trade cost for an additional road kilometer traveled between two markets is 1.2 Ugandan shillings (standard error 0.289), which implies a cost of about \$0.5 per kilometer for one ton of shipments. This is consistent with additional survey data from Bergquist *et al.* (2022) documenting that fuel costs for a fully-loaded 5-ton is 0.3 Ugandan shillings per kg per km (standard error 0.024). This would imply that fuel costs account for about 25% of total trade costs, which is consistent with existing findings (e.g., Hummels (2007)). If we replace the

[33] For local trade costs we do not include potential yields.

[34] We do so for power reasons. The dataset covers two crops, maize and beans. Including a crop-month FE in the regression above yields very similar results.

specification above to be in logs on both left and right-hand sides, the distance elasticity is .0258 (standard error 0.0057), which is close to existing recent evidence for within-country African trade flows by e.g. Atkin & Donaldson (2015). We use this distance elasticity to calibrate ad valorem trade costs τ_{od} for trade in the manufacturing good.

To calibrate the local trading frictions between farmers and their local market ($t_{im,g}$), we implement a similar strategy, using gaps between selling farmers' farm-gate prices and local market prices as reported in the UNPS.[35] We first estimate:

$$p_{i,g,t} = p_{m,g,t} - t_{im,g,t} = \theta_{m,g,t} - t_{im,g,t}$$

where $p_{i,g,t}$ is the farm-gate price of good g of farmer i in market (parish) m at year-month t and $p_{m,g,t}$ is the local market price that we do not directly observe and capture with parish-by-crop-by-harvest time fixed effects ($\theta_{m,g,t}$). The farmer-by-crop-by-time specific residual is $-t_{im,g,t}$, the negative of the local trade cost.[36]

The estimated average farmer-level trade friction to their local markets ranges between 23 at the 1st and 90 shilling at the 99th percentile in the population, with an average of about 66 Ugandan shilling per kilogram, which amounts to roughly 8 percent of the average crop price.[37] Finally, we use the UNPS microdata to estimate the trading frictions farmers face when hiring or selling labor in the local market in the same way as for crop trade costs. We replace $p_{i,g,t}$ on the left-hand side above with "farm-gate" wages (paid by farmer i to hired labor, i.e., inclusive of transaction costs).[38] On average, hiring farmers is subject to labor trading frictions of 248 shilling (or 10 US cents) per day for hiring a worker, or around 5% of the daily wage.

5 Counterfactual Analysis: Local vs. At-Scale Policy Impacts

Bringing together the model and solution method from Section 2, the key parameters estimated from local experiments in Section 3, and the calibration to the granular economic geography described in Section 4, we now proceed to quantify local vs. at-scale counterfactuals for one of the most widespread agricultural support policies in low and middle-income countries: a subsidy for modern inputs.[39] In Section 6, we further discuss other agricultural policies to which our approach is immediately relevant, and others to which it could be tailored.

[35] To ensure we are capturing farm-gate prices we restrict the sample to transactions by farmers to private traders. Bergquist *et al.* (2022) document that these transactions occur at the farm-gate.

[36] Since the distribution of trade costs is therefore mechanically centered at zero, after predicting trade costs for the full Ugandan population (see the next step), we shift the distribution rightwards such that a farmer in the bottom 0.1 percentile faces trade costs to the local market that are close to zero (1 Ugandan shilling).

[37] In the upper panel of table A.10 we also report corroborating evidence that the estimated trade costs are significantly related to other measures of remoteness at the farmer level in the UNPS data.

[38] Farmers report hired person-days and expenditure on hired labor, which we use to compute daily wages on the farm.

[39] These policies are widespread: in a survey of 10 African countries, Jayne & Rashid (2013) find that input subsidy programs account for on average 28.6% of total public expenditure on agriculture. They estimate that over 60% of sub-Saharan Africa's population lives in a country with a major input subsidy program.

We proceed with four main sets of results. We first present the analysis of how the welfare impacts of a modern input subsidy differ between a local intervention and one at scale – among the same sample of farmers – and quantify the underlying mechanisms. Second, we use our framework to investigate how the sign and extent of GE forces can differ as a function of saturation rates at different geographical scales, with new implications for randomized saturation designs in the RCT literature. Third, we investigate the role of capturing a realistic, granular economic geography for counterfactual analysis. Fourth, we explore the sensitivity of our findings across alternative parameter values, highlighting the importance of using local experiments for parameter estimation, and present additional model validation results.

Local Effects vs Scaling Up

To fix ideas, we focus on the effects of a subsidy for modern inputs (chemical fertilizers and hybrid seed varieties in the data). We investigate the effects of an intervention that gives a 75 percent cost subsidy for these inputs across all crops.[40] We run two types of counterfactuals in the calibrated model. As depicted in Figure 1, households are located in roughly 4,500 rural parish markets and 70 urban centers. In the local intervention, we randomly select a 2.5 percent sample in each of the rural parishes (roughly 100,000 households nationwide). For each of these markets, we then shock this random sample of households with the subsidy for modern inputs and solve for the counterfactual equilibrium as stated in Section 2. This is akin to running 4500 separate small-scale RCTs. For the intervention at scale, we offer the subsidy to all farming households in the economy (including the original 2.5 percent sample). In both types of counterfactuals, we solve for changes in household-level outcomes across all 4.5 million Ugandan households. We then compare the changes in economic outcomes for the sample of households treated in the original, local-only intervention to their economic outcomes when the intervention is also scaled to the rest of the Ugandan countryside.

Figures 2-4 present the main counterfactual results. In Figure 2 we start by documenting the difference in welfare effects between the at-scale and local interventions across all $\sim 100,000$ national sample households. The left panel shows the at-scale impact minus the local intervention impact, in percentage points, for these households. The right panel aggregates to average effects at the level of parish markets, to facilitate comparison between the average treatment effect that a given parish would experience at scale to the average treatment effect that would be typically measured in a local experiment.[41] The black lines plot the distribution of these differences, with the vertical bar showing the average difference. To shed light on distributional impacts, the blue and red lines show the same effects for the top and bottom quintiles (roughly 20,000 households each) of land shares in initial household income. Those in the bottom quintile – whom we refer

[40] To simplify the exercise, we leave aside for the moment the public finance dimension of the subsidy. It would be straightforward to have this financed by a lump-sum tax in the model.

[41] Changes in welfare are changes in real incomes, with the price index defined as the ideal price index over manufacturing and agricultural consumption given by the nested Stone-Geary preferences stated in the model's parameterization at the end of Section 2.

to as "land-poor" – are smallholder farmers whose land profits from agricultural production are relatively small and who therefore get a larger fraction of their income from labor (including the implicit value used on their own farm, as well as any explicit value they receive from selling their daily labor to other, larger farms). Those in the top quintile – whom we refer to as "land-rich" – are larger landowners who have greater crop income and who tend to be net buyers of labor.[42]

Two main insights emerge. First, the distribution is wide, with households experiencing more than +/- 5 percentage point changes in their welfare impact when the intervention is scaled-up (with the average household experiencing a decrease of about 1 percentage point, or about 20% of the average local welfare effect in appendix Table A.11). Second, scaling up the intervention has very different effects on land-rich vs. land-poor households. We see that the mass of land-rich households lies to the left of zero, suggesting that they tend to lose at scale relative to how they fare under the local intervention, while the mass of land-poor households lies to the right, on average gaining at scale. Table A.11 shows the point estimates of both local and at-scale effects across these different groups.

To further investigate the distributional implications of scaling in this context, Figure 3 presents non-parametric estimates of the local and at-scale welfare effect as a function of initial land income shares. We see in the left panel that while the local intervention strongly benefits land-rich households more than the land-poor (by on average up to 5.5 percentage points), the at-scale intervention significantly flattens this gradient (reducing this gap by more than half, to 2 percentage points). Driving this compression is the fact that land-poor households experience gains that are on average larger at scale than they are under the local intervention, with the poorest households experiencing welfare gains that are 1.5 percentage points larger at scale. In contrast, land-rich households fare worse at scale, with the richest experiencing a 2 percentage point drop in their welfare gains relative to the local intervention. Qualitatively similar differences are present in the right panel when comparing land-rich and -poor households within markets, after conditioning on parish market fixed effects, suggesting that these effects are not driven purely by differences across locations.

Appendix Figures A.5-A.9 and Table A.12 and further investigate the underlying mechanisms driving these differences at scale. In Figure A.5, we decompose the difference between the at-scale effect and the local effect into different underlying components for both the effects on nominal incomes and household price indices. Table A.12 also presents point estimates of the average effects on these various components different groups of households. We find that GE forces on average decrease the positive effect on land income at scale compared to the local intervention for both land-rich and land-poor households, as the price of the local non-traded factor of production (labor) appreciates and (most) crop output prices fall (see Table A.12). Wages and labor income increase on average as a result. Both effects favor (relatively) the initially land-poor, who experience larger increases in their labor earnings and lesser reductions in their land

[42]Appendix Figure A.4 also presents flexibly estimated (positive) relationships between our measure of land shares and households' land ownership in acres or households' calibrated total incomes. Fink *et al.* (2020) document similar patterns in another African context (Zambia).

earnings.[43] Price index effects are more muted at scale vs. locally, because reductions in the relative price of food (favoring poorer households) are offset by relative price changes within food that tend to favor richer households.

Figure 4 provides evidence on the role of remoteness. Theory suggests that the gap between the effects of a local vs. at-scale intervention increase in market remoteness, as GE effects on crop output prices are strongest in more remote markets, where local prices are less pinned down by world prices at border crossings or proximity to cities. The left panel of 4 confirms this hypothesis for access to other markets within Uganda (measured by the log of the inverse distance-weighted sum of population in all other markets and cities d in Uganda for each origin parish o on the x-axis: $\sum_{d \neq o} \frac{Pop_d}{Distance_{od}}$). The right panel plots the same relationship with respect to the log distance to the nearest border crossing in km on the x-axis. Both panels present the difference in households' welfare impact (at scale - local) in percentage points on the y-axis. We find that deviations between local and at-scale effects tend to be more pronounced in relatively remote rural market places.

Appendix Figures A.7-A.9 provide additional evidence on the roles played by initial technology usage and crop planting decisions in shaping effects at-scale vs. the local intervention. We document that land-rich households benefit more from the local intervention in part because of significantly higher pre-existing usage of modern technology (higher cost shares for fertilizer and hybrid seeds). Average crop prices fall most at-scale among crops with higher pre-existing usage of modern technology, and farmers planting these crops gain more in the local intervention (and relatively lose more at scale).

GE Forces as a Function of the Intervention's Scale

Experimental approaches to capturing GE effects often employ "randomized saturation" designs, in which the fraction of individuals treated is randomized across geographic areas or "clusters" in order to study the market-level outcomes that emerge (see e.g. Baird *et al.* (2011); Burke *et al.* (2019); Egger *et al.* (2022)). Here we use our approach to investigate how the GE effects in our context evolve as the intervention is scaled up to an increasingly large fraction of households nationwide, and as the geographic scale of the cluster is varied. Both have implications for the optimal design and lessons that can be learned from randomized saturation designs.

Panel A of Figure 5 presents the welfare impact of the subsidy on the original national farmer sample as a function of the nationwide fraction of the rural population that is also treated. The left-most point on the x-axis corresponds to the local intervention, where only parish-level samples of 2.5% of the local population are treated. The right-most point on the x-axis corresponds to the at-scale intervention above where 100% of rural Ugandan households receive the subsidy

[43] This is driven both by higher pre-existing labor income shares among the land-poor as well as slight differences in average wage and crop price effects due to differences in crop and technology usage across households and markets. Appendix Figure A.6 shows the same graph without the initial income share weighting (no longer summing up to the total income effect), documenting about 1 percentage point more positive wage effects at scale (compared to the local effect) among the land-poor, but also about 1 percentage point more negative land earning effects at scale.

treatment. The point estimates going from left to right plot the average treatment effect on the same initial 2.5% household sample across increases in the national saturation rate in steps of 10 percentage points of the rural population.[44]

The left figure in Panel A traces the average welfare impact, while the right figure displays the average effect separately for the bottom and top quintiles of the initial land income shares. The main insight that emerges is that the extent of GE forces appears to be a monotonic and roughly linear function of the national saturation rate, both for the average effect in the left figure and the distributional implications of the policy on the right in Panel A. These findings are reassuring, as they would in principle support comparisons between just two discrete levels of saturation, as has become common practice in randomized saturation designs.

That said, Panel A varies the saturation at the national level. In practice, randomized saturation designs typically randomize the saturation within some smaller geographic unit ("cluster"). Panel B of Figure 5 explores the role played by the size of these clusters. To illustrate, we consider the case of a study design that uses subcounties (of which there are 811 in Ugdanda during our study period) as the unit at which saturation is randomized. These are relatively large geographical units compared to the typical "clusters" in the literature as we discuss below. For example, Egger *et al.* (2022) randomize treatment saturation at the level of sublocations in Kenya (groups of 10-15 villages), which are smaller than Uganda's subcounties.

Consider, specifically, a design that randomly selects 51 subcounties in which to implement this design (each randomly picked within one of the 51 districts of Uganda). First, just to demonstrate that these 51 subcounties are not distinct in some important way, we replicate the exercise from Panel A (increasing saturation rates nationwide in increments of 10 percentage points) and plot results for this random subset of subcounties (including roughly 6500 households of the same national 2.5% sample as in Panel A above); the blue line in Panel B shows results that closely mirror those in Panel A. Next, we consider the more feasible randomized saturation design in which – rather than varying the saturation rate at the national level – the saturation rate is varied at the *subcounty level*, with the rate of saturation goes from 0% to 100% *just within the 51 study subcounties* as we move from left to right along the x-axis. Results are presented in orange in Panel B.

Two main insights emerge from this exercise. First, in contrast to changes in national saturation rates, for which we see the impact of the program decreasing monotonically with scale, we find almost no changes in the average impact of the program as a function of subcounty-level saturation rates, even at 100% saturation within these areas (see left side of Panel B). This means that a design that randomizes the saturation at the subcounty level, even with extreme differences in saturation rates, would not be able to measure GE-driven changes in the average impact across these rates. One might then incorrectly conclude there is no change to the program's average impact from scaling up. Second, one would also draw the wrong distributional

[44]We solve for counterfactual outcomes after randomly selecting additional fractions of households within all parishes in increments of 10% until reaching full saturation. The first 10% national saturation treats an additional 7.5% of the local population in all parishes.

implications from a randomized saturation design at the subcounty level. While at the national level, declines in the average welfare impact are predominantly driven by a reduction in welfare gains among the top quintile of land-rich households, a design that randomized saturation at the sublocation level would find weaker reductions among the land-rich and stronger increases in gains among the land-poor as a function of local saturation rates – offsetting one another so that the average effect across farmers is close to constant. The forces behind these trends are that farmers' crop prices react differently to saturation rates at more or less local geographical scales: increasing nationwide saturation rates has significant implications on output prices (see Table A.12), whereas changes in the saturation within sub-county populations have much more muted implications on output prices. As a result, local increases in saturation mainly imply that parts of the land revenue gains are capitalized into the local non-traded factor of production (labor) – explaining why averages are close to unaffected, while land-poor farmers gain more (and land-rich farmers lose less) as a function of local saturation compared to nationwide saturation.

These results suggest some caution in extrapolating from the reduced-form results observed in a randomized saturation design what welfare impacts would look like under a nationwide program. Even when randomizing saturation at the subcounty level – which in Uganda encompasses on average 32 villages and 30,000 individuals, and therefore is larger than most units used in the existing randomized saturation literature[45] – this may still be too "local" in scale, and therefore unable to generate the type of GE forces that would emerge under a nationwide roll-out. This by no means implies that these designs are not useful for informing predictions of impacts at scale, but rather that the variation they generate may need to be combined with approaches such as the one described here in order to make predictions for impacts at national scale, a point we turn to in Section 6.

The Role of a Granular Economic Geography

In Section 2, we emphasized several features of the economic geography of rural agricultural markets that are typically absent from existing quantitative models, but which may matter for the propagation of shocks across markets and sectors. These include: (i) a granular economic geography with trade costs between household locations within markets and transportation routes across markets; (ii) homogeneous goods, allowing for extensive margin impacts across trading pairs; and (iii) additive trade costs, allowing for incomplete and heterogeneous pass-through of price shocks. Our approach captures these and a number of additional salient features. How much do these innovations matter, quantitatively, for the implied effects at-scale?

To this end, we compare the effect of the at-scale intervention across models with alternative geographies. In the first alternative model, we follow the tradition in CGE analysis and most of macroeconomics, and estimate GE counterfactuals in a single integrated national market – assuming no trade costs for output or inputs within Uganda. In the second alternative model, we instead follow the literature in international trade and assume the Ugandan economy is sub-

[45]See e.g. Baird *et al.* (2011), Burke *et al.* (2019), Egger *et al.* (2022).

ject to iceberg (ad-valorem) trade costs and structural gravity in a standard Armington model at the level of parish markets trading crops.[46] Except for changing assumptions on the nature of trade frictions and product differentiation in agriculture, we keep the rest of the model and its calibration as in our baseline.[47]

Figure 6 shows the comparison to a single integrated market in the left panel, and the comparison to the Armington model in the right panel. In both graphs, the y-axis displays percentage point differences in the welfare impact of the at-scale intervention between the baseline model minus the effect in the alternative model across the $\sim 100,000$ households as a function of initial land income shares on the x-axis as before. The dashed red lines indicate the sample average of these differences. On average, the single integrated market would overestimate the welfare gains at scale by 15%. In terms of distributional implications, the single-market economy would miss the reversal of the policy's regressivity at scale revealed by our model: land-poor households on the left of the x-axis experience higher gains at scale under realistic, granular economic geography compared to a world without trading frictions, whereas land-rich households on the right experience significantly smaller gains. Comparing this to the left panel in Figure 3, the single market would capture less than half the GE adjustment on the distributional implications at scale compared to the local effect. This is because crop price adjustments are muted in a single national market place, as world market prices at the border are more binding across markets. This decreases the asymmetry between the local intervention (at unchanged initial output prices) and the intervention at scale – benefiting land-rich households at scale whose output prices decrease less compared to a world with a granular economic geography.

The comparison to the Armington model in the right panel of Figure 6 documents that both the average and distributional welfare implications differ meaningfully when assuming ad-valorem iceberg trade costs and structural gravity with product differentiation, we typically do for manufacturing varieties. We find that average welfare gains under our preferred model with homogeneous goods and additive trade costs are about 50% greater than implied under the standard Armington model, which underestimates gains to land-rich households in particular. The weaker average effects in the Armington model are due to the implied lower elasticity of substitution (i.e., finite) between the varieties of a given crop produced in different parishes. The weaker response of the demand for crops leads to a bigger drop in prices but smaller effects on wages.

These results indicate that embracing a granular and realistic economic geography matters for counterfactual analysis at scale, both for average effects and distributional implications.

[46]We treat each parish as a single integrated markets and assume that each crop is differentiated across parishes but that farmers within a parish produce homogeneous crops. Following the literature, we use a trade elasticity of 5 (i.e., the elasticity of substitution in consumption across varieties of each crop across different parishes).

[47]This Armington specification is another special case of our framework where each location produces a different good, akin to how we model the manufacturing sector. In this specification, we use our estimated iceberg trade costs to calibrate trade shares in the baseline equilibrium, and we can use the exact hat algebra to describe the counterfactual equilibrium.

Importance of Local Experiments in Identifying Key Elasticities and Model Validation

In Section 4, we emphasized the role that local experiments could play in rigorous identification of key elasticities. The more sensitive the counterfactuals are to these elasticities, the more critical clean identification becomes, as biased estimates generated from observational variation could substantially distort the implied policy impacts. In this section, we explore how alternative parameter estimates would alter the implications of effects at-scale. This exercise offers both greater intuition about how key elasticities drive impacts at scale, and guidance on which parameters are most critical to identify accurately with credible approaches. Beyond parameter sensitivity, we also present additional model validation results.

Figure 7 presents the counterfactual results for the intervention at scale under alternative parameter assumptions on the supply side (κ and μ) and the demand side (σ). In the upper left panel, we see that the magnitude of the lower-tier supply elasticity, κ, is quite important for our estimates. Higher values of κ increase the estimated welfare effects at-scale, as farmers are more responsive to price changes in how they allocate their land across technology choices within a given crop. This may help explain why some RCTs have found larger effects over the long-run, as greater time for adjustment may imply larger elasticities (Bouguen *et al.* , 2019).

Higher values of κ also lead to larger differences between the local and at-scale intervention in GE, as greater responsiveness on the part of others leads to larger output and factor price changes at scale compared to local intervention (at original prices). This highlights the importance of careful identification of this parameter. Using exogenous variation in prices coming from experiments, as we do here with an experimental fertilizer subsidy (Carter *et al.* , 2020), can increase our confidence in our estimate of this key parameter for a given policy context. This therefore represents an important role that can be played by experiments, a point we return to in Section 6.

Conversely, our estimates are less sensitive to the upper-tier supply elasticity μ (across crops) or the value of the demand elasticity σ (upper right and lower panels). In our setting, cost shares of modern inputs do not differ substantially across crops, and while we find above that crops in GE are affected differently by the subsidy policy, cost share differences remain relatively minor (compared to shifting across production regimes within crops). How households trade off these crops in consumption is therefore also less critical for the changes in the policy's impact locally vs at scale. However, in other contexts (e.g. with more strongly differing input suitability in production across crops, or with an intervention targeted at one particular crop), both μ and σ could play more important roles in shaping the effects at scale and their difference relative to the local effects.

Beyond parameter sensitivity, we present additional model validation results. One important innovation of our theory is to use the model-based price discovery algorithm to solve the model with the new economic features we allow for in this setting. This involves solving for farm-gate prices (at the level of household locations) and trade flows that rationalize the observed consumption and production decisions given a graph of trade costs. For model validation, we are able use data on crop prices and trade flows between 260 Ugandan markets in the trader surveys collected by Bergquist *et al.* (2022). Comparing these market places in our baseline model and in the data, we can assess to what extent the model-based estimates of local crop prices and predicted trading relationships between markets capture variation in prices and trade flows of those same markets in the survey data.

Panel A of Figure 8 compares the variation in local market prices for maize and beans across the Ugandan markets in data vs. model. For each of 38 months of the trader survey data, we take the median market price for each crop and market in a given month. The y-axis of the binned scatter plot shows the residuals from a regression of the log median market prices in the trader surveys on month-by-crop fixed effects. The x-axis displays mean deviations of log prices for the same two crops across the same markets in the baseline equilibrium – the results from the price discovery algorithm. Reassuringly, the model-based price variation – based entirely on observed information on crop production, consumption and trade costs on a connected graph of household locations in Uganda – presents a rather tight, positive and roughly linear relationship to observed price variation in the same crops and markets pooled over the 38 months of survey data.[48]

But part of the price variation across markets in the trader survey data was used in our calibration of trade costs – in particular price gaps between trading pairs. To ensure that the model's relationship to the survey data is not partly mechanical in that respect, Panel B converts the data to bilateral origin-destination price gaps (with each bilateral pair counted only once for a given crop and month of data). We then exclude all pairs with positive trade flows (which were used to quantify trade costs in the model calibration). The remaining bilateral price gaps in the data are then plotted against the same market-to-market price gaps from our price discovery algorithm. Panel B confirms a roughly linear and rather tight positive relationship between price variation in the model to price variation in the survey data, even when excluding any moments used in the calibration of trade costs in the model.

Finally, Panel C of Figure 8 compares the observed active trading routes in the data to the ones predicted by our model's price discovery algorithm. Of the 1256 bilateral trade flows for maize observed in the data (stacked across 12 months), the model captures 968 active trading relationships (77%). For beans, the model predicts 75% of the observed bilateral trading relationships (392/522). The reverse proportions – the fraction of crop-by-market pair relationships predicted in the model that are captured in the trader surveys for the same markets and crops – are somewhat lower (71% for maize and 37% for beans). One explanation for this is that the

[48]There are, of course, many reasons why the price deviations can differ in data vs. model. On the survey data side, there could be measurement error, unobserved variation in crop quality, as well as temporary idiosyncratic shocks on the day that information was collected across different market places in Uganda. On the model side, household locations, expenditure shares and crop production moments are partly extrapolated to the population with likely significant degrees of measurement error. Parish markets in the model are based on centroids, whereas real-world market places that are assigned to the same parish identifier do not necessarily coincide geographically. All of these factors would imply a somewhat noisy and attenuated relationship between model fundamentals and real world data.

trader surveys are based on a sample of traders, whereas the model captures aggregate trade flows between markets. We view the high proportion of observed bilateral trading relationships for a given crop that the model correctly predicts as another piece of reassuring evidence that the model-based price discovery algorithm reveals meaningful economic variation across markets.

6 Discussion

This paper develops a toolkit that can be combined with field and quasi-experiments to investigate GE treatment effects at scale. We see these two approaches as complementary and hope that, in combination, one can expand what can be learned from (quasi-)experiments or quantitative GE models alone. In the following discussion, we explore some concrete ways in which we view these toolkits as complementary. We also discuss some practical considerations for combining the two approaches, including how to broaden our approach to a wider range of agricultural interventions and concrete tips for data collection and implementation of our approach. As a complement to our paper, we are also creating a streamlined coding toolkit that can be combined with data for calibration and experimental estimates of local interventions to implement our methodology in different empirical and policy settings.

6.1 Complementary Tools

What do approaches such as ours bring to experiments? Muralidharan & Niehaus (2017) discuss three ways in which the impact of policies implemented at scale can differ from those measured in small-scale RCTs: (1) GE and spillover effects: factor and output prices or other market-level features may shift in ways that alter treatment effects and their distribution; (2) external validity: treatment heterogeneity may mean that results measured among the study sample differ from those that would be experienced by the broader population; and (3) implementation differences: program logistics may be different at scale, as implementation moves from a researcher-run or pilot program to a large-scale operation run by governments or other big organizations.

Our approach provides a new toolkit to investigate and quantify the first two issues. On GE effects, the quantitative model developed here is explicitly targeted at analyzing how input and output prices adjust – and the resulting ripple effects on factor usage, production, consumption, and ultimately household welfare – when policies are implemented at scale. By simulating effects in the whole population or among areas not in the study sample, this toolkit also speaks to external validity, to the extent that treatment effects and GE forces vary based on dimensions that are modeled in our framework (such as heterogeneity in revenue or consumption impacts driven by variation in initial crop allocations, technology and factor usage in production, expenditure shares in consumption, or local trade costs and linkages to other markets). Our approach does not have much to say about the third issue of implementation differences, other than noting that estimates will be more accurate the closer the experiment's implementation is to the final at-scale policy.[49]

[49]In principle, one could investigate counterfactuals with alternative assumptions on how implementation at scale

Finally, in addition to helping us to learn more from experiments ex-post, our toolkit can also provide guidance for experiments "ex-ante" to inform the experimental design and data collection (including questions of stratification and power calculations), as we discuss below.

Conversely, what do smaller-scale experiments bring to quantitative GE models like the one developed in this paper? We see three important roles. The first, which we demonstrate here, is to use exogenous variation from RCTs or quasi-experiments to more credibly identify some of the key parameters both on the supply and demand sides of the model. As documented in the previous section, these elasticities matter for the extent and incidence of GE forces at scale. Long-run RCTs are particularly useful here, as they give time for adjustment and are therefore more likely to capture long-run elasticities.

A second benefit from RCTs is that the fieldwork and data collection can provide key moments for the model calibration that are frequently outside the scope of available administrative or other microdata. For example, in our analysis above we brought to bear knowledge of bilateral market-to-market trade flows for trade cost estimation.

A third role for RCTs is model validation. Randomized saturation designs, like the ones explored in the previous section, can be particularly useful here as they can provide empirical counterparts to model-predicted GE forces. Although we show that randomized saturation designs do not necessarily, in reduced form, yield GE impacts at a broader scale of program roll-out (i.e. beyond the level of clusters as defined in the RCT), they can still be very useful for estimating "sublocation GE effects" – changes in crop and factor prices and other market-level features driven by local differences in saturation – that can be compared to model-based counterfactuals based on the same geographical clusters to validate the model. Such validation can then lend credibility to predicted effects at a larger geographical scale, at which saturation randomization may not be feasible.

6.2 Combining the Toolkits in Practice

Assuming one wants to combine an experiment with an approach like ours, how does one go about it in practice?

Applying the Approach to Different Agricultural Policies

Our approach is well-suited to be off-the-shelf applicable to three common types of agricultural interventions. The first is shocks to agricultural productivity, which enter through the $\hat{b}_{j,k,\omega}$ in the model. These can include weather and climate change-driven shocks, the introduction of new seed varieties and other technologies, and complementary inputs such as irrigation that alter agricultural productivity. As demonstrated above, this can also include input subsidies, one of the largest and most common agricultural policy interventions in developing countries. Second, our approach is readily applicable to demand-side shocks, such as policies that affect non-agricultural income, incomes in different regions (e.g. urban households) or those that

may change the direct incidence or take-up of the subsidy, and quantify implications at scale based on those assessments. In practice more research may be needed in this space to learn about such differences (Duflo (2017)).

alter consumption preferences, such as nutritional awareness campaigns (entering the model through $\hat{a}_{j,k}$). Finally, our approach is relevant to policies that affect trade costs ($\hat{\tau}_{od,k}$ or $\hat{t}_{od,k}$ in our model). Most obviously, this includes road building initiatives, railway network extensions, and other policies that reduce transportation costs. But other policies aimed at reducing trading frictions – for example those targeted at reducing search frictions or improving competition of the transportation sector – could be considered subject to a well-defined mapping between these policies and reductions in trading frictions.

Of course, there is also a range of agricultural policies for which our current model would not be off-the-shelf appropriate and for which it would need to be substantially tailored or extended to speak to effects at-scale. One important example is land market reforms such as those that title land (D. Ali & Ferrara, 2015), privatize land (Manysheva, 2022), legalize land sales (Chaoran Chen, 2022), or put caps on farm sizes (Adamopoulous & Restuccia, 2020). Our current model would also not be readily applicable to policies that aim to reduce risk (Donovan, 2021); (Emerick *et al.* , 2016) or alleviate the impact of inter-temporal shifts in preferences (Duflo *et al.* , 2011) or prices (Basu & Wong, 2015); (Burke *et al.* , 2019). We consider our setup and solution method as a first and important step to unlock quantitative analysis paired with rich and granular microdata in this important policy setting, and see these and other extensions as promising avenues for future research in this context.

Data collection and research design considerations

In the following, we offer some practical considerations for both data collection and research design. In terms of data collection, researchers will want to collect data on production and consumption of all major crops, not just those directly targeted by the intervention, as in GE multiple output and factor markets can be affected. These data are crucial for estimating both supply- and demand-side elasticities, as well as for calibrating cost shares or technology use in production functions across crops. Given that wage effects can play an important role, capturing input expenditures on labor (including own labor) is important, albeit often difficult to measure. For the model calibration at scale, collecting similar covariates to those included in nationwide administrative datasets (ideally captured using similarly-worded survey questions) can support the extrapolation step of the model calibration in cases in which not all household outcomes in the initial equilibrium are observed in national census data. Finally, collecting data on market prices and trade flows is useful for calibrating trade costs between markets as well as between households and markets. A large literature in international trade and economic geography has documented that (easier-to-observe) freight rates typically only account for a fraction of overall trading frictions across space (e.g., Allen (2014)). As we lay out in Section 4, knowledge of where trade flows occur, their direction and the market prices at both origin and destination can be used to estimate total trade costs in a theory-consistent way.

Our toolkit also offers guidance in terms of the research design. When randomized saturation designs are planned, researchers can use estimates of parameter values (drawn from our study or others in the literature), to calibrate the model ahead of time in an exercise mimicking a power calculation. Such model-based simulations could inform decisions about, for example, the level at which to randomize saturation, the degree of cross-cluster spillovers or the degree of saturation needed to detect treatment effects on GE outcomes. A calibrated version of our model can also be used for stratification to make the estimated treatment effects representative of the overall population. In particular, our model embraces a number of sources for heterogeneous treatment effects that are not generally included among the standard demographic characteristics used for stratification – such as measures of a market's trading costs for farmers within the market region or to other destinations (market access/remoteness), differences in regional production functions or household expenditure shares for the same crops. In particular, rather than merely stratifying on a number of factors, our model would allow researchers to stratify on predicted treatment effects (both locally and at scale). Finally, clarity about the parameters to be used in model estimation ex-ante may point researchers to additional experimental variation that can be used for estimation.[50]

7 Conclusion

We propose a new approach to quantify large-scale GE counterfactuals in the context of agricultural policies that can both complement evidence from field and quasi-experiments and be informed by it. We develop a rich but tractable quantitative GE model of farm production, consumption and trading. To capture a number of salient features that we document in this setting, the model departs from the workhorse "gravity" structure in international trade and economic geography in several dimensions. We then propose a new solution method that allows us to study GE counterfactuals in this rich environment, without imposing additional data requirements that would be practically infeasible. To showcase our approach, we bring to bear administrative microdata on household locations, production, consumption and the transportation network within and across local markets to calibrate the model to the roughly 4.5 million households populating Uganda in 2002. We use a combination of existing RCTs and variation from natural experiments to estimate the model's key parameters.

We find that the average effect of a subsidy for chemical fertilizers and hybrid seed varieties on rural household real incomes can differ substantially when implemented at scale compared to results from a local intervention that leaves output and factor prices largely unaffected. We show that this difference extends to the policy's distributional implications, which are regressive according to results from the local intervention, but much less so when implemented at scale. We also use our framework to document new insights about the sign and extent of GE impacts as a function of saturation rates at different geographical scales. We find that while GE forces appear to be a monotonic and approximately linear function of saturation rates within a given

[50]For example, even with randomized saturation designs that generate variation in agricultural prices, one may not be able to use this variation to estimate demand for these goods, as many consumers of these products are also producers and therefore price changes can generate changes in income. Separate experiments to identify demand-side elasticities may be needed, such as e.g. the randomized price experiment used in Bergquist & Dinerstein (2020).

geographical area, both their average size and distributional impact depend on the geographical scale at which saturation is being implemented.

The framework we lay out in this paper is aimed at providing a toolkit that can be used to complement the empirical findings from experiments and quasi-experiments related to agriculture. While we hope to break new ground in this context, this paper by no means exhausts the interesting dialogue between reduced-form evidence and model-based counterfactuals. For example, from theory to field work, that dialogue could be used to inform the design of future RCTs to include data collection targeted at estimating key supply and demand elasticities in a given context. From fieldwork to theory, on the other hand, that dialogue could yield additional results on model validation, with a focus not just on the local effects in a given market place, but also using experimental estimates of GE forces from randomized saturation designs. These and related questions provide an exciting agenda for future research in this area.

References

Acampora, Michelle, Casaburi, Lorenzo, & Willis, Jack. 2022. Land Rental Markets: Experimental Evidence from Kenya. *Working Paper*.

Adamopoulous, Tasso, & Restuccia, Diego. 2020. Land reform and productivity: a quantiative analysis with micro data. *American Economic Journal: Macroeconomics*.

Allen, Treb. 2014. Information frictions in trade. *Econometrica*, **82**(6), 2041–2083.

Atkin, David, & Donaldson, Dave. 2015. Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs. *NBER Working Paper*.

Baird, S., McIntosh, C., & Ozler, B. 2011. Cash or Condition? Evidence from a Cash Transfer Experiment. *Quarterly Journal of Economics*, **126**, 1709–1753.

Baqaee, David, & Farhi, Emmanuel. 2019. Networks, barriers, and trade.

Basu, Karna, & Wong, Maisy. 2015. Evaluating seasonal food storage and credit programs in east Indonesia. *Journal of Development Economics*, **115**(200-216).

Bergquist, Lauren Falcao, & Dinerstein, Michael. 2020. Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya. *American Economic Review*, **110**(12), 3705–47.

Bergquist, Lauren Falcao, McIntosh, Craig, & Startz, Meredith. 2022. Search Costs, Intermediation, and Trade: Experimental Evidence from Ugandan Agricultural Markets. *Working Paper*.

Berry, Steven, Gandhi, Amit, & Haile, Philip. 2013. Connected substitutes and invertibility of demand. *Econometrica*, **81**(5), 2087–2111.

Bouguen, Adrien, Yue Huang, Michael, Kremer, Michael, & Miguel, Edward. 2019. Using RCTs to Estimate Long-Run Impacts in Development Economics. *Annual Review of Economics*, **11**, 523–561.

Brooks, Wyatt, & Donovan, Kevin. 2020. Eliminating uncertainty in market access: The impact of new bridges in rural Nicaragua. *Econometrica*, **88**(5), 1965–1997.

Buera, Francisco, Kaboski, Joseph, & Shin, Yongseok. 2017. The Macroeconomics of Microfinance. *Notre Dame Working Paper*.

Burfisher, Mary E. 2021. *Introduction to computable general equilibrium models*. Cambridge University Press.

Burke, Marshall, Bergquist, Lauren Falcao, & Miguel, Edward. 2019. Sell Low and Buy High: Arbitrage and Local Price Effects in Kenyan Markets. *Quarterly Journal of Economics*.

Caldwell, Richard, Lambert, Rachel, Magruder, Jeremy, McIntosh, Craig, & Suri, Tavneet. 2019. Improving agricultural extension and information services in the developing world. *VoxDev Article*.

Carter, Michael, Laajaj, Rachid, & Yang, Dean. 2020. Subsidies and the African Green Revolution: Direct Effects and Social Network Spillovers of Randomized Input Subsidies in Mozambique. *American Economic Journal: Applied Economics*.

Castaneda, Andres, Doan, Dung, Newhouse, David, Nguyen, Ming Cong, Uematsu, Hiroki, & Pedro Azevedo, Joao. 2016. Who are the poor in the developing world? *World Bank Policy Research Working Paper*.

Chaoran Chen, Diego Restuccia, Raul Santaeulalia-Llopis. 2022. The effects of land markets on resource allocation and agricultural productivity. *Review of Economic Dynamics*, **45**, 41–54.

Costinot, Arnaud, & Donaldson, Dave. 2016. How Large are the Gains from Economic Integration? Theory and Evidence from US Agriculture 1880-1997. *NBER Working Paper*.

Costinot, Arnaud, & Rodríguez-Clare, Andrés. 2014. Trade theory with numbers: Quantifying the consequences of globalization. **4**, 197–261.

D. Ali, K. Deininger, M. Goldstein, & Ferrara, E. La. 2015. Impact of the land tenure regularization program in Rwanda. *World Bank Development Research Group Case Study*.

de Janvry, Alain, & Sadoulet, Elisabeth. 1995. *Quantiative Development Policy Analysis*. The John Hopkins University Press.

De Melo, Jaime. 1988. Computable general equilibrium models for trade policy analysis in developing countries: A survey. *Journal of Policy Modeling*, **10**(4), 469–503.

Deininger, Klaus, Ayalew Ali, Daniel, & Alemu, Tekie. 2008. Assessing the functioning of land rental markets in Ethiopia. *Economic Development and Cultural Change*, **57**(1), 67–100.

Dekle, Robert, Eaton, Jonathan, & Kortum, Samuel. 2007. Unbalanced Trade. *The American Economic Review*, 351–355.

Donaldson, Dave, & Hornbeck, Richard. 2016. Railroads and American economic growth: A "market access" approach. *The Quarterly Journal of Economics*, **131**(2), 799–858.

Donovan, Kevin. 2021. The Equilibrium Impact of Agricultural Risk on Intermediate Inputs and Aggregate Productivity. *Review of Economic Studies*.

Duflo, Esther. 2017. Richard T. Ely Lecture: The Economist as Plumber. *American Economic Review*, **107**(5), 1–26.

Duflo, Esther, Kremer, Michael, & Robinson, Jonathan. 2011. Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya. *American Economic Review*, **101**(6), 2350–90.

Egger, Dennis, Haushofer, Johannes, Miguel, Edward, Neihaus, Paul, & Walker, Michael. 2022. General Equilibrium Effects of Cash Transfers: Experimental Evidence from Kenya. *Econometrica*.

Emerick, Kyle. 2018. Agricultural productivity and the sectoral reallocation of labor in rural India. *Journal of Development Economics*.

Emerick, Kyle, & Burke, Marshall. 2016. Adaptation to climate change: Evidence from US Agriculture. *American Economic Journal: Economic Policy*.

Emerick, Kyle, de Janvry, Alain, Sadoulet, Elisabeth, & Dar, Manzoor H. 2016. Technological Innovations, Downside Risk, and the Modernization of Agriculture. *American Economic Review*, **106**(6), 1537–61.

Faber, Benjamin, & Gaubert, Cecile. 2019. Tourism and economic development: Evidence from Mexico's coastline. *American Economic Review*, **109**(6), 2245–93.

Farrokhi, Farid, & Pellegrina, Heitor S. 2022. Trade, technology, and agricultural productivity. *Forthcoming, Journal of Political Economy*.

Fink, Gunther, Jack, B Kelsey, & Masiye, Felix. 2020. Seasonal liquidity, rural labor markets, and agricultural production. *American Economic Review*, **110**(11), 3351–92.

Fujimoto, Junichi, Lagakos, David, & Vanvuren, Mitchell. 2019. The Aggregate Effects of Free Secondary Schooling in the Developing World. *Working Paper*.

Gollin, Doug, Hansen, Casper Worm, & Wingender, Asger Moll. 2021. Two Blades of Grass: The Impact of the Green Revolution. *Journal of Political Economy*.

Head, Keith, & Mayer, Thierry. 2014. Gravity equations: Workhorse, toolkit, and cookbook. *Pages 131–195 of: Handbook of international economics*, vol. 4. Elsevier.

Holden, Stein T, Otsuka, Keijiro, & Place, Frank M. 2010. The emergence of land markets in Africa: Impacts on poverty, equity, and efficiency.

Hummels, David. 2007. Transportation costs and international trade in the second era of globalization. *Journal of Economic perspectives*, **21**(3), 131–154.

Huntington, Heather, & Shenoy, Ajay. 2021. Does Insecure Land Tenure Deter Investment? Evidence from a Randomized Controlled Trial. *Journal of Development Economics*.

Jayne, T.S., & Rashid, Shahidur. 2013. Input subsidy programs in sub-Saharan Africa: a synthesis of recent evidence. *Agricultural Economics*, **44**, 547–562.

Lagakos, David, Mobarak, Ahmed Mushfiq, & Waugh, Michael E. 2021. The welfare effects of encouraging rural-urban migration. *Econometrica*.

Manysheva, Kristina. 2022. Land property rights, financial frictions, and resource allocation in developing countries. *Working Paper*.

Muralidharan, Karthik, & Niehaus, Paul. 2017. Experimentation at Scale. *Journal of Economic Perspectives*, **31**(4), 103–24.

Porzio, Tommaso, Rossi, Federico, & Santangelo, Gabriella. 2022. The human side of structural transformation. *American Economic Review*, **112**(8), 2774–2814.

Powell, Alan A, & Gruen, FHG. 1968. The constant elasticity of transformation production frontier and linear supply system. *International economic review*, **9**(3), 315–328.

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

Sraer, David, & Thesmar, David. 2018. A Sufficient Statistics Approach for Aggregating Firm-Level Experiments. *NBER Working Paper*.

8 Figures and Tables

Figures

Figure 1: Ugandan Markets and Transportation Network

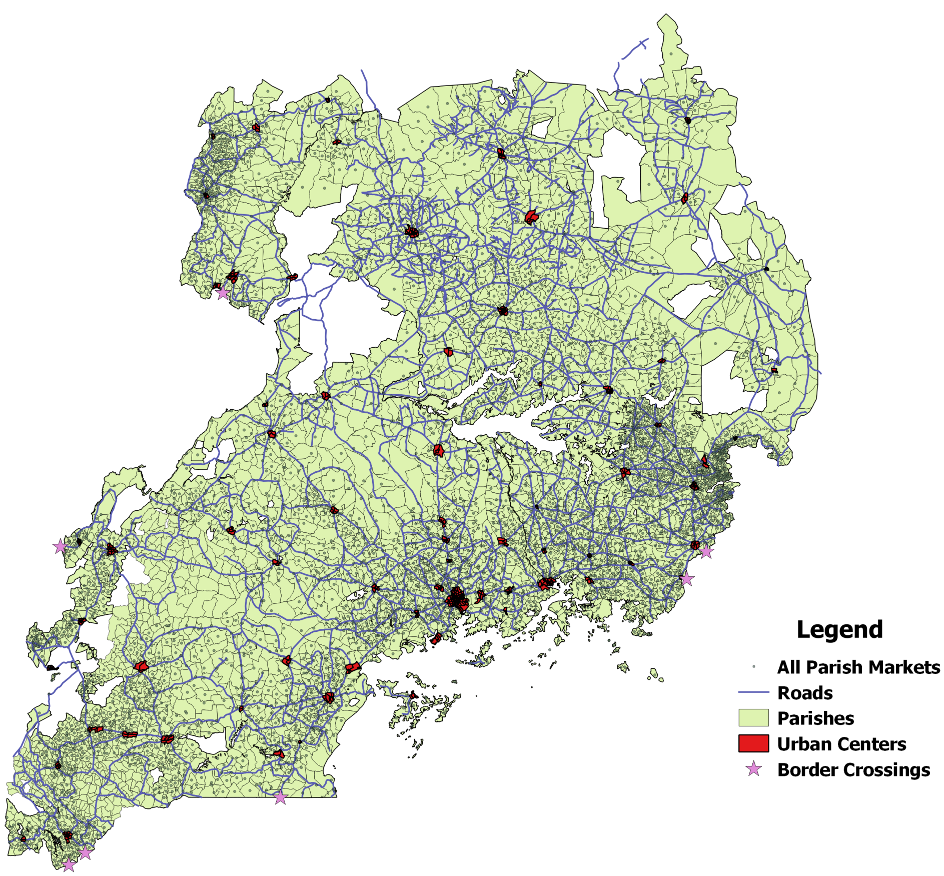

The figure displays the location of local parish markets, urban markets, border crossings and the road network in Uganda. See Section 3 for discussion of the data and Section 5 for the counterfactual analysis based this geography.

Figure 2: Difference in the Effect at Scale vs. Local Interventions

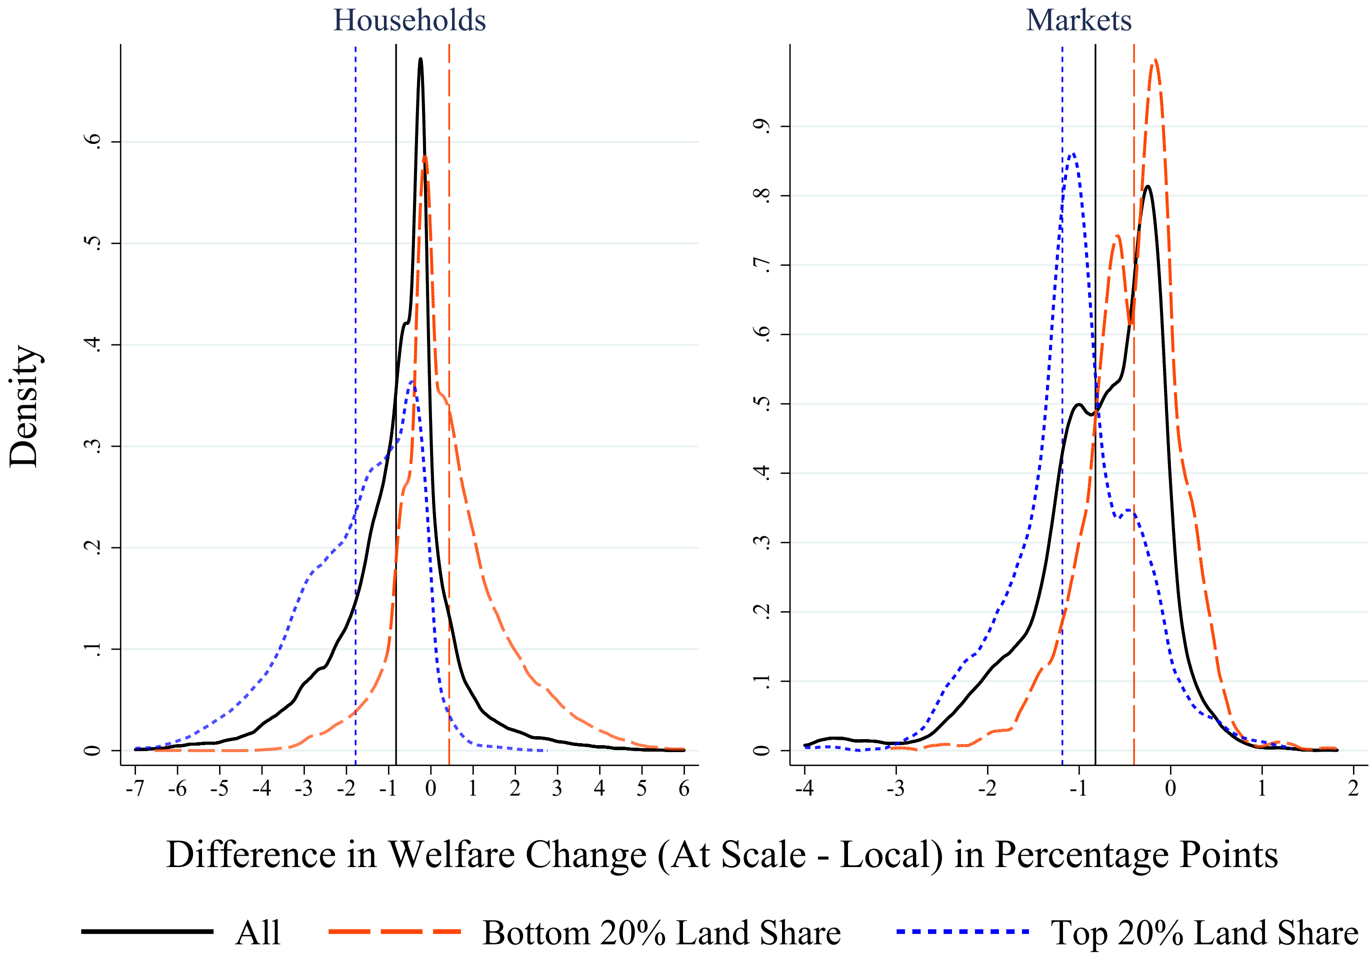

The figure plots distributions of the difference in welfare changes from at-scale versus local interventions in percentage points for the identical representative sample of roughly 100k randomly selected rural households (left panel), and their averages across parishes (right panel). Vertical bars indicate mean differences. Overall differences are plotted in black. The blue line shows the effects for the top quintile of land shares in initial household income, while the red line shows the same effects for the bottom quintile. See Section 5 for discussion.

Figure 3: Distributional Implications

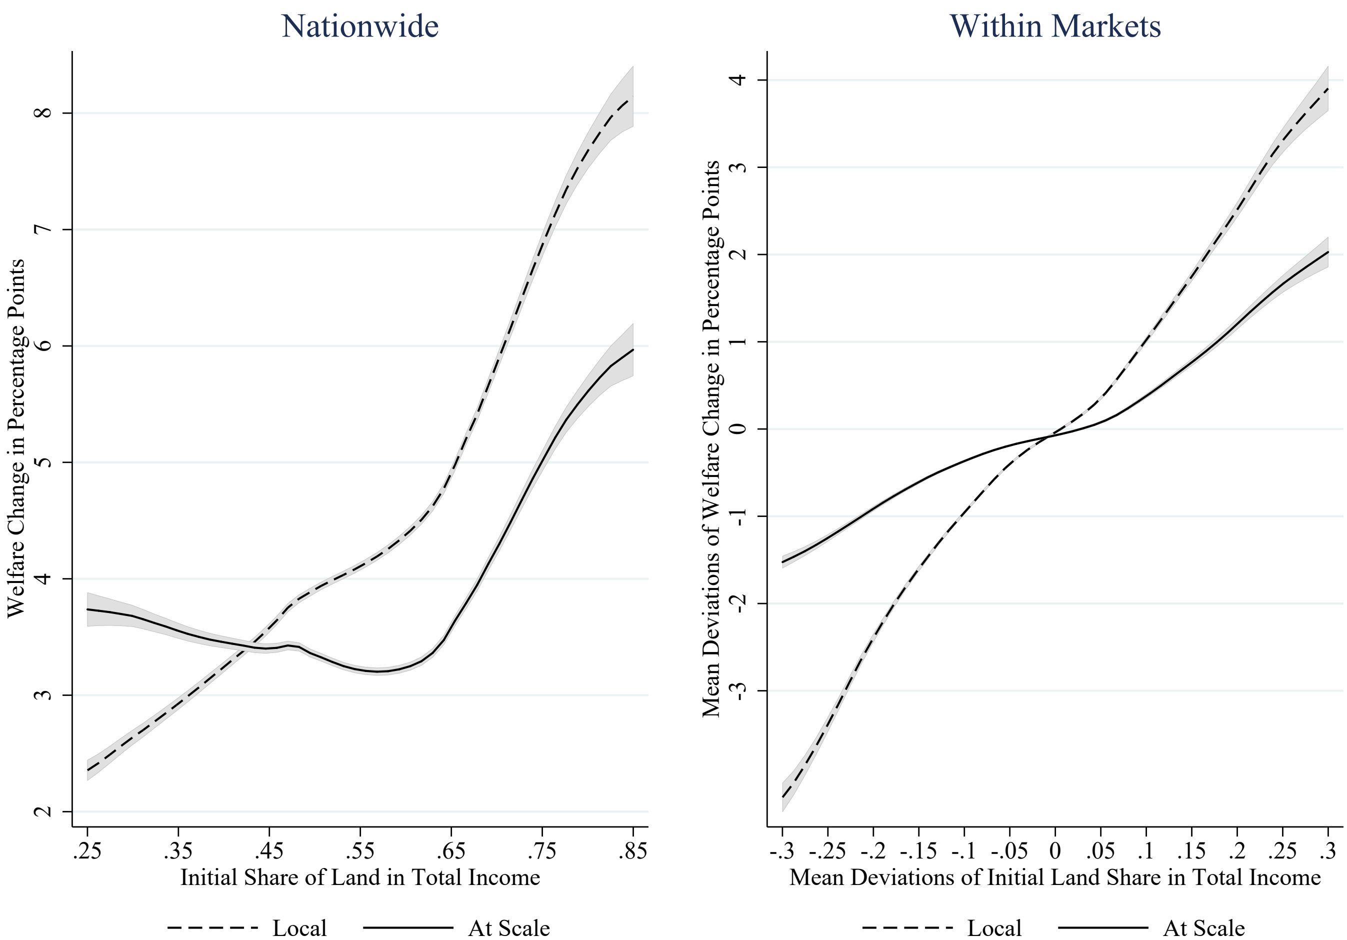

The figure plots the welfare changes resulting from the local and at-scale interventions, respectively, as a function of initial land shares. Estimates are for an identical representative sample of roughly 100k rural Uganda households. The right panel uses deviations from the parish means on both axes instead of the levels plotted in the left panel. Estimates are from local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 5 for discussion.

Figure 4: Effects as a Function of Remoteness

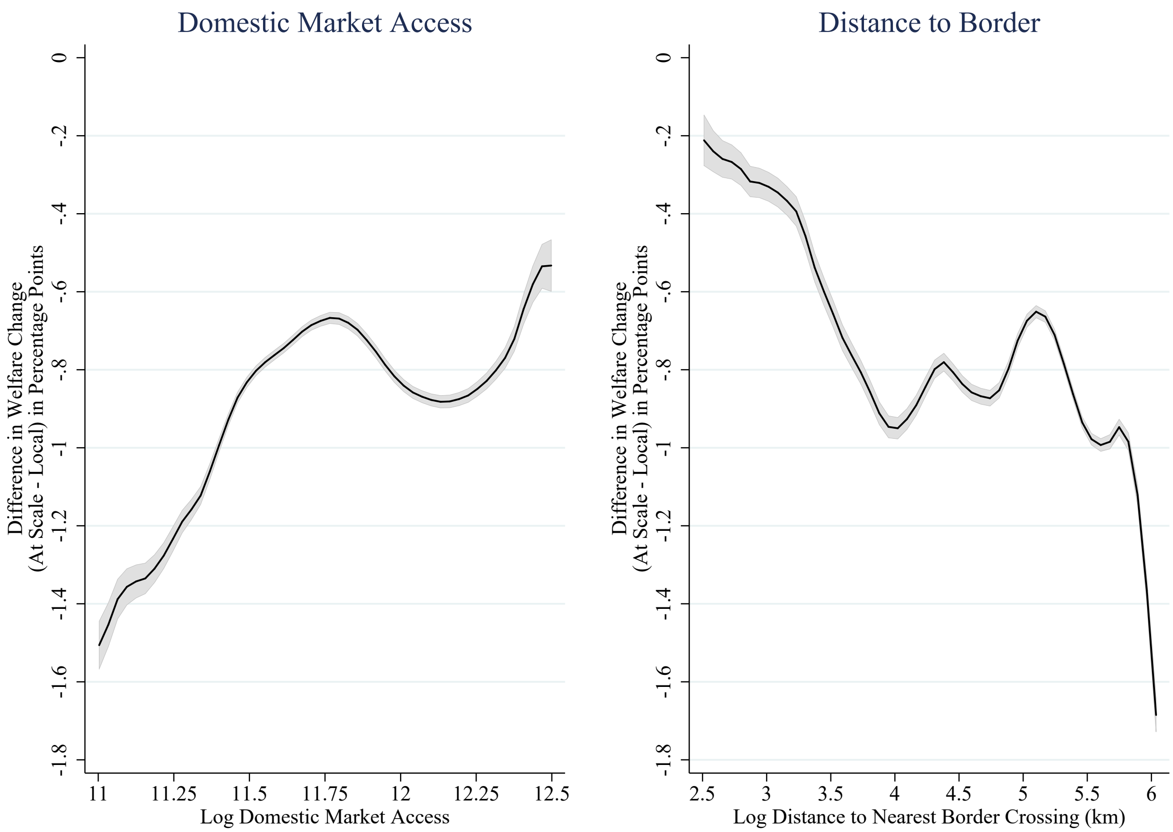

The figure plots the difference in households welfare impact (at scale - local) in percentage points on the y-axis. The left panel uses the log of the inverse-distance weighted sum of populations in all other markets and cities in Uganda ($\sum_{d \neq o} \frac{Pop_d}{Distance_{od}}$) on the x-axis, with distance measures in km. The right panel plots the same relationship with respect to the log distance to the nearest border crossing in km on the x-axis. Estimates are from local polynomial regressions based on the representative sample of roughly 100k rural Ugandan households. Shaded areas indicate 95 percent confidence intervals. See Section 5 for discussion.

Figure 5: GE Forces as a Function of Saturation

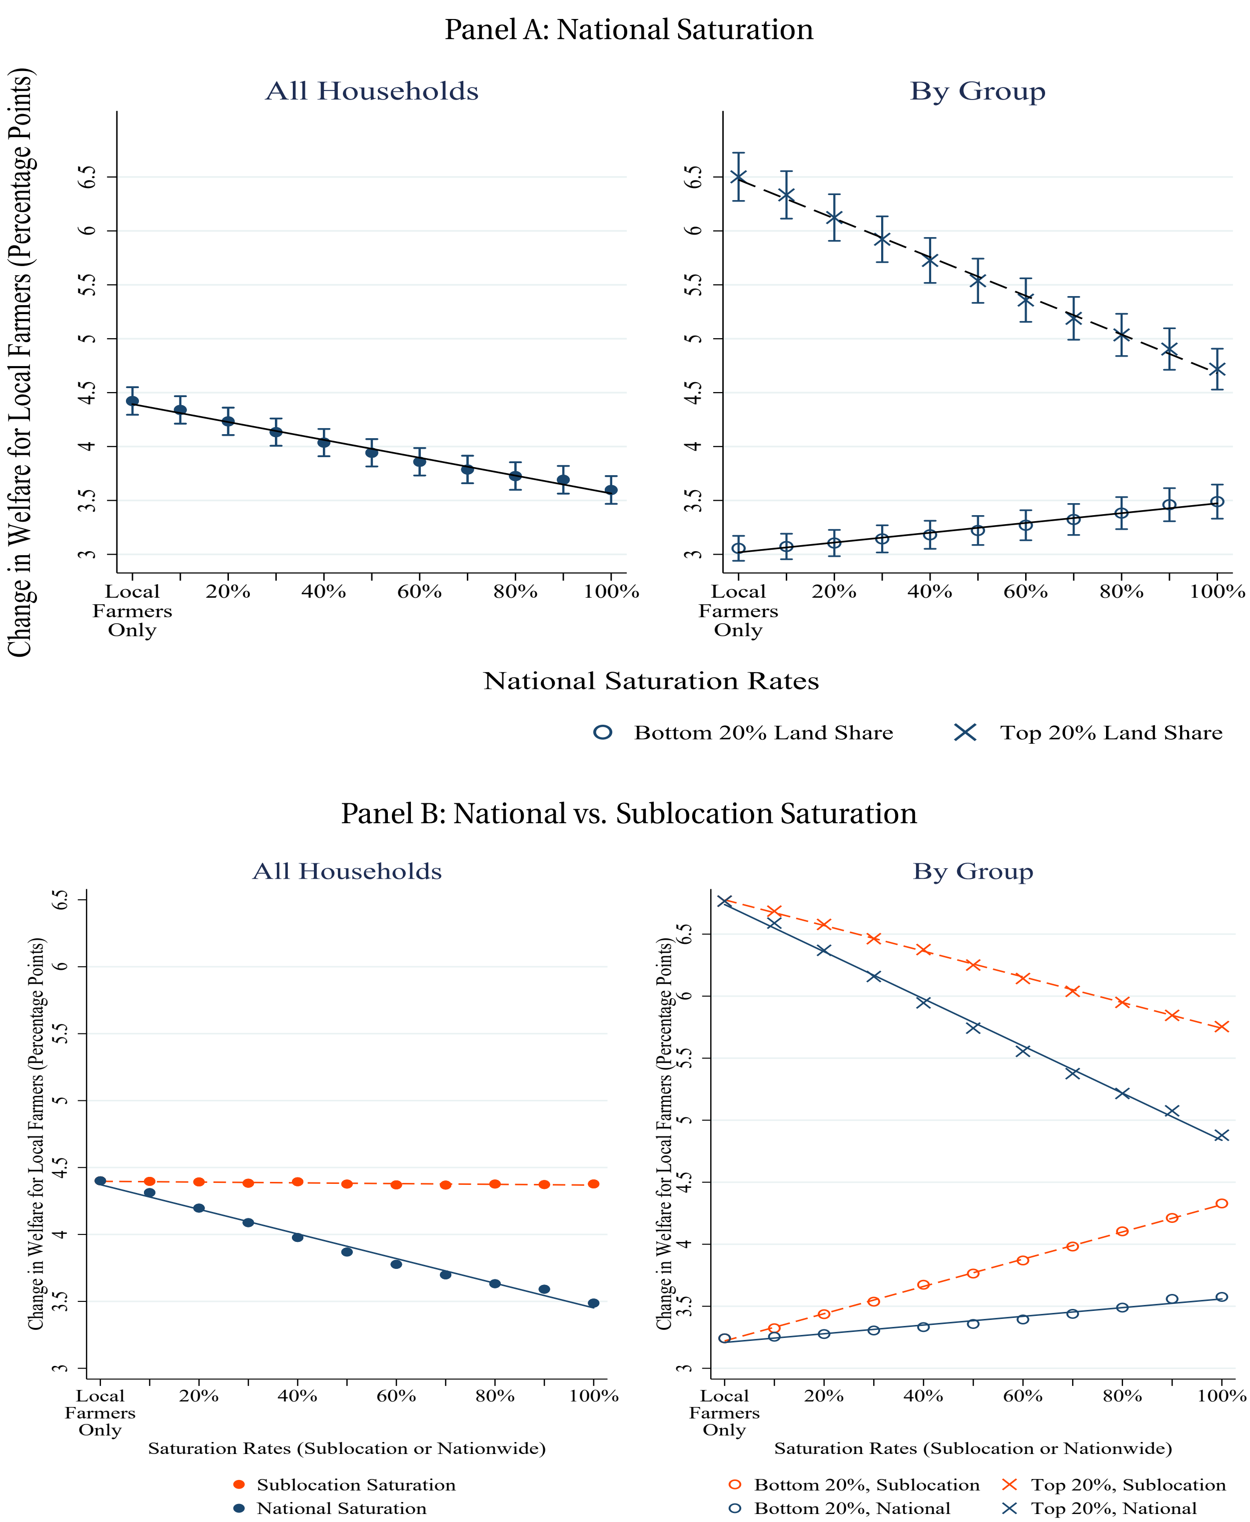

Panel A presents average welfare effects among the sample of roughly 100k Ugandan rural households as a function of national saturation rates (steps of 10% of the rural population randomly selected from each parish). Confidence intervals are at the 95% level. Panel B shows results for roughly 6500 households from the same 100k sample located in 51 randomly selected subcounties (one in each district). Blue markers depict average effects for this group across nationwide saturation rates as in Panel A. Orange markers depict the average effects across saturation rates only within the 51 subcounties, leaving the rest of Uganda untreated.

Figure 6: Role of a Granular Economic Geography

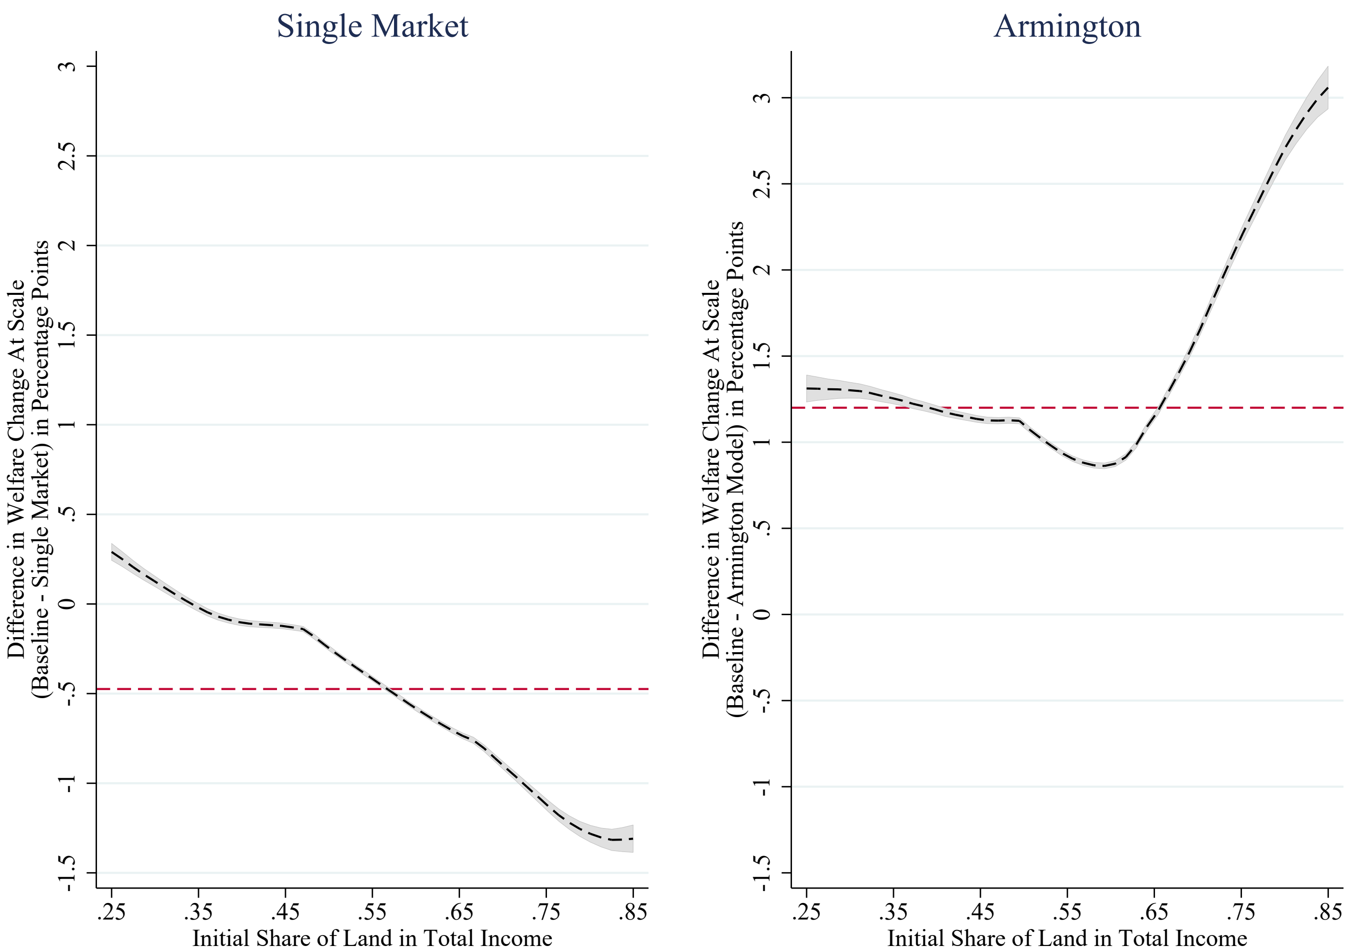

The figure plots local polynomial regressions of the percentage point difference in welfare effects across the roughly 100k rural households between our baseline model and the alternative model as a function of initial land income shares. Panel A is based on the alternative assumption of a single integrated national market. Panel B is based on the alternative assumption of an Armington model with iceberg trade costs and parish-level product differentiation in agriculture. Shaded areas indicate 95 percent confidence intervals. The dotted red lines indicate the sample average of these differences. See Section 5 for discussion.

Figure 7: Sensitivity to Alternative Parameters

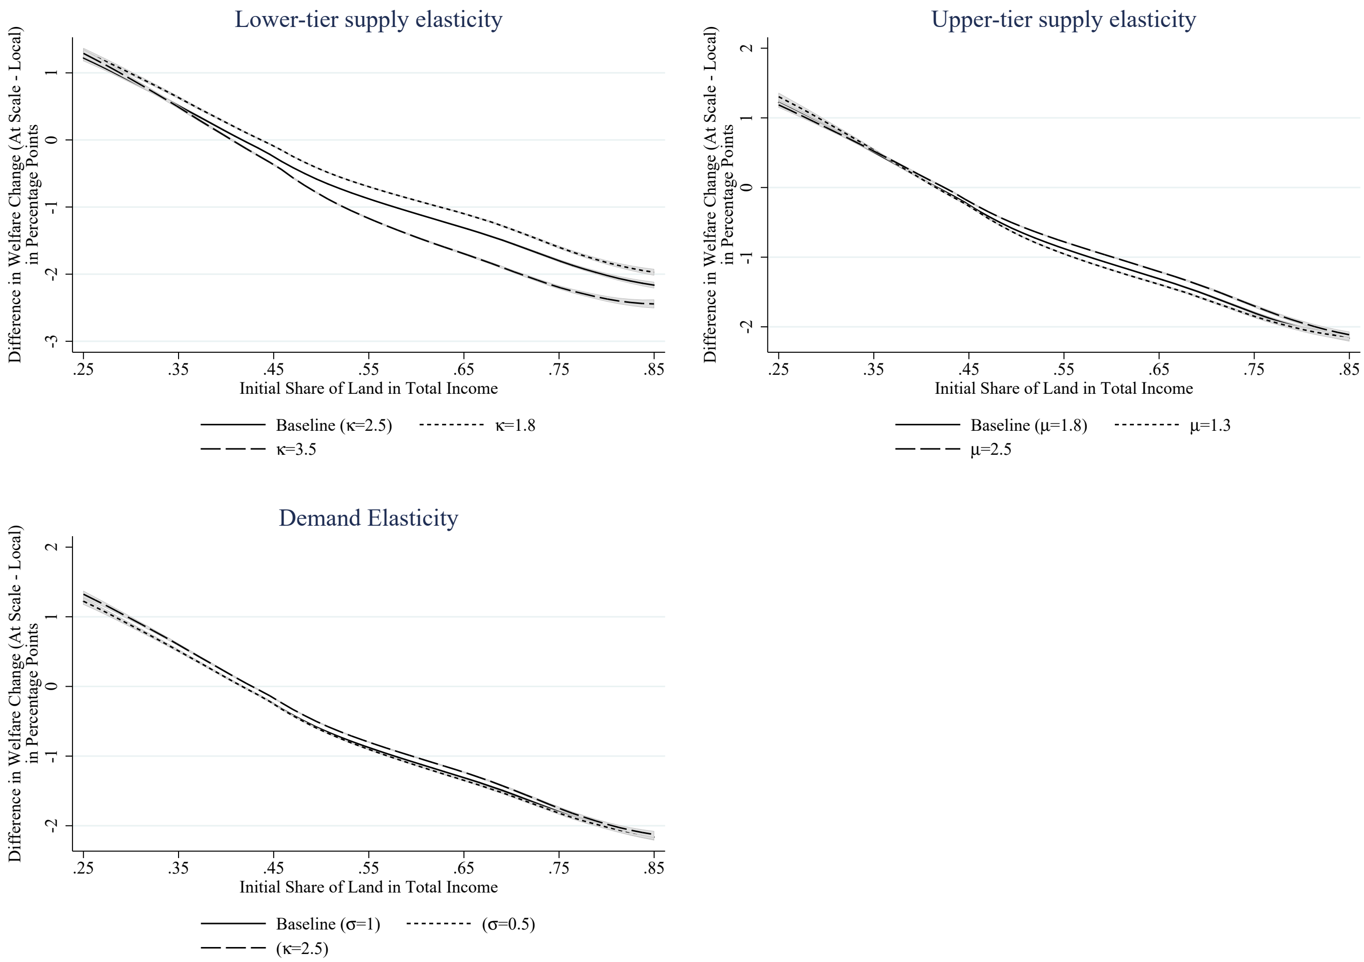

The figure plots estimates from local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 5 for discussion.

Figure 8: Model Validation Using Price Data and Trade Flows from Trader Surveys

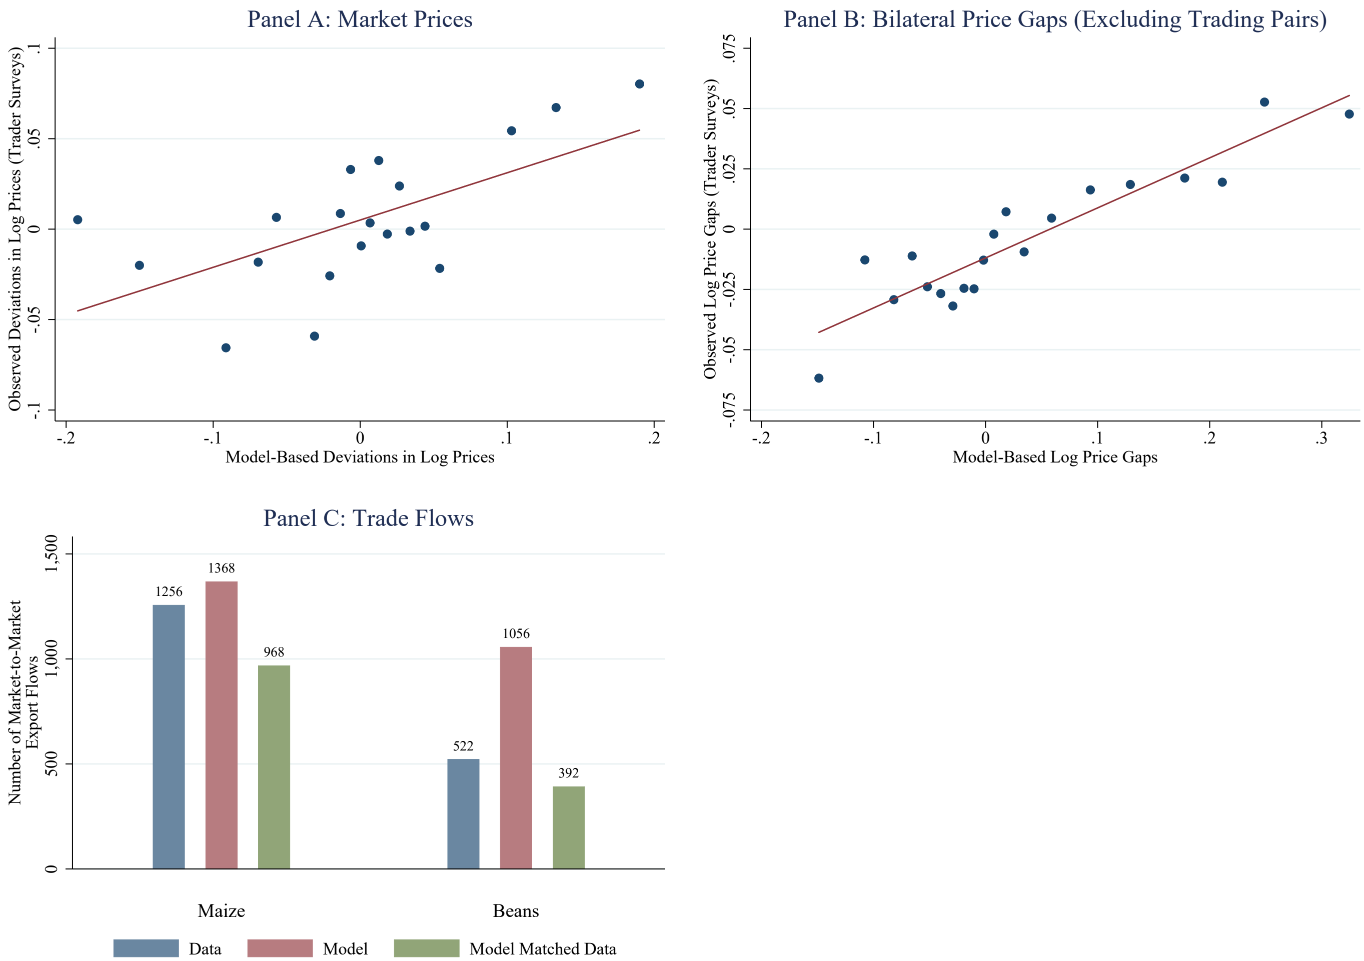

Panel A is based on 260 markets with price data in the trader surveys (Bergquist *et al.* , 2022) collected for each of 38 months during 2015-2018 and for two crops (maize and beans). We take the median market price for each crop and market in a given month. The y-axis of the binned scatter plot shows the residuals from a regression of the log median market prices on month-by-crop fixed effects. The x-axis displays mean deviations of log prices for the same two crops across the same markets in the baseline equilibrium (the results from the price discovery algorithm). Panels B and C make use of an additional dataset from the trader surveys, covering trade flows between the 260 markets for maize and beans for each of the 12 months in 2016. In Panel B, we convert the market price dataset from Panel A for the 12 months of 2016 into bilateral price pairs (counting each pair only once (not twice) per month and crop). The y-axis of the binned scatter plot in Panel B are origin-destination bilateral log price gaps for each month and crop, excluding trading pairs. The x-axis are model-based bilateral log price gaps for the same markets and crops based on the price discovery algorithm of our model. In Panel C we compare the trade flows reported across the markets in the trader surveys for either maize or beans for each month of data in 2016 to the bilateral trade flows from the models' price discovery algorithm. See Section 5 for discussion.

Tables

Table 1: Estimation of σ

Dependent Variable is Log Quantities (Instrument is Randomized Subsidy Amounts)

	(1)	(2)	(3)	(4)
VARIABLES	OLS	IV	OLS	IV
Log P	-4.8067***	-0.9446	-5.0225***	-1.0020*
	(0.2686)	(0.6230)	(0.3362)	(0.5473)
Observations	1,247	1,247	1,247	1,247
Market-Day FX	Yes	Yes	No	No
Market-Day-Seller FX	No	No	Yes	Yes
1st Stage F-Stat		321		659

See Section 4 for discussion. Standard errors clustered at level of communities. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2: Estimation of κ

Dependent Variable is $\log \frac{\pi_{i1|kt}}{\pi_{i0|kt}}$ (Instrument is RCT Treat Indicator)

	First Stage		Reduced Form		IV	
	(1)	(2)	(3)	(4)	(5)	(6)
	Cross-Section	Panel	Cross-Section	Panel	Cross-Section	Panel
Treat	-0.75***	-0.75***	0.62*	0.64*		
	(0.05)	(0.05)	(0.36)	(0.36)		
Log Input Price					-0.83*	-0.85
					(0.49)	(0.50)
Observations	63	127	63	127	63	127
Community FX		Yes		Yes		Yes
Round FX		Yes		Yes		Yes
F-Stat					204.57	204.51

See Section 4 for discussion. Standard errors clustered at level of communities. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 3: Estimation of μ

Panel A: First Stage Regressions. Dependent Variable is $log(\pi_{i,k,t})$

	(1)	(2)	(3)	(4)
	log$(\pi_{i,k,t})$	log$(\pi_{i,k,t})$	log$(\pi_{i,k,t})$	log$(\pi_{i,k,t})$
VARIABLES	All Years	All Years	2005-13	2005-13
IV0	-0.4644***	-0.3663**	-0.9344	-1.8073*
	(0.1216)	(0.1719)	(0.6402)	(1.0427)
Observations	27,650	27,647	4,580	4,580
HH-Crop FX	yes	yes	yes	yes
Crop-Year FX	yes	.	yes	.
Region-Crop-Year FX	no	yes	no	yes
Number of clusters	135	135	92	92

Panel B: Dependent Variable is Log Harvest $(log(q_{i,k,t}))$

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
VARIABLES	All Years	All Years	All Years	All Years	2005-13	2005-13	2005-13	2005-13
log$\pi_{i,k,t}$	0.3574***	0.7969*	0.3569***	0.6325	0.4146***	0.9248***	0.4253***	0.9038***
	(0.0164)	(0.4221)	(0.0164)	(0.4952)	(0.0341)	(0.3097)	(0.0325)	(0.2163)
Observations	27,966	27,650	27,963	27,647	4,486	4,282	4,480	4,276
HH-Crop FX	yes	yes	yes	yes	yes	yes	yes	yes
Crop-Year FX	yes	yes	.	.	yes	yes	.	.
Region-Crop-Year FX	no	no	yes	yes	no	no	yes	yes
Number of clusters	135	135	135	135	95	95	95	95
1st Stage F-Stat		14.60		4.543		32.81		17.93

Panel C: Dependent Variable is Log Adjusted Output

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
VARIABLES	All Years	All Years	All Years	All Years	2005-13	2005-13	2005-13	2005-13
log$\pi_{i,k,t}$	0.4108***	0.4007	0.4061***	0.7895	0.4411***	0.5529	0.4382***	0.7537**
	(0.0358)	(0.5423)	(0.0362)	(0.7251)	(0.0601)	(0.4214)	(0.0620)	(0.3154)
Observations	27,966	27,650	27,963	27,647	4,486	4,282	4,480	4,276
HH-Crop FX	yes	yes	yes	yes	yes	yes	yes	yes
Crop-Year FX	yes	yes	.	.	yes	yes	.	.
Region-Crop-Year FX	no	no	yes	yes	no	no	yes	yes
Number of clusters	135	135	135	135	95	95	95	95
1st Stage F-Stat		14.60		4.543		32.81		17.93

See Section 4 for discussion. Standard errors clustered at level of counties. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Scaling Agricultural Policy Interventions: Appendix (Not for Publication)

Appendix 1 describes the database used in the estimation. Appendix 2 uses the data to document stylized facts that inform our theory in Section 2. Appendix 3 provides additional figures and tables that we reference in the main text and in the stylized facts below. Appendix 4 presents additional details of the model and solution method.

Appendix 1 Data

Our analysis makes use of six main datasets. This appendix provides additional details and descriptive statistics.

Uganda National Panel Survey (UNPS)

The UNPS is a multi-topic household panel collected by the Ugandan Bureau of Statistics as part of the World Bank's Living Standards Measurement Survey. The survey began as part of the 2005/2006 Ugandan National Household Survey (UNHS). Then starting in 2009/2010, the UNPS set out to track a nationally representative sample of 3,123 households located in 322 enumeration areas that had been surveyed by the UNHS in 2005/2006. The UNPS is now conducted annually. Each year, the UNPS interviews households twice, in visits six months apart, in order to accurately collect data on both of the two growing seasons in the country. In particular, the main dataset that we assembled contains 77 crops across roughly 100 districts and 500 parishes for the periods 2005, 2009, 2010, 2011 and 2013. It includes detailed information on agriculture, such as crop production, the amount of land allocated to each crop, labor and non-labor inputs used in each plot and technology used at the household-parcel-plot-season-year. Data on consumption of the household contains disaggregated information on expenditures broken up across crops and other consumption.

Uganda Population and Housing Census 2002

The Ugandan Census has been conducted roughly every ten years since 1948. Collected by the Ugandan Bureau of Statistics, it is the major source of demographic and socio-economic statistics in Uganda. Over the span of seven days, trained enumerators visited every household in Uganda and collected information on all individuals in the household. At the household level, the Census collects the location (down to the village level), the number of household members, the number of dependents, and ownership of basic assets. Then for each household member, the Census collects information on the individual's sex, age, years of schooling obtained, literacy status, and source of livelihood, among other indicators. We have access to the microdata for the 100 percent sample of the 2002 Census.

GIS Database and Border Prices

We use several geo-referenced datasets. We use data on administrative boundaries and detailed information on the transportation network (covering both paved and non-paved feeder roads) from Uganda's Bureau of Statistics. We complement this database with geo-referenced information on crop suitability from the Food and Agricultural Organization (FAO) Global Ago-Ecological Zones (GAEZ) database. This dataset uses an agronomic model of crop production to convert data on terrain and soil conditions, rainfall, temperature and other agro-climatic conditions to calculate the potential production and yields of a variety of crops. We use this information as part of the projection from the UNPS sample to the Ugandan population at large. Finally, we use information on world prices of crops and intermediate inputs at Uganda's border from the FAO statistics database.

Survey Data on Cross-Market Trade Flows and Trade Costs

The survey data collected by Bergquist *et al.* (2022) captures cross-market trade flows and can be used to calibrate between-market transportation costs. They collect trade flow data in a survey of maize and beans traders located in 260 markets across Uganda (while not nationally representative, these markets are spread throughout the country). Traders are asked to list the markets in which they purchased and sold each crop over the previous 12 months. This information can be used to limit the calibration of cross-market trade costs to market pairs between which there were positive trade flows over a given period. They complement this data with a panel survey, collected in each of the 260 markets every two weeks for three years (2015-2018), in which prices are measured for maize, beans, and other crops. A greater description of the data collection can be found in Bergquist *et al.* (2022).

Demand Estimation

To estimate the slope of the demand curve for crops in Sections 2 and 4, we bring to bear transaction-level microdata from maize markets in rural Kenya that was collected as part of an experiment in Bergquist & Dinerstein (2020). Though for our purposes these subjects would ideally be representatively drawn from the same area in which the at-scale policy will be implemented, rural areas across East Africa share many features, including crops grown, farming methods (mostly rain-fed agriculture), and overall levels of development. This is especially true for the rural area of western Kenya studied in Bergquist & Dinerstein (2020), which takes place 30km from the Ugandan border. In their experiment, which took place in open-air maize markets, individual consumers who approached maize traders to make a purchase were offered a surprise price discount, the size of which was randomized across ten possible amounts. The value of the discount ranged from from roughly 0-15% of the baseline price and was randomized across customers within a given market-day. Using the subsidy as exogenous variation in consumer prices, the experiment measured resulting quantities purchased. We use these experimental data to estimate our key demand elasticity.

Supply Estimation

To estimate the key supply elasticity governing farmers' choice of land allocation across modern or traditional planting technologies, we exploit experimental variation from Carter *et al.* (2020). In this RCT, randomly selected farmers in Mozambique were offered fertilizer and improved seeds at a subsidized price. Data collected on farmers' use of modern technologies and output by plot allows estimation of the impact of changing input prices (instrumented by treatment) on land allocations across technologies. We complement this RCT with a natural experiment in the UNPS microdata that allows us to estimate the upper-tier supply elasticity in our model for substitution of land allocations across crops.

Appendix 2 Stylized Facts

In this appendix, we use the data described above to document the empirical context and a number of well-known stylized facts about agricultural trade across markets.

Major Crops, Regional Specialization and Price Gaps, Subsistence, Trading and Land Allocations

Appendix Figures A.1, A.2 and Tables A.1-A.5 present a number of basic stylized facts about the empirical context. Unless otherwise stated, these are drawn from the UNPS panel data of farmers. First, Table A.1 documents that the 9 most commonly grown crops (matooke (banana), beans, cassava, coffee, groundnuts, maize, millet, sorghum and sweet potatoes) account for 99 percent of the land allocation for the median farmer in Uganda (and for 86 percent of the aggregate land allocation).

Second, Figure A.1 and Table A.2 document a significant degree of regional specialization in Ugandan agricultural production across regions. Table A.2 provides information that these regional differences translate into meaningful variation in regional market prices across crops: the across-district variation in average crop prices accounts for 20-60 percent of the total variation in observed farm-gate prices.

Third, Table A.3 documents that the majority of all farmers are either net sellers or net buyers, rather than in subsistence, and this holds across each of the 9 major crops. The table also presents evidence that there are significant movements in and out of subsistence, conditional on having observed subsistence at the farmer level in a given season. Fourth, Table A.4 documents that farmers buy and sell their crops mostly in local markets, which in turn are connected to other markets through wholesale traders. Finally, Table A.5 documents that farmers frequently reallocate their land allocations across crops over time.

Product Differentiation Across Farmers

Appendix Table A.6 looks at evidence on product differentiation across farmers. The canonical approach in models of international trade sets focus on trade in manufacturing goods across countries, where CES demand coupled with product differentiation across manufacturing varieties imply that all bilateral trading pairs have non-zero trade flows. In an agricultural setting, however, and focusing on households instead of entire economies, this assumption would likely be stark. Consistent with this, the survey data collected by Bergquist *et al.* (2022) suggest that less than 5 percent of possible bilateral trading connections report trade flows in either of the crops covered by their dataset (maize and beans). This finding reported in Table A.6 provides corroborating evidence that agricultural crops in the Ugandan empirical setting are unlikely well-captured by the assumption of product differentiation across farmers who produce the crops. Our solution method will explicitly account for these zero trade flows and allow for endogenous switching on and off of trade flows as a result of treatment at-scale.

Nature of Trade Costs

The magnitude and nature of trade costs between farmers and local markets and across local markets play an important role for the propagation of output and factor price changes between markets along the transportation network. The canonical assumption in models of international trade is that trade costs are charged ad valorem (as a percentage of the transaction price). Ad valorem trade costs have the convenient feature that they enter multiplicatively on a given bilateral route, so that the pass-through of cost shocks at the origin to prices at the destination is complete (the same percentage change in both locations). In contrast, unit trade costs –charged per unit of the good, e.g. per sack or kg of maize– enter additively and have the implication that price pass-through is a decreasing function of the unit trade costs paid on bilateral routes. Market places farther away from the origin of the cost shock experience a lower percentage change in destination prices, as the unit cost makes up a larger fraction of the destination's market price.

To explore the nature of trade costs across Ugandan markets, we replicate results reported in Bergquist *et al.* (2022). Specifically, we estimate:

$$t_{odkt} = (p_{dkt} - p_{okt}) = \alpha + \beta p_{okt} + \theta_{od} + \phi_t + \epsilon_{odkt}$$

where t_{odkt} are per-unit trade costs between origin o and destination d for crop k (maize or beans) observed in month t, p_{okt} are origin unit prices, θ_{od} are origin-by-destination fixed effects, and ϕ_t are month fixed effects. Alternatively, origin-by-destination-by-month fixed effects (θ_{odt}) can be included.

Following Bergquist *et al.* (2022), we estimate these specifications conditioning on market pairs for which we observe positive trade flows in a given month. If trade costs include an ad valorem component, we would expect the coefficient β to be positive and statistically significant. On the other hand, if trade costs are charged per unit of the shipment (e.g. per sack), we would expect the point estimate of β to be close to zero.

One concern when estimating these specifications is that the origin crop price p_{okt} appears both on the left and the right-hand sides of the regression, giving rise to potential correlated measurement errors. This would lead to a mechanical negative bias in the estimate of β. To address this concern, we also report IV estimation results in which we instrument for the origin price in a given month with the price of the same crop in the same market observed in the previous month.

As reported in Table A.7, we find that β is slightly negative and statistically significant in the OLS regressions, but very close to zero and statistically insignificant after addressing the concern of correlated measurement errors in the IV specification. Taken together with existing evidence from field work (e.g. Bergquist & Dinerstein (2020)), these results suggest that trade costs in this empirical setting are best-captured by per-unit additive transportation costs.

Household Preferences

Appendix Figure A.2 reports a non-parametric estimate of the household Engel curve for food consumption. We estimate flexible functional forms of the following specification:

$$FoodShare_{it} = f\left(Income_{it}\right) + \theta_{mt} + \epsilon_{it}$$

where θ_{mt} is a parish-by-period fixed effect and $f(Income_{it})$ is a potentially non-linear function of household i's total income in period t. The inclusion of market (parish)-by-period fixed effects implies that we are comparing how the expenditure shares of rich and poor households differ while facing the same set of prices and shopping options. As reported in the figure, the average food consumption share ranges from 60 percent among the poorest households to about 20 percent among the richest households within a given market-by-period cell. In our model, these nonhomothetic preferences will allow for distributional effects due to changing food prices that result from the scaled intervention.

Modern Technology Adoption

Many policy interventions that are run through agricultural extension programs are aimed at providing access, information, training and/or subsidies for modern technology adoption among farmers. One important question in this context is whether adopting modern production techniques could be captured by a Hicks-neutral productivity shock to the farmers' production functions for a given crop. Alternatively, adopting modern techniques could involve more complicated changes in the production function, affecting the relative cost shares of factors of production, such as land and labor.

To provide some descriptive evidence on this question, we run specifications of the following form:

$$LaborShare_{ikt} = \alpha + \beta ModernUse_{ikt} + \theta_m + \phi_k + \gamma_t + \epsilon_{ikt}$$

where $LaborShare_{ikt}$ is farmer i's the cost share of labor relative to land (including both rents paid and imputed rents) for crop k in season t (there are two main seasons per year), $ModernUse_{ikt}$ is an indicator whether the farmer uses modern inputs for crop k in season t (defined as chemical fertilizer or hybrid seeds), and θ_{mkt}, ϕ_k and γ_t are district, crop and season fixed effects. Alternatively, we also include individual farmer fixed effects (θ_i).

As reported in appendix Table A.8, we find that the share of labor costs relative to land costs increases significantly as a function of whether or not the farmer uses modern production techniques. This holds both before and after the inclusion of farmer fixed effects (using variation only within-farmer across crops or over time). These results suggest that modern technology adoption is unlikely to be well-captured by a simple Hicks-neutral productivity shift in the production function. As a result, interventions at scale that affect the use of modern technologies may also have knock-on effects on local labor demand and wages. Our model will allow for such effects.

Appendix 3 Additional Figures and Tables

Table A.1: Main Crops

	(1)	(2)
	Aggregate	Median
VARIABLES	Share of Land	Share of Land
cropID==Beans	0.1442	0.1072
	(0.0086)	(0.0078)
cropID==Cassava	0.1908	0.0917
	(0.0121)	(0.0063)
cropID==Coffee	0.0718	0.0000
	(0.0048)	(0.0000)
cropID==Groundnuts	0.0541	0.0000
	(0.0052)	(0.0000)
cropID==Maize	0.1723	0.0923
	(0.0119)	(0.0052)
cropID==Matooke	0.1646	0.0089
	(0.0040)	(0.0089)
cropID==Millet	0.0315	0.0000
	(0.0021)	(0.0000)
cropID==Sorghum	0.0524	0.0000
	(0.0037)	(0.0000)
cropID==Sweet Potatoes	0.0886	0.0259
	(0.0061)	(0.0070)
Observations	45	45
Total Share	.859	.986

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Aggregate and median shares for each of the 9 crops are computed for each of four years of data from the UNPS. The table reports the means and standard deviations across the 4 rounds of data. See Appendix 1 for discussion and Section 3 for description of the data.

Figure A.1: Regional Specialization

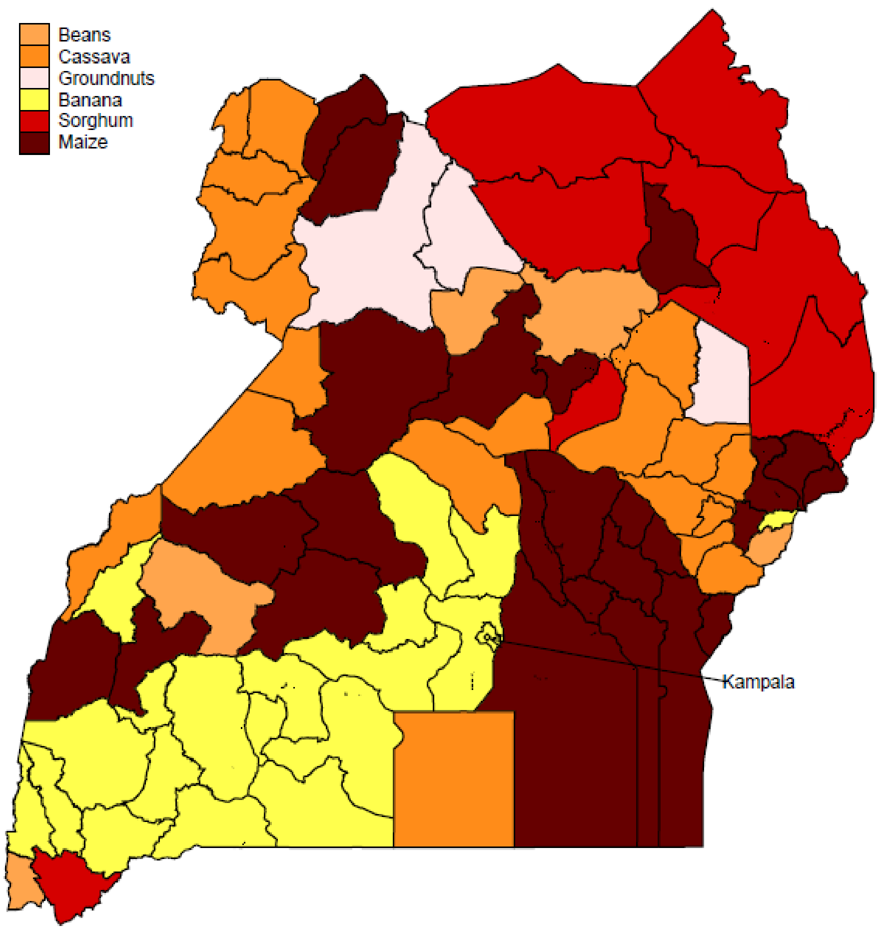

The figure displays the crop with the highest land allocation in each Ugandan district. We use the UNPS data to compute the mean of each crop's land shares across 4 rounds of data. See Appendix 1 for discussion and Section 3 for description of the data.

Table A.2: Regional Price Gaps

	District Dummies		Urban dummy	
Crop	F-statistic	Adjusted R-sq	Urban coefficient	p-value
Maize	6.83***	0.29	0.32	0.00
Millet	2.59***	0.36	0.30	0.00
Sorghum	2.71***	0.30	0.16	0.25
Cassava	5.68***	0.22	0.07	0.09
Beans	4.75***	0.29	0.21	0.00
Groundnuts	2.22***	0.26	0.10	0.22
Simsim	3.69***	0.19	-0.01	0.88
Sweet Potatoes	7.95***	0.33	0.10	0.07
Banana	4.10***	0.13	0.01	0.87
Coffee	5.65***	0.62	0.12	0.63
District FE	yes	yes	yes	yes

See Appendix 1 for discussion and Section 3 for description of the data.

Table A.3: Farmer Trading vs Subsistence

Panel A

Crop	Subsistence	Net buyer	Net seller	Total
Maize	33.65	22.50	43.85	1,049
Millet	31.12	38.07	30.82	331
Sorghum	31.02	34.98	33.99	303
Beans	44.87	10.73	44.40	1,081
Groundnuts	32.38	22.61	45.01	491
Simsim	25.47	26.71	47.83	161
Sweet Potatoes	21.60	63.03	15.37	898
Cassava	43.91	33.54	22.56	1,157
Banana	44.11	15.71	40.18	764
Coffee	0.97	9.95	89.08	412
Total	34.27	27.88	37.85	6,647

Panel B

Year	Subsistence to Trade	Trade to Subsistence
2009	24.90	38.83
2010	22.38	30.65
2011	24.61	31.32
2013	21.28	39.53
Average	23.35	35.00

See Appendix 1 for discussion and Section 3 for description of the data.

Table A.4: Farmers Sell Their Crops to Local Markets

Selling_Mode	Count_in_1000	Share
Government/LC	285.8	0.00400
Private trader in local village/market	44269	0.672
Private trader in district market	7081	0.107
Consumer at market	9744	0.148
Neighbor/ Relative	3907	0.0590
Other (specify)	610.6	0.00900
Total	65898	1

See Appendix 1 for discussion and Section 3 for description of the data.

Table A.5: Farmers Re-Allocate Their Land Across Crops Over Time

Panel A

Crop	Entry rate	Exit rate
Maize	46.79	16.98
Millet	13.03	42.21
Sorghum	7.87	45.30
Beans	34.10	9.78
Groundnuts	19.01	42.59
Simsim	5.07	45.19
Sweet Potatoes	37.39	31.07
Cassava	44.85	17.10
Banana Food	17.69	11.18
Coffee	9.66	18.84
Total	17.53	22.47

Panel B

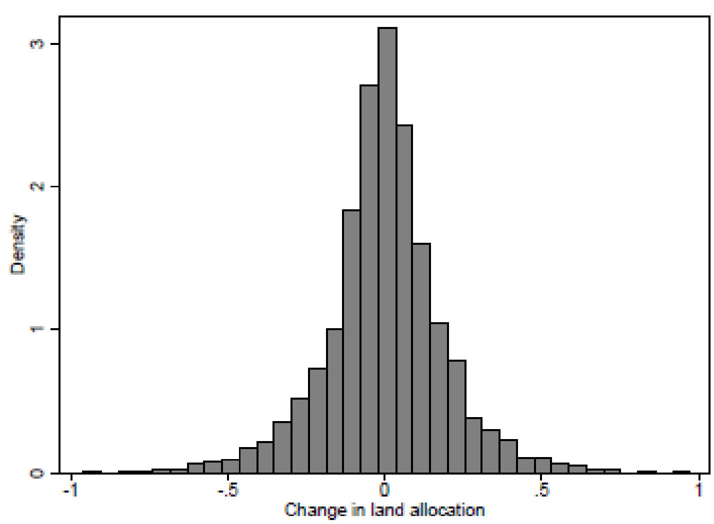

See Appendix 1 for discussion and Section 3 for description of the data.

Table A.6: Product Differentiation (Missing Trade Flows

	(1)	(2)
VARIABLES	Buying Dummy	Selling Dummy
Proportion_Trading	0.0429***	0.0432***
	(0.0021)	(0.0021)
Observations	9,146	9,146

*** p<0.01, ** p<0.05, * p<0.1

See Appendix 1 for discussion and Section 3 for description of the data.

Figure A.2: Household Preferences (Non-Homotheticity)

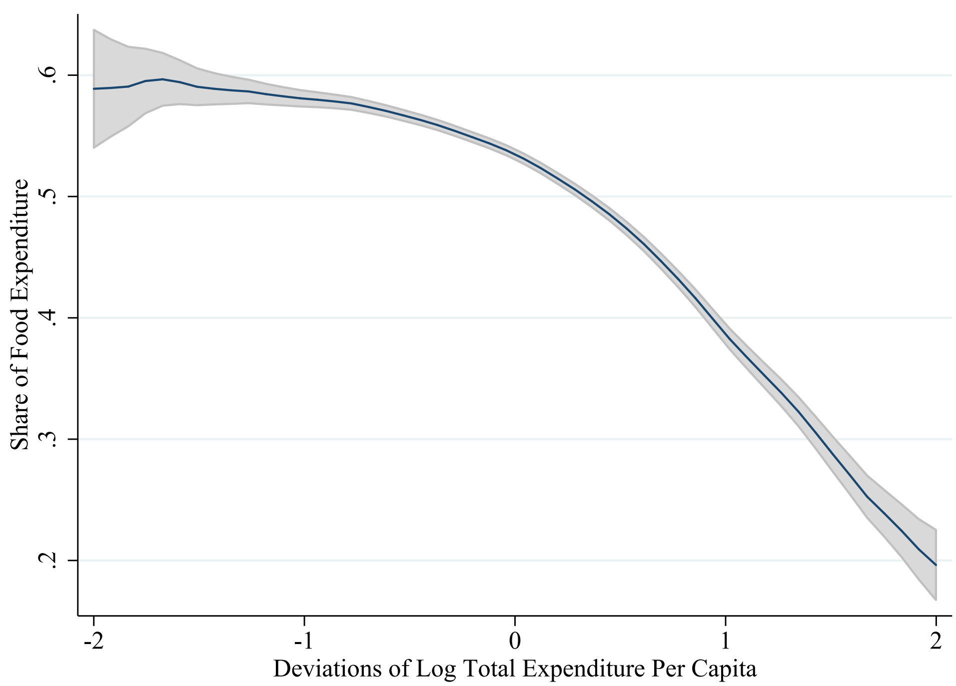

See Appendix 1 for discussion and Section 3 for description of the data.

Table A.7: Nature of Trade Costs

	(1)	(2)	(3)	(4)
	Price Gap	Price Gap	Price Gap	Price Gap
VARIABLES	OLS	OLS	IV (Lagged Price)	IV (Lagged Price)
Origin Price	-0.0605***	-0.0419**	-0.0081	-0.0002
	(0.0188)	(0.0206)	(0.0256)	(0.0274)
Observations	8,524	8,430	7,153	7,079
Pair FX	yes	.	yes	.
Month FX	yes	.	yes	.
Pair-by-Month FX	no	yes	no	yes

Standard errors clusterd at level of bilateral pairs.

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

See Appendix 1 for discussion and Section 3 for description of the data.

Table A.8: Technology Adoption and Production Cost Shares

	(1)	(2)
VARIABLES	Labor Share	Labor Share
Use Modern	0.1056***	0.0423***
	(0.0126)	(0.0112)
Observations	26,037	25,889
District FX	yes	.
Crop FX	yes	yes
Season FX	yes	yes
Farmer FX	no	yes

Standard errors clusterd at level of farmers.

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

See Appendix 1 for discussion and Section 3 for description of the data.

Table A.9: Calibrated Cost Shares in Production

VARIABLES	(1) Land Share Traditional	(2) Labor Share Traditional	(3) Intermediate Share Traditional	(4) Land Share Modern	(5) Labor Share Modern	(6) Intermediate Share Modern
cropID1==Beans	0.5107	0.4893	0.0000	0.4607	0.3852	0.1541
	(0.0259)	(0.0259)	(0.0000)	(0.0041)	(0.0139)	(0.0154)
cropID1==Cassava	0.5566	0.4434	0.0000	0.4429	0.3785	0.1786
	(0.0503)	(0.0503)	(0.0000)	(0.0180)	(0.0187)	(0.0176)
cropID1==Coffee	0.6777	0.3223	0.0000	0.5428	0.2683	0.1889
	(0.0571)	(0.0571)	(0.0000)	(0.0164)	(0.0202)	(0.0122)
cropID1==Groundnuts	0.5134	0.4866	0.0000	0.4204	0.4253	0.1543
	(0.0231)	(0.0231)	(0.0000)	(0.0190)	(0.0450)	(0.0271)
cropID1==Maize	0.5000	0.5000	0.0000	0.4153	0.4335	0.1512
	(0.0272)	(0.0272)	(0.0000)	(0.0520)	(0.0559)	(0.0159)
cropID1==Matooke	0.6343	0.3657	0.0000	0.6180	0.2564	0.1256
	(0.0455)	(0.0455)	(0.0000)	(0.0394)	(0.0275)	(0.0119)
cropID1==Millet	0.5285	0.4715	0.0000	0.5485	0.3381	0.1134
	(0.0174)	(0.0174)	(0.0000)	(0.0074)	(0.0039)	(0.0035)
cropID1==Sorghum	0.5563	0.4437	0.0000	0.5774	0.3321	0.0905
	(0.0216)	(0.0216)	(0.0000)	(0.0062)	(0.0060)	(0.0051)
cropID1==Sweet Potatoes	0.5088	0.4912	0.0000	0.4721	0.3642	0.1637
	(0.0258)	(0.0258)	(0.0000)	(0.0735)	(0.0800)	(0.0107)

See Section 4 for discussion and Section 3 for description of the data.

Figure A.3: Relative World Price Changes Over the Sample Period

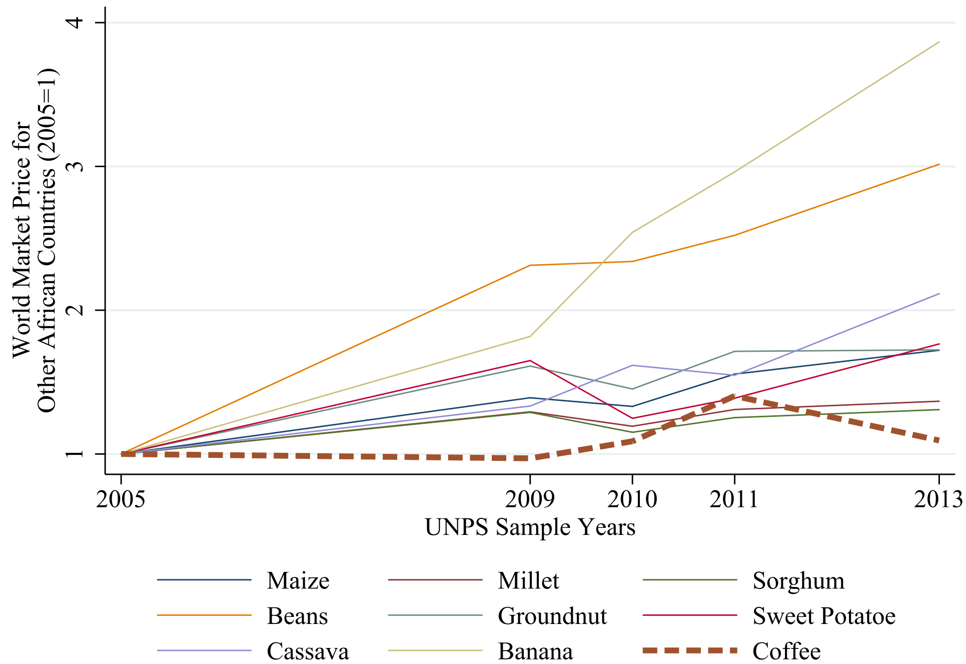

See Section 4 for discussion of the data.

Table A.10: Predicted Local Trade Costs and Measures of Remoteness

	(1) Transport Cost per unit	(2) Distance to Community Road	(3) Distance to District Road	(4) Distance to Gravel Road	(5) Distance to Tarmac Road	(6) Hiring Dummy Out-of-Sample
Crop Trade Costs						
Predicted $t_{im}/100$	0.358**	0.503***	2.001***	3.859***	6.120***	
	(0.181)	(0.135)	(0.635)	(0.975)	(2.344)	
Observations	544	6,331	5,460	2,282	805	
Labor Trade Costs						
Predicted $t^L_{im}/100$		0.024	0.093	0.092	-0.015	-0.061***
		(0.022)	(0.068)	(0.168)	(0.369)	(0.005)
Observations		6,317	5,448	2,275	803	7,853

See Section 4 for discussion. All distances are measured in km. Mean share of HHs hiring-in labor is 42% outside estimation sample. Standard errors clustered at level of households. *** p<0.01, ** p<0.05, * p<0.1

Figure A.4: Land Income Shares, Land Ownership and Household Incmomes

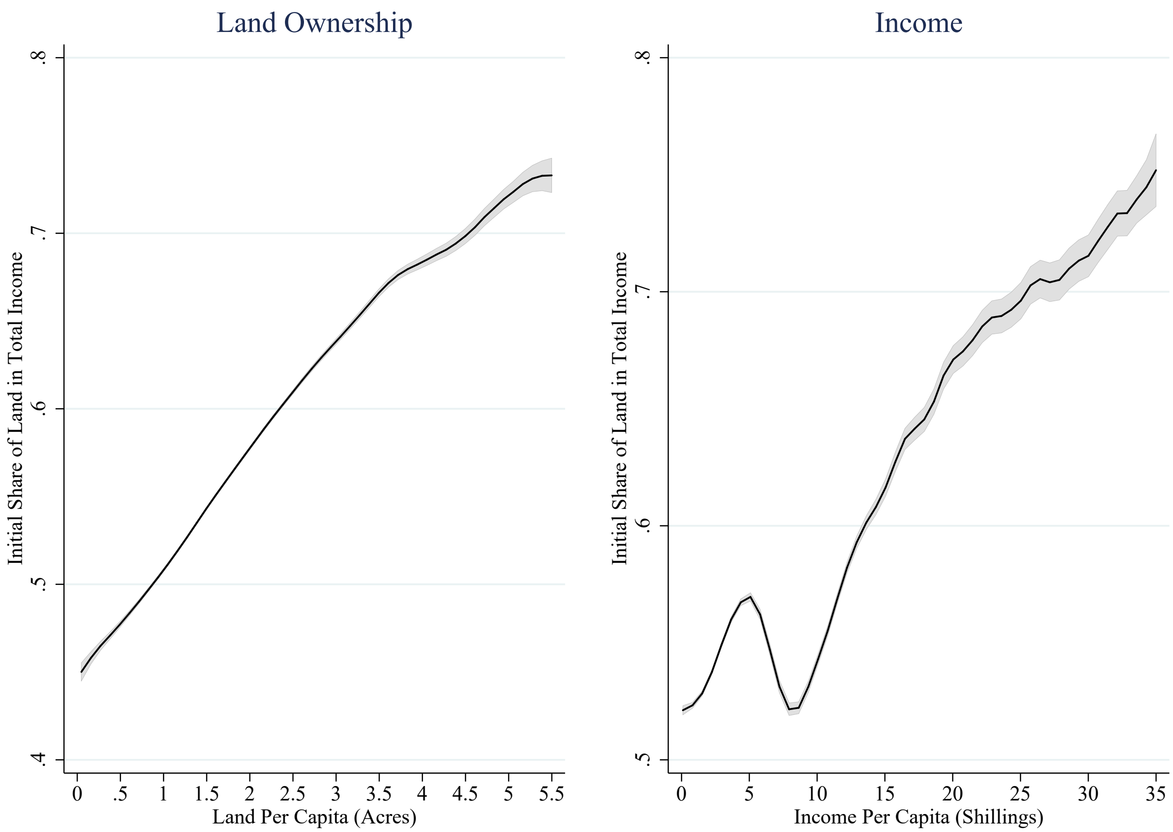

The figure plots estimates from local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 5 for discussion.

Table A.11: Effect on Household Welfare

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Welfare	Welfare	Welfare	Welfare	Welfare	Welfare	Welfare	Welfare
	Local	At Scale	Local	At Scale	Local	At Scale	Local	At Scale
VARIABLES	All Households	All Households	Bottom 20%	Bottom 20%	Middle 20%	Middle 20%	Top 20%	Top 20%
Percentage Point Change	4.42***	3.60***	3.06***	3.49***	4.05***	3.02***	6.50***	4.72***
	(0.07)	(0.07)	(0.06)	(0.08)	(0.07)	(0.08)	(0.11)	(0.10)
Observations	104,361	104,361	19,829	19,829	19,828	19,828	20,872	20,872
No Clusters	4502	4502	3577	3577	4130	4130	4087	4087

Standard errors clustered at market-level.

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

See Section 5 for discussion.

Table A.12: Channels

Panel A: Local Effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Income	Wage	P_manu	P_banana	P_bean	P_cassava	P_coffee	P_groundnut	P_maize	P_millet	P_sorghum	P_sweetpot
VARIABLES	Local	Local	Local	Local	Local	Local	Local	Local	Local	Local	Local	Local
Effect	4.3325***	0.6550***	0.0000	-0.0481***	-0.0314***	-0.0150**	-0.0012***	-0.0260***	-0.5200***	0.0259***	0.0101***	0.1064***
	(0.0648)	(0.0159)	(0.0000)	(0.0038)	(0.0026)	(0.0072)	(0.0003)	(0.0085)	(0.0194)	(0.0020)	(0.0008)	(0.0038)
Observations	104,361	104,361	104,361	104,361	104,361	104,361	104,361	104,361	104,361	104,361	104,361	104,361
No Clusters	4502	4502	4502	4502	4502	4502	4502	4502	4502	4502	4502	4502

Panel B: At-Scale Effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Income	Wage	P_manu	P_banana	P_bean	P_cassava	P_coffee	P_groundnut	P_maize	P_millet	P_sorghum	P_sweetpot
VARIABLES	At Scale	At Scale	At Scale	At Scale	At Scale	At Scale	At Scale	At Scale	At Scale	At Scale	At Scale	At Scale
Effect	3.4595***	2.8838***	0.9739***	-0.1075***	-0.3004***	-0.1369***	-0.0069***	0.1102***	-4.4588***	0.2085***	0.0146***	0.8932***
	(0.0698)	(0.0569)	(0.0035)	(0.0053)	(0.0102)	(0.0327)	(0.0006)	(0.0115)	(0.0455)	(0.0072)	(0.0030)	(0.0125)
Observations	104,361	104,361	104,361	104,361	104,361	104,361	104,361	104,361	104,361	104,361	104,361	104,361
No Clusters	4502	4502	4502	4502	4502	4502	4502	4502	4502	4502	4502	4502

Standard errors clustered at market-level.

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

The table presents effects from the local and from the intervention at scale for the identical representative sample of 10k randomly selected rural households. See Section 5 for discussion.

Figure A.5: Decomposition of Difference At Scale vs. Local Effect

The left panel presents the difference between the at-scale effect and the local effect on components of nominal incomes across the initial land share distribution, while the right panel presents the same for components of the household price index. Estimates are from local polynomial regressions based on the representative sample of roughly 100k rural Ugandan households. Shaded areas indicate 95 percent confidence intervals. See Section 5 for discussion.

Figure A.6: Wage and Land Income Effects (Without Income-Share Weights)

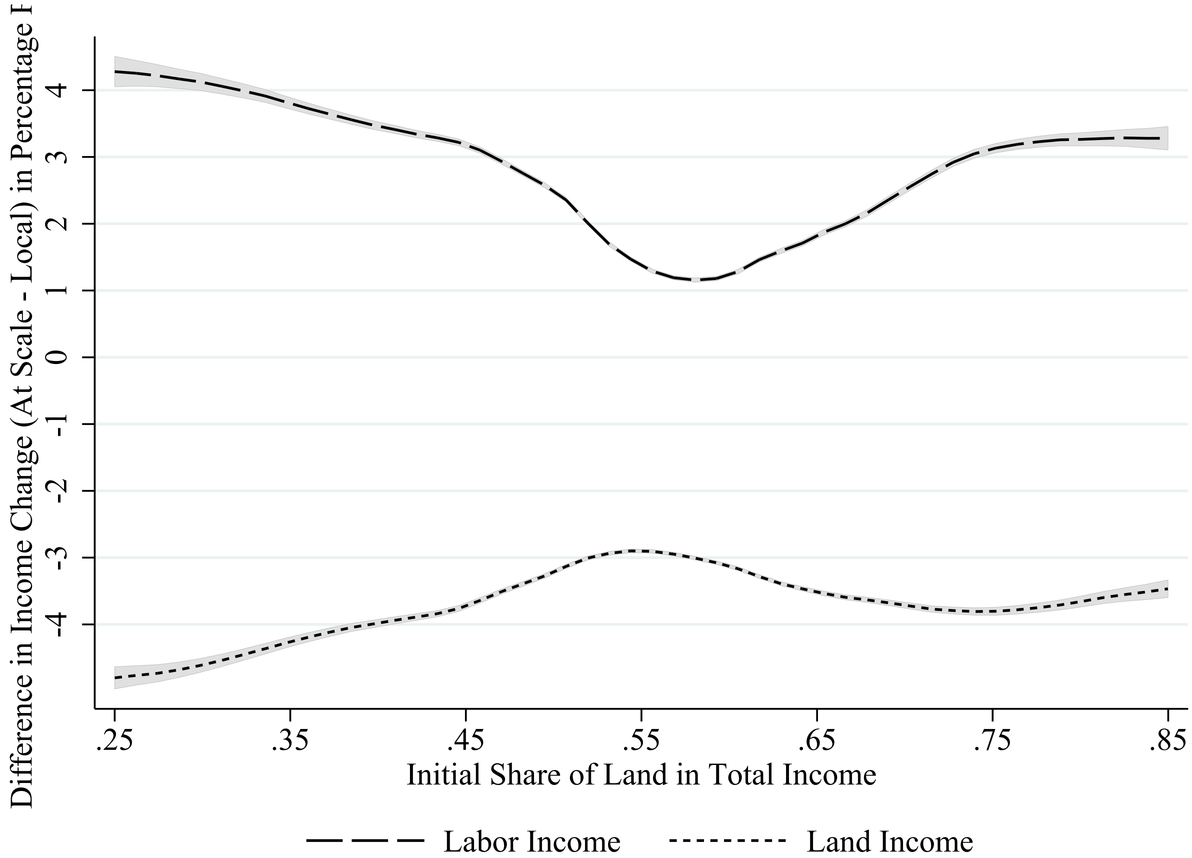

The figure plots estimates from local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 5 for discussion.

Figure A.7: Initial Usage of Modern Inputs Across Land-Poor vs Land-Rich Households

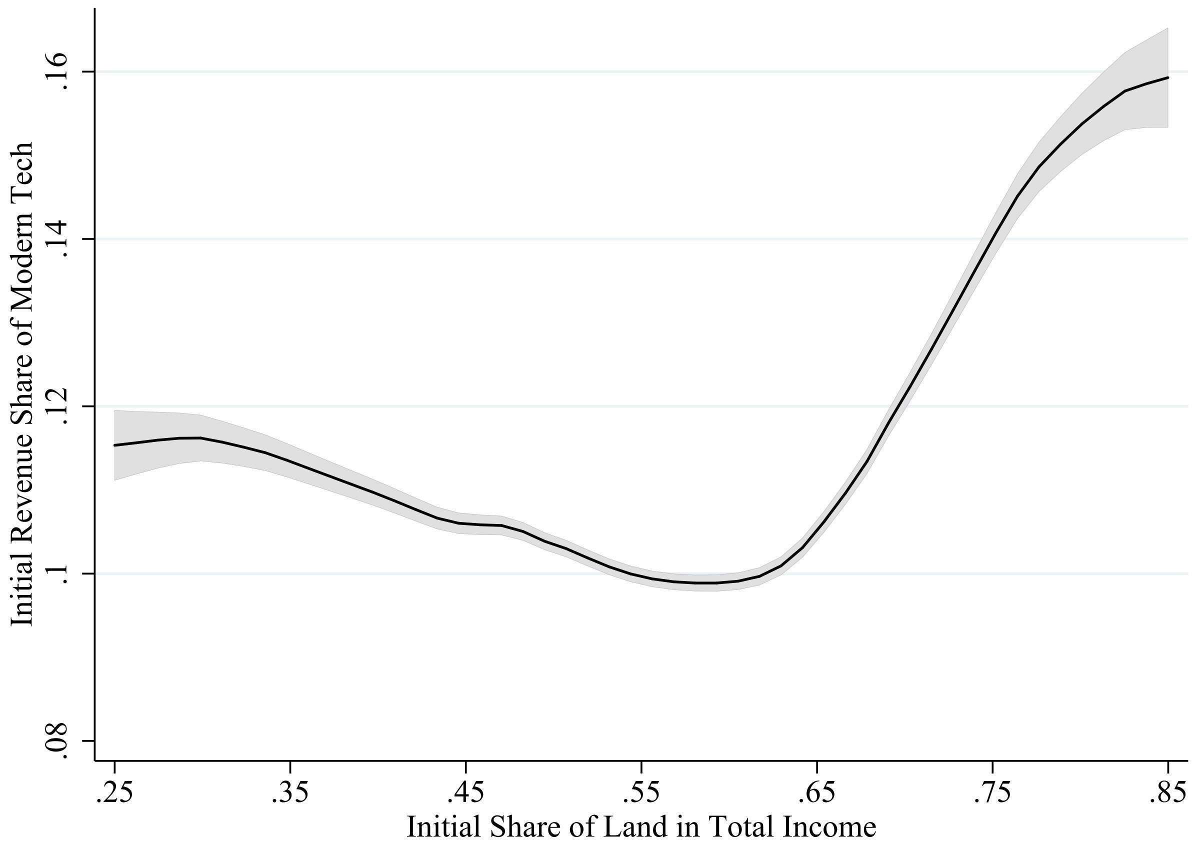

The figure plots estimates from local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 5 for discussion.

Figure A.8: Effects as a Function of Initial Usage of Modern Inputs

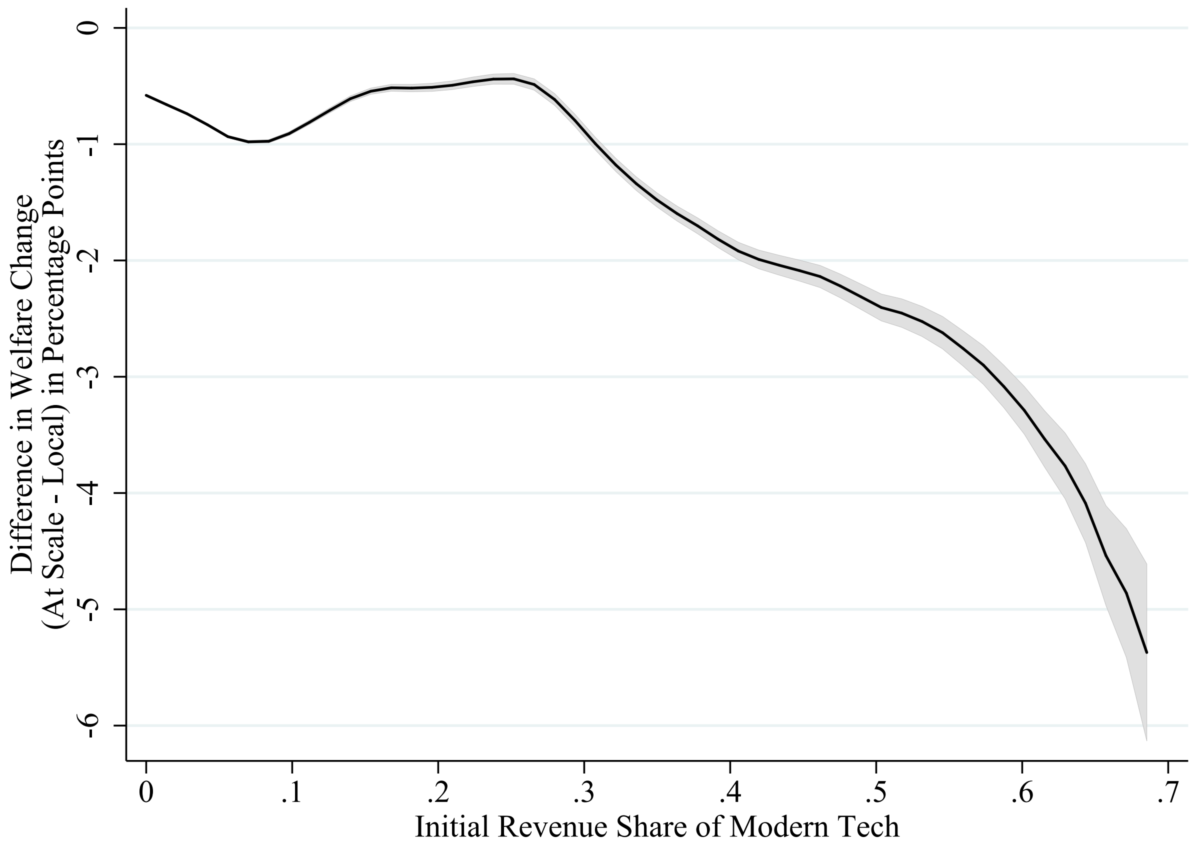

The figure plots estimates from local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 5 for discussion.

Figure A.9: Effects as a Function of Initial Crop Shares

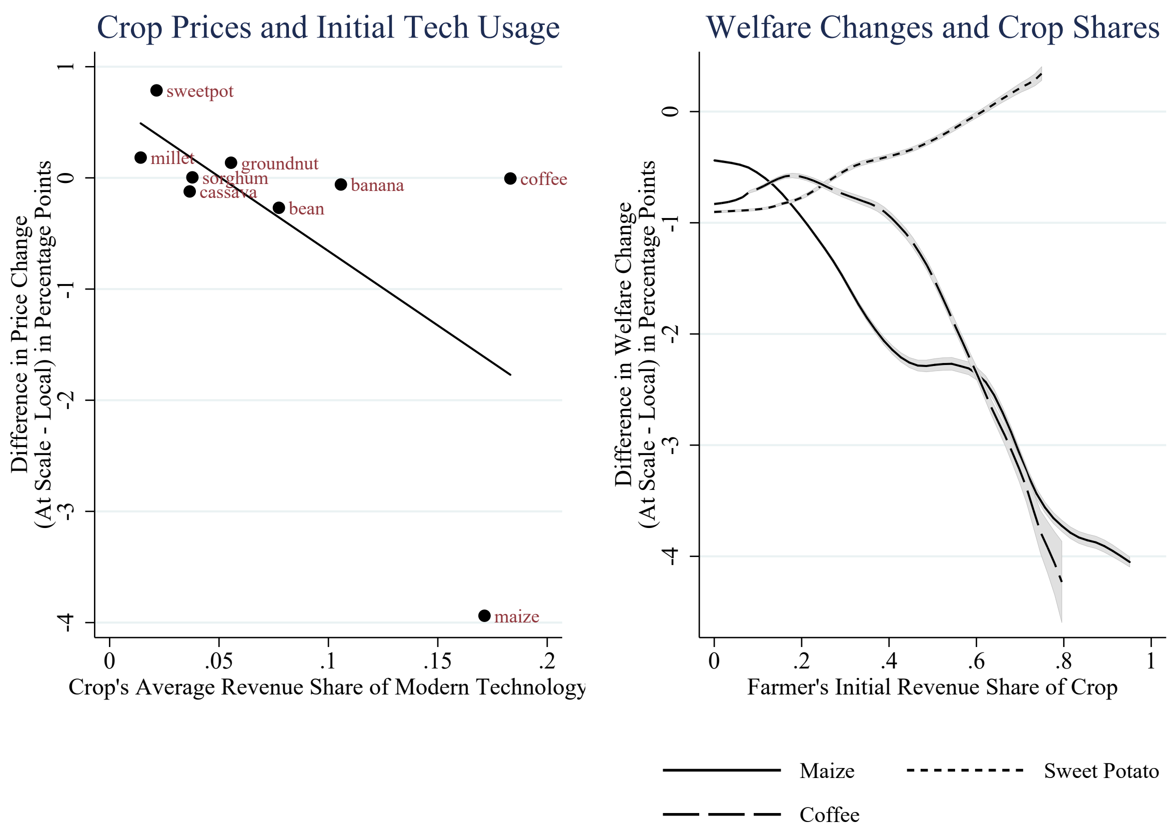

The figure plots estimates from local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 5 for discussion.

Appendix 4 Model and Solution Method

In Appendix 4.A, we first present the excess demand functions $\chi_{j,g}(\bullet)$ used in the text to define the equilibrium, and we then present the excess demand functions for the "price discovery" step. In Appendix 4.B, we present the model allowing for general functional forms on preferences and technology for which exact hat algebra is feasible. In Appendix 4.C, we formally describe this class of functions. In Appendix 4.D, we develop the proof for uniqueness in price discovery for the special case with iceberg trade costs. In Appendix 4.E, we provide additional details on recovering trade shares in manufacturing. In Appendix 4.F, we show that the introduction of hub-and-spoke trade costs leads to well-defined market prices. Finally, in Appendix 4.G we extend the model to incorporate cross-market migration of labor.

Appendix 4.A Excess Demand Functions

The excess demand function for farmers are given by

$$\chi_{i,g}\left(\left\{b_{i,k}p_{i,k}\right\}_i, \left\{r_{i,k,\omega}\right\}_i, \left\{p_{i,g}\right\}_i, I_i\right) = \begin{cases} \xi_g\left(\left\{b_{i,k}p_{i,k}\right\}_i, I_i\right) I_i - p_{i,g}\sum_\omega q_{i,g,\omega}\left(\{p_{i,g}\}_i, \{r_{i,k,\omega}\}_i\right) & \text{for } g \in \mathcal{K}_A, \\ \xi_g\left(\left\{b_{i,k}p_{i,k}\right\}_i, I_i\right) I_i & \text{for } g \in \mathcal{K}_\mathcal{M}, \\ \sum_{k\in\mathcal{K}_A,\omega} \alpha_{i,g,k,\omega}p_{i,k}q_{i,k,\omega}\left(\{p_{i,g}\}_i, \{r_{i,k,\omega}\}_i\right) - p_{i,g}L_i & \text{for } g = L. \end{cases}$$

The excess demand functions for urban households are given by

$$\chi_{h,g}\left(\left\{b_{h,k}p_{h,k}\right\}_h, \left\{r_{h,k,\omega}\right\}_h, \left\{p_{h,g}\right\}_h, I_h\right) = \begin{cases} \xi_g\left(\left\{b_{h,k}p_{h,k}\right\}_h, I_h\right) I_h & \text{for } g \in \mathcal{K}_A, \\ \left[\xi_g\left(\left\{b_{h,k}p_{h,k}\right\}_h, I_h\right) - \mathbb{1}\left(g = g(h)\right)\right] I_h & \text{for } g \in \mathcal{K}_\mathcal{M}, \end{cases}$$

where expenditure share function $\xi_g(\bullet)$ and crop output function $q_{i,g,\omega}(\bullet)$ are defined in the main text. Indicator function $\mathbb{1}\left(g = g(h)\right)$ is equal to one only if manufacturing variety g belongs to urban household h and zero otherwise.

Finally, for Foreign we have

$$\chi_{F,g}\left(\left\{b_{F,k}p_{F,k}\right\}_F, \left\{r_{F,k,\omega}\right\}_F, \left\{p_{F,g}\right\}_F, I_F\right) = \begin{cases} -\infty & \text{if} \quad p_{F,g} < p^*_{F,g} \\ ]-\infty, \infty[& \text{if} \quad p_{F,g} = p^*_{F,g} \\ \infty & \text{if} \quad p_{F,g} > p^*_{F,g}, \end{cases} \qquad \text{for } g \in \mathcal{K}_A,$$

$$\chi_{F,g}\left(\left\{b_{F,k}p_{F,k}\right\}_F, \left\{r_{F,k,\omega}\right\}_F, \left\{p_{F,g}\right\}_F, I_F\right) = X_{F,g}(p_{F,g}), \qquad \text{for } g \in \mathcal{K}_\mathcal{M} \setminus \{g(F)\}.$$

We include some variables (e.g., $\left\{r_{j,k,\omega}\right\}_j$ for $j \in \mathcal{H} \cup \{F\}$) that are not defined as arguments in the excess demand functions so that they cover all agents – this is not a problem because if the function does not depend on these arguments then there is no need to define them.

Excess demand as functions of data $\mathbb{D}_A$ and prices $\{p_{j,g}\}_{g\in\mathcal{K}_A\cup\{L\}}$ for farmers and urban households (used for the price discovery step) are given by

$$\chi_{i,g}\left(\{p_{i,g}\}_{g\in\mathcal{K}_A\cup\{L\}};\mathbb{D}_A\right)=\xi_{i,g}I_i\left(\{p_{i,g}\}_{g\in\mathcal{K}_A\cup\{L\}};\mathbb{D}_A\right)-\sum_{\omega}p_{i,g}q_{i,g,\omega},\qquad \text{for } g\in\mathcal{K}_A,$$

$$\chi_{i,g}\left(\{p_{i,g}\}_{g\in\mathcal{K}_A\cup\{L\}};\mathbb{D}_A\right)=\sum_{k\in\mathcal{K}_{\mathcal{A}},\omega}\alpha_{i,g,k,\omega}p_{i,k}q_{i,k,\omega}-p_{i,g}L_i\qquad \text{for } g=L,$$

$$\chi_{h,g}\left(\{p_{h,g}\}_{g\in\mathcal{K}_A\cup\{L\}};\mathbb{D}_A\right)=\xi_{h,g}I_h,\qquad \text{for } g\in\mathcal{K}_A,$$

$$\chi_{F,g}\left(\{p_{F,g}\}_{g\in\mathcal{K}_A\cup\{L\}};\mathbb{D}_A\right)=\begin{cases}-\infty & \text{if}\quad p_{F,g}<p^*_{F,g}\\ ]-\infty,\infty[& \text{if}\quad p_{F,g}=p^*_{F,g}\\ \infty & \text{if}\quad p_{F,g}>p^*_{F,g}\end{cases}\qquad \text{for } g\in\mathcal{K}_{\mathcal{A}},$$

where

$$I_i\left(\{p_{i,g}\}_{g\in\mathcal{K}_A\cup\{L\}};\mathbb{D}_A\right)=\sum_{k\in\mathcal{K}_A,\omega}\left(1-\sum_n\alpha_{i,n,k,\omega}\right)p_{i,k}q_{i,k,\omega}+p_{i,L}L_i.$$

Appendix 4.B General Functional Forms on Preferences and Technology

We now restate the assumptions on preferences and technology, but allowing for general functional forms that satisfy certain assumptions needed for exact hat-algebra (after the price discovery step), discussed formally in Appendix 4.C. The model equilibrium and solution to counterfactuals in the main text and excess demand functions in Appendix 4.A also apply for these more general functional forms.[1] The purpose of this exercise is to allow researchers to customize the model by choosing alternative preferences and technology, depending on their application of the model.

Preferences

Agents $j\neq F$ have indirect utility function $V\left(\{b_{j,k}p_{j,k}\}_j,I_j\right)$, where I_j denotes income, $\{p_{j,k}\}_j$ denotes prices and $\{b_{j,k}\}_j$ denotes demand shifters. Let $\xi_{j,k}$ denote the expenditure share of agent j on good k and $\xi_k\left(\{b_{j,k}p_{j,k}\}_j,I_j\right)$ the corresponding expenditure share function (assumed common across all agents in Home). Roy's identity implies that

$$\xi_{j,k}=\xi_k\left(\{b_{j,k}p_{j,k}\}_j,I_j\right)=-\frac{\frac{\partial\ln V\left(\{b_{j,k}p_{j,k}\}_j,I_j\right)}{\partial\ln p_{j,k}}}{\frac{\partial\ln V\left(\{b_{j,k}p_{j,k}\}_j,I_j\right)}{\partial\ln I_j}}.$$

Turning to Foreign, our small-open economy assumption for Home implies that Foreign's demand (in value) for manufacturing good $g(h)$ can be specified directly as a function of this

[1] With general functional forms for preferences and technology, the only change in excess demand functions presented above is for farmers' excess demand for labor ($g=L$), where we replace $\alpha_{i,g,k,\omega}$ with the cost share function $\alpha_{i,g,k,\omega}(\{p_{i,n}\}_i,r_{i,k,\omega})$.

goods's individual price, $X_{F,g(h)}(p_{F,g(h)})$.

Technology

Farmers produce agricultural goods $k \in \mathcal{K}_A$ using land, labor and intermediate goods with techniques $\omega \in \Omega$. Assuming constant returns to scale in agriculture, letting $r_{i,k,\omega}$ denote the return to an effective unit of land allocated by farmer i to produce agricultural good k with technique ω, and letting $c_{i,k,\omega}(\{p_{i,n}\}_i, r_{i,k,\omega})/a_{i,k,\omega}$ denote the corresponding unit cost function – with $a_{i,k,\omega}$ a Hicks-neutral productivity shifter – then at an interior solution to the farmer's optimization problem we must have

$$p_{i,k} = c_{i,k,\omega}(\{p_{i,n}\}_i, r_{i,k,\omega})/a_{i,k,\omega}.$$

This determines $r_{i,k,\omega}$ as an implicit function of prices, $p_{i,k}$ and $\{p_{i,n}\}_i$, and productivity $a_{i,k,\omega}$. In turn, letting $\alpha_{i,n,k,\omega}(\{p_{i,n}\}_i, r_{i,k,\omega})$ denote the cost share of input n for farmer i producing crop k using technique ω, an envelope result implies that

$$\alpha_{i,n,k,\omega}(\{p_{i,n}\}_i, r_{i,k,\omega}) = \frac{\partial \ln c_{i,k,\omega}(\{p_{i,n}\}_i, r_{i,k,\omega})}{\partial \ln p_{i,n}}.$$

Farmer i allocates their land endowment Z_i across different agricultural goods and techniques to maximize their total land returns, $\sum_{k,\omega} r_{i,k,\omega} Z_{i,k,\omega}$, where $Z_{i,k,\omega}$ measures the effective units of land allocated by farmer i to produce crop k with technique ω. We allow for decreasing marginal productivity in how physical units of land Z_i can be converted into efficiency units of land for different crops and techniques. Specifically, we assume that the feasible set for the allocation of efficiency units of land across crops and techniques is $\{\{Z_{i,k,\omega}\}_i | f(\{Z_{i,k,\omega}\}_i) \leq Z_i\}$, with $f(\bullet)$ homogeneous of degree one and strictly quasi-convex. Total land returns of farmer i are then given by

$$Y\left(\{r_{i,k,\omega}\}_i\right) Z_i \equiv \max_{\{Z_{i,k,\omega}\}_i} \sum_{k,\omega} r_{i,k,\omega} Z_{i,k,\omega} \quad \text{s.t.} \quad f(\{Z_{i,k,\omega}\}_i) \leq Z_i.$$

Letting $\pi_{i,k,\omega}$ denote the share of land returns of farmer i coming from production of crop k with technique ω, and $\pi_{k,\omega}\left(\{r_{i,k,\omega}\}_i\right)$ the corresponding function, an envelope result implies that

$$\pi_{i,k,\omega} = \pi_{k,\omega}\left(\{r_{i,k,\omega}\}_i\right) = \frac{\partial \ln Y\left(\{r_{i,k,\omega}\}_i\right)}{\partial \ln r_{i,k,\omega}}.$$

Finally, letting $q_{i,k,\omega}$ denote output of crop k for farmer i with technique ω, then

$$q_{i,k,\omega}\left(\{p_{i,g}\}_i, \{r_{i,k,\omega}\}_i\right) = \frac{\pi_{k,\omega}\left(\{r_{i,k,\omega}\}_i\right) Y\left(\{r_{i,k,\omega}\}_i\right) Z_i}{\left[1 - \sum_n \alpha_{i,n,k,\omega}(\{p_{i,n}\}_i, r_{i,k,\omega})\right] p_{i,k}}.$$

Turning to urban households, we assume that each urban area is associated with a single representative urban household who produces a differentiated manufacturing good. We keep the technology simple by assuming that manufacturing production is linear in labor, so that the quantity of manufacturing good $g(h)$ produced by urban household h is given by $a_h L_h$. Given that labor supply is perfectly inelastic, we can then treat $q_h \equiv a_h L_h$ as the urban households' endowment of manufacturing good $g(h)$.

It remains to parameterize all the relevant functions, namely $\xi_g(\bullet)$, $X_{F,g}(\bullet)$, $Y(\bullet)$, and $c_{i,k,\omega}(\bullet)$, and ensure that these functions are conducive to exact hat-algebra, as defined in the next section.

Appendix 4.C Functional Forms for Exact Hat Algebra

For a function $f(p)$ (e.g., expenditure shares, shares of land returns), exact-algebra entails writing $f(p') = g(f(p), \hat{p})$, where $g(,)$ is some function and $\hat{p} = p'/p$ denotes the vector of ratios (element-wise), so that we can solve for counterfactual $f(p')$ as a function of $f(p)$ without necessarily knowing p. Not all functions f, however, allow us to write $f(p')$ in this way. The goal of this appendix is to describe the class of such functions.

Definition *Let f be a smooth function from $\mathbb{R}^n$ to its image $Im(f) \subset \mathbb{R}^m$. We say that this function is "conducive to exact hat algebra" if we can write:*

$$f(p.\hat{p}) = g(f(p), \hat{p})$$

for all $p, \hat{p} \in \mathbb{R}^n_+$, for some function $g : Im(f) \times \mathbb{R}^n_+ \to \mathbb{R}^m$, and where $p.\hat{p}$ is the element-wise product of p and $\hat{p}$.

The following proposition provides a characterization of such functions:

Proposition *Suppose that f is a smooth function from $\mathbb{R}^n_+$ to $\mathbb{R}^m$. Then these three properties are equivalent:*

- *i) f is conducive to exact hat algebra.*
- *ii) For all $p_0, p_1, \hat{p} \in \mathbb{R}^n_+$,*

  $$f(p_0) = f(p_1) \implies f(p_0.\hat{p}) = f(p_1.\hat{p})$$

  (where $p.\hat{p}$ denotes the element-wise product).
- *iii) Consider $F(x) = f(\exp(x))$, where $\exp(x)$ denotes the vector of elements $\exp(x_i)$. There is a linear subspace E of $\mathbb{R}^n$ on which F is injective, and a linear function $\pi : \mathbb{R}^n \to E$, equal to the identity on E, such that*

  $$F(x) = F(\pi(x)), \forall x \in \mathbb{R}^n.$$

  This implies that level sets of F are affine, and that f can be written as a combination of Cobb-Douglas functions (exponential of π) and an invertible function.

Note that such definition and results may apply to the derivatives instead of the output function itself. For instance, with a production function featuring constant returns to scale, we can observe the initial values of the gradient (in log), which corresponds to the shares of the different inputs entering the production function. In such cases, we can use a similar approach if the gradient is itself conducive to exact hat algebra, according to the definition above. By integrating, we can then retrieve the total changes in the output function as a function of the initial values of the log-gradient and the changes in the arguments. Let $J(p) = \left\{ \frac{\partial \log f}{\partial \log p_i} \right\}$ denotes the gradient of $\log f$ in $\log p$, and assume that we can write $J(p.\hat{p}) = G(J(p), \hat{p})$ equal to a function G of the initial values of J and the changes in prices, $\hat{p}$, we have then:

$$\log f(p.\hat{p}) - \log f(p) = \int_{x=0}^{\log \hat{p}} J(p.\exp(x))\, dx = \int_{x=0}^{\log \hat{p}} G(J(p), \exp(x))\, dx.$$

The proposition above can then be applied to characterize the class of such function J and their primitives, $\log f$.

Proof of the Proposition For the proof, it is more convenient to take the log of each argument. Let us denote by $x = \log p$ the log of inputs and by $\delta = \log(p'/p)$ the log change, so that a relative change in variables becomes additive. Consider $F(x) = f(\exp(x))$, where $\exp(x)$ denotes the vector of elements $\exp(x_i)$.

Proof of i) implies ii) If i) is satisfied then we can write $F(x + \delta) = G(F(x), \delta)$. Suppose that $F(x_0) = F(x_1)$, we have then

$$F(x_0 + \delta) = g(F(x_0), \exp(\delta)) = g(F(x_1), \exp(\delta)) = F(x_1 + \delta)$$

Similarly, in terms of function f, with $p = \exp(x)$ and $\hat{p} = \exp(\delta)$, $f(p_0) = f(p_1)$ implies:

$$f(p_0.\hat{p}) = g(f(p_0).\hat{p}) = g(f(p_1).\hat{p}) = f(p_1.\hat{p})$$

Proof of ii) implies i) To prove the converse property, let's construct a function $K : Im(f) \to \mathbb{R}^n$ such that $F(K(y)) = y$ for all $y \in Im(F)$. Then, for all $y \in Im(f)$ and all $x \in \mathbb{R}^n$, define g as

$$g(y, \delta) = F(K(y) + \delta)$$

Mechanically, by definition of K, we have: $F(K(F(x))) = F(x)$ for any $x \in \mathbb{R}^n$. Property ii) implies that $F(K(F(x)) + \delta) = F(x + \delta)$ for any $\delta \in \mathbb{R}^n$. Hence we obtain

$$g(F(x), \delta) = F(K(F(x)) + \delta) = F(x + \delta)$$

for any $x, \delta \in \mathbb{R}^n$. In terms of function f, with $p = \exp(x)$ and $\hat{p} = \exp(\delta)$ this implies:

$$f(p.\hat{p}) = g(f(p), \hat{p}).$$

Proof of iii) implies ii) If there is such a projection,

$$F(x_0) = F(x_1)$$

implies:

$$\pi(x_0) = \pi(x_1)$$

and as such:

$$\begin{aligned} F(x_0 + \delta) &= F(\pi(x_0 + \delta)) \\ &= F(\pi(x_0) + \pi(\delta)) \\ &= F(\pi(x_1) + \pi(\delta)) \\ &= F(\pi(x_1 + \delta)) \\ &= F(x_1 + \delta) \end{aligned}$$

Proof of ii) implies iii) To prove the converse property, first notice that each level set is a translation of any other one since for any shift δ, two points x_0 and x_1 are on the same level set if and only if $x_0 + \delta$ and $x_1 + \delta$ are on the same level set:

$$F(x_0) = F(x_1) \iff F(x_0 + \delta) = F(x_1 + \delta)$$

Hence we just need to describe the shape of a single level set to find the shape of all other ones. In the case where a level set is a point, all level sets are points and F is injective and property iii) is trivial; so for the remainder we will assume that level sets are not points.

Let's consider a function $\pi : \mathbb{R}^n \to \mathbb{R}^n$ such that $F(\pi(x)) = F(x)$ for all $x \in \mathbb{R}^n$. For any $x_0, x_1 \in \mathbb{R}^n$, $F(\pi(x_0)) = F(x_0)$ and property ii) imply:

$$F(\pi(x_0) + \pi(x_1)) = F(x_0 + \pi(x_1))$$

when we shift both sides by $\pi(x_1)$. Again using property ii) applied to $F(x_1) = F(\pi(x_1))$ and shifting by x_0, we obtain:

$$F(x_0 + \pi(x_1)) = F(x_0 + x_1)$$

Combining, we obtain:

$$F(\pi(x_0) + \pi(x_1)) = F(\pi(x_0 + x_1))$$

Similarly, as it implies that $F(2\pi(x)) = F(\pi(2x))$, we obtain:

$$F\left(\frac{\pi(x_0) + \pi(x_1)}{2}\right) = F\left(\pi\left(\frac{x_0 + x_1}{2}\right)\right)$$

If, in addition, F is injective on the image of π (i.e. π projects on at most a single point per level set), then we have

$$\pi\left(\frac{x_0+x_1}{2}\right)=\frac{\pi(x_0)+\pi(x_1)}{2} \tag{A.1}$$

for all x_0, x_1. For any F,we can construct such a projection π by chosing an arbitrary point on each level set.

Let us pick a point x_0 where the derivative of F has its maximal rank over a neighborhood of x_0. Assuming property ii), the derivative is the same on all points of the level set $\{x; F(x) = F(x_0)\}$ associated with point x_0. We can thus define an open set around x_0 that includes the level set $\{x; F(x) = F(x_0)\}$ and define a projection π that is continuous on that open set. Property A.1 then implies that π is linear on that set and thus that it is an affine set in $\mathbb{R}^n$.[2]

Since all level sets are translations of each other, all level sets are parallel affine sets of $\mathbb{R}^n$. The level set crossing the origin is then a linear subspace of $\mathbb{R}^n$. Denote by E its complement. E is crossing each level set only once, hence F is then injective on E. Denote by $\pi : \mathbb{R}^n \to E$ the projection of all points of a level set onto its intersection with E, we obtain that π is a linear function satisfying the conditions laid out in iii).

Examples and counter-examples Cobb-Douglas production functions provide an extreme example where we can combine the changes in output without even knowing the initial level of output (just knowing the functional form and the relative change in inputs). Level sets for Cobb-Douglas (in log) are planes and are thus affine as described above.

Next, consider expenditure shares across goods (depending on prices) when preferences are CES. Based on expenditure shares, we can identify relative prices up to a common constant. Knowledge of such relative prices is then sufficient to compute the change in expenditure shares depending on the change in prices, as it is well documented in the literature. In this case, level sets (in log) are all the lines parallel to the $(1, ..., 1)$ vector.

With Stone-Geary preferences exhibiting strictly positive minimum consumption requirements ϕ_i for each good i, expenditure shares are given by:

$$f_i(p/w) = \phi_i p_i/w + \alpha_i\left(1-\sum_j \phi_j p_j/w\right)$$

depending on normalized prices p_i/w. In this case, f is not conducive to exact hat algebra. For instance, if $n = 2$, $\phi_i = 1$ and $\alpha_i = 1/2$, we have:

$$f_1(p_1/w, p_2/w) \quad = \quad \frac{1}{2}\left[1 + p_1/w - p_2/w\right]$$

for $i = 1, 2$. We can see that $f_1 = f_2 = 1/2$ implies $p_1/w = p_2/w$, but we cannot identify

[2]Note that we cannot have a disconnected level sets (e.g. the union of two affine subsets) as the average between any two points of that level sets is again in the level set.

its value. However, the overall level of $p_1/w = p_2/w$ matters for the counterfactual outcome $f(\hat{p}_1 p_1/w, \hat{p}_2 p_2/w)$ as soon as $\hat{p}_2 \neq \hat{p}_1$. The same issue arises even if we consider expenditures instead of expenditure shares as observable outcome.

To fix this issue, a solution is to assume that one good (manufacturing good, say good $i = 1$) does not have a minimum consumption requirement, i.e. $\phi_1 = 0$, such that:

$$f_1(x) = \alpha_1 \left(1 - \sum_{j \neq 1} \phi_j p_j / w\right)$$

for manufacturing and:

$$f_i(x) = \phi_i p_i / w + \alpha_i \left(1 - \sum_{j \neq 1} \phi_j p_j / w\right)$$

for other goods. Function f is now invertible up to x_1, noticing that x_1 does not influence any expenditure share, and is now conducive to exact hat algebra. Note that other counter-examples can be found for homogeneous (homothetic) functions.

Appendix 4.D Price Discovery

In this subsection, we show that, in the case with only iceberg trade costs (i.e., $t_{od,g} = 0$ for all o, d, g), the price discovery step described in Section 2 is well defined in the sense that there is a unique set of prices $\{p_{j,g}\}$ that solves the system of equations (12)-(13) (for a given set of Foreign prices) and excess demand functions in Appendix 4.A. To do so, we think of that system of equations as characterizing the equilibrium of a competitive exchange economy, and so the goal is to prove that this economy has a unique equilibrium.

We consider an equivalent economy where there is a single market with an expanded set of goods (which we now call varieties) given by

$$\mathcal{V} \equiv \{(o, g) \in \mathcal{J} \times \mathcal{K}_A \cup \{L\} \mid q_{o,g} > 0\},$$

where $\mathcal{J}$ is the set of all agents excluding Foreign. A variety of good g produced by agent o is indexed by $(o, g) \in \mathcal{J} \times \mathcal{K}_A \cup \{L\}$. Agent $o's$ endowment of (o, g) is $q_{o,g}$. Naturally, no other agent $o' \neq o$ has a positive endowment of (o, g) and so $q_{o,g}$ is also the total endowment of variety (o, g) in the economy.

Letting $p_{o,g}$ denote the price of variety $(o, g) \in \mathcal{V}$, the price at which agent d has access to variety (o, g) is then $\tau_{od,g} p_{o,g}$. Letting $\boldsymbol{p} \equiv \{p_{o,g}\}_{(o,g) \in \mathcal{V}}$, the excess demand function (in value) for a variety $(o, g) \in \mathcal{V}$ is given by

$$\chi_{o,g}(\boldsymbol{p}) = \sum_{d \in \mathcal{J} \cup \{F\}} X_{d,o,g}(\boldsymbol{p}) - p_{o,g} q_{o,g},$$

where $X_{d,o,g}(\bullet)$ is the expenditure of agent d on variety (o, g). For $d \in \mathcal{J}$, and letting $\xi_{d,g} \in [0, 1]$

denote the expenditure share of gross income of agent $d \in \mathcal{J}$ (i.e., $\sum_g p_{d,g} q_{d,g}$) on good g,[3] we have

$$X_{d,o,g}(\boldsymbol{p}) \in \begin{cases} [0, \xi_{d,g} I_d] & \text{if } o \in \arg\min_{o' \in \mathcal{J} \cup \mathcal{F}} p_{o',g} \tau_{o'd,g} \\ 0 & \text{if } o \notin \arg\min_{o' \in \mathcal{J} \cup \mathcal{F}} p_{o',g} \tau_{o'd,g} \end{cases}.$$

$$I_d = \sum_g p_{d,g} q_{d,g},$$

In turn, for $d = F$ we have

$$X_{F,o,g}(\boldsymbol{p}) \in \begin{cases} 0 & \text{if } p_{o,g} > p^*_{F,g} \\ [0, \infty[& \text{if } p_{o,g} = p^*_{F,g} \\ \infty & \text{if } p_{o,g} < p^*_{F,g} \end{cases}.$$

We henceforth follow the convention that $q_{o,g} = 0 \implies p_{o,g} = \infty$ and $X_{d,o,g}(\boldsymbol{p}) = 0$, and also let

$$X_F(\boldsymbol{p}) \equiv \sum_{d \in \mathcal{J}, g} X_{d,F,g}(\boldsymbol{p})$$

denote the aggregate expenditure on goods from Foreign (imports).

The equilibrium is a set of prices $\boldsymbol{p}$ such that the excess demand (in value) for all varieties in $\mathcal{V}$ is zero,

$$\chi_{o,g}(\boldsymbol{p}) = 0, \quad \forall (o, g) \in \mathcal{V}. \tag{A.2}$$

We further assume that each agent $j \in \mathcal{J}$ produces at least one good (to ensure positive income) and has a positive expenditure share on each good that it produces:

Assumption A1: Endowments and demand.

1. $\sum_{g \in \mathcal{K}} q_{o,g} > 0, \quad \forall o \in \mathcal{J}$.

2. $q_{o,g} > 0 \implies \xi_{o,g} > 0, \ \forall o \in \mathcal{J}, g \in \mathcal{K}_A \cup \{L\}$.

For future purposes, note that the second part of this assumption implies that an increase in any price $p_{o,g'}, \ (o, g') \in \mathcal{V}$ leads to a strict increase in the value of excess demand $\chi_{o,g}(\boldsymbol{p})$ for any variety (o, g) with $\xi_{o,g} > 0$.

We say that a set of prices $\boldsymbol{p}$ is connected if there is only one trading block, i.e. there is no partition $\{\mathcal{J}_1, \mathcal{J}_2\}$ of $\mathcal{J}$ such that for all $g \in \mathcal{K}_A$ we have (i) $X_{d,o,g}(\boldsymbol{p}) = X_{o,d,g}(\boldsymbol{p}) = 0, \ \forall o \in \mathcal{J}_1, d \in \mathcal{J}_2$ (i.e., no trade between the two blocks) and (ii) $X_{F,o,g}(\boldsymbol{p}) = 0, \ \forall o \in \mathcal{J}_1$ or $X_{F,o,g}(\boldsymbol{p}) = 0, \ \forall o \in \mathcal{J}_2$

[3]Recall that the set of goods includes labor and crops. Gross income for a household is composed of the value of endowment of crops plus labor income. Subtracting the cost of intermediate goods (which are not included in the set of goods because prices are exogenous) and labor (as an input) yields disposable income, which is spent on consumption goods.

(i.e., it is not the case that both trade blocks trade with Foreign). Given Assumption A1, we now show that there can be at most one connected $\boldsymbol{p}$ that solves the system of equations A.2. We do so by appealing to the result in Corollary 1 of Berry *et al.* (2013) – henceforth BGH – which states sufficient conditions under which a function is injective on a set. Applying this result to our excess demand function $\{\chi_{o,g}(\boldsymbol{p})\}_{o,g}$ over the set of connected $\boldsymbol{p}$, we then get our desired result.

To apply the results of BGH we need to define "good 0," which is critical for the concept of "connected substitutes." We do this by considering each variety $(o,g) \in \mathcal{V}$ as a regular good and by thinking of the value of imports, $X_F(\boldsymbol{p})$, as the "demand for good 0." Trade balance then implies that

$$X_F(\boldsymbol{p}) = -\sum_{o,g} \chi_{o,g}(\boldsymbol{p}),$$

as in equation (2) of BGH.[4] We next show that Assumptions 1-3 in Corollary 1 of BGH are satisfied in our setting.

Translated to our context and notation, Assumption 1 in BGH states that the set of possible prices $\mathcal{P}$ is a Cartesian product.[5] This is immediately satisfied since $X_{o,g}(\boldsymbol{p})$ is satisfied for all prices $\boldsymbol{p}$ with $p_{o,g} \in [0,\infty[$.

Given that expenditure shares in demand are fixed and that higher prices lead to higher income (weakly), it is then easy to verify that import demand, $X_F(\boldsymbol{p})$, increases weakly with the price of any domestic variety in $\mathcal{V}$ while demand for variety (o,g), $\chi_{o,g}(\boldsymbol{p})$, increases weakly with the price of any other variety $(o',g') \in \mathcal{V}$ with $(o',g') \neq (o,g)$. This shows that varieties in our context are weak substitutes, and hence Assumption 2 in BGH is satisfied.

To verify that Assumption 3 in BGH is satisfied, we use the equivalent condition stated in BGH's Lemma 1. Translated to our context, this condition states that for any nonempty subset $\mathcal{V}_0$ of $\mathcal{V}$ either (i) there is a variety $(o,g) \in \mathcal{V}_0$ such that $X_F(\boldsymbol{p})$ increases strictly in $p_{o,g}$ or (ii) there is a variety $(o',g') \in \mathcal{V}\backslash\mathcal{V}_0$ such that $\chi_{o',g'}(\boldsymbol{p})$ increases strictly in $p_{o,g}$. We now show that this condition is satisfied by considering the three possible cases.

First, if there is an agent o and two goods g and g' such that $(o,g) \in \mathcal{V}_0$ and $(o,g') \in \mathcal{V}\backslash\mathcal{V}_0$ then an increase in $p_{o,g}$ leads to an increase in revenues for agent o and an increase in demand for variety (o,g') through an income effect given our Assumption A1.

Second, suppose that for any agent o either all or none of the varieties are in $\mathcal{V}_0$ (otherwise we are back to case one just above). Suppose also that there is a variety $(o,g) \in \mathcal{V}_0$ and a variety $(o',g') \in \mathcal{V}\backslash\mathcal{V}_0$ such that agent o purchases good g' from o', i.e. such that $X_{o,o',g}(\boldsymbol{p}) > 0$. In that case, an increase in the price $p_{o,g}$ leads to an increase in revenues for agent o and an increase in demand for variety (o',g') again through an income effect.

Finally, the third case is one where, for any agent o, either all or none of the varieties are in $\mathcal{V}_0$,

---

[4] BGH add $+1$ to demand for good "0," but this does not affect any results nor assumptions on monotonicity.

[5] Here we look at prices, thus reversing all signs of the slopes in BGH, who focus instead on demand shifters (denoted with x). Our set $\mathcal{P}$ corresponds to the set $\mathcal{X}$ in BGH, while the set of all connected prices $\mathcal{P}^* \in \mathcal{P}$ corresponds to $\mathcal{X}^* \subset \mathcal{X}$ in BGH.

and where no agent o purchases goods from agents that have varieties outside $\mathcal{V}_0$. As we focus on connected price vectors, this implies that there is non-zero demand for some Foreign good by some agent o that has some varieties (o, g) in $\mathcal{V}_0$. As such, an increase in the price $p_{o,g}$ leads to an increase in the demand for Foreign goods $X_F(\boldsymbol{p})$.

Appendix 4.E Recovering Trade Shares in Manufacturing

In Section 2, we lay out our solution method when available data include expenditure shares $\xi_{j,g(h)}$ for manufacturing goods for all $h \in \mathcal{H}$ and agents $j \in \mathcal{I} \cup \mathcal{H}$. As in our case, such data are not always available at such level of aggregation. Here we provide details on how to recover expenditure shares $\xi_{j,g(h)}$ following a method similar to Donaldson & Hornbeck (2016) and Faber & Gaubert (2019). We assume that we can obtain aggregate information on international trade deficit in manufacturing.

First, we need to separately infer aggregate imports and aggregate exports of manufacturing with Foreign. Given income levels of farmers (inferred along with agricultural crop prices) and urban households in Home (observed), we can compute overall expenditures on manufacturing by each agent in Home as $I_j \cdot (1 - \sum_{k \in \mathcal{K}_A} \xi_{j,k})$ for $j \in \mathcal{I} \cup \mathcal{H}$. Total revenues in manufacturing in Home are $\sum_h I_h$, and the difference between total expenditures and revenues in manufacturing gives us Home's overall deficit in manufacturing. Assuming that we can observe (e.g. from international trade data) the ratio of this deficit to Home's manufacturing imports, we can then deduce the value of manufacturing imports by Home, $\sum_{j \in \mathcal{I} \cup \mathcal{H}} X_{j,g(F)}$, as well as its manufacturing exports to Foreign, $\sum_{h \in \mathcal{H}} X_{F,g(h)}$.

Next, we assume that the demand shifter in manufacturing (which can be interpreted as quality or productivity) may vary across sources (urban households and Foreign) but is not specific to each destination, i.e. $b_{j,g(h)} = b_{M,g(h)}, \forall j \in \mathcal{I} \cup \mathcal{H} \cup \{F\}$ and $\forall h \in \mathcal{H} \cup \{F\}$. Excess demand for the manufacturing good of urban household h satisfies:

$$\sum_{j \in \mathcal{J}} \chi_{j,g(h)}(\{b_{M,k} p_{j,k}\}_j, I_j) = 0.$$

In this expression, note again that we can simplify the arguments of function $\chi_{j,g(h)}$ since the demand for a manufacturing good does not depend on land rents once we know farmer income (and their expenditure share in manufacturing). For the manufacturing good produced in Foreign, we have

$$\sum_{j \in \mathcal{I} \cup \mathcal{H}} \chi_{j,g(F)}(\{b_{M,k} p_{j,k}\}_j, I_j) = \sum_{j \in \mathcal{I} \cup \mathcal{H}} X_{j,g(F)}$$

where the right-hand side is observed or inferred as discussed above. Combined with $p_{j,g(h)} = \tau_{hj,g(h)} p_{h,g(h)}$ for $h \in \mathcal{H}$ and $p_{j,g(F)} = \tau_{Fj,g(F)} p_{F,g(F)}$, the previous displayed equations constitute a system of equations in $b_{M,g(h)} p_{h,g(h)}$ for $h \in \mathcal{H}$ and $b_{M,g(F)} p_{F,g(F)}$, which has a unique solution as long as demand features gross substitutes, as is the case in most of the trade literature (e.g., with CES demand). Given the solution in $b_{M,g(h)} p_{h,g(h)}$ (up to a common constant), we can recover expenditure shares $\xi_{j,k}$ for each agent $j \in \mathcal{I} \cup \mathcal{H}$ and each manufacturing variety $k \in \mathcal{K}_M$.

Appendix 4.F Hub-and-Spoke Trade Costs

In this subsection, we want to show that condition (3) leads to well defined market prices once we make the hub-and-spoke assumption on trade costs in expressions (14) and (15). To simplify notation we ignore the subindex for g and focus on one particular agriculture good. Since we are assuming away iceberg trade costs, then (3) entails

$$p_o + t_{od} \geq p_d \perp x_{od}. \tag{A.3}$$

We define the market price associated with a farmer $i \in \mathcal{J}(m)$ by

$$p_m(i, sells) \equiv p_i + t_{im}$$

if the farmer is a seller of the good and by

$$p_m(i, buys) \equiv p_i - t_{mi}$$

if the farmer is a buyer of the good. Consider three farmers i_1, i_2 and i_3 connected to market m (i.e., $i_1, i_2, i_3 \in \mathcal{J}(m)$), and assume that i_1 and i_2 are sellers and i_3 is a buyer. We first show that $p_m(i_1, sells) = p_m(i_2, sells)$ and then show that $p_m(i_1, sells) = p_m(i_3, buys)$, implying that there is a well defined market price p_m.

To prove $p_m(i_1, sells) = p_m(i_2, sells)$, assume by contradiction that $p_m(i_1, sells) \neq p_m(i_2, sells)$. This would imply that

$$p_{i_1} + t_{i_1 m} \neq p_{i_2} + t_{i_2 m}.$$

Without loss of generality, assume that

$$p_{i_1} + t_{i_1 m} < p_{i_2} + t_{i_2 m}.$$

Let j be the agent that buys the good from farmer i_2, and let t_{mj} be the trade cost from market m to agent j. Combining this with (A.3) (which holds with equality for j and i_2) we get

$$p_{i_1} + t_{i_1 m} + t_{mj} < p_{i_2} + t_{i_2 m} + t_{mj} = p_j,$$

which indicates that j could instead buy the same good from i_1 at a lower price, contradicting condition (A.3) for j and i_1, which implies

$$p_{i_1} + t_{i_1 m} + t_{mj} \geq p_j.$$

To prove $p_m(i_1, sells) = p_m(i_3, buys)$, assume by contradiction that $p_m(i_1, sells) \neq p_m(i_3, buys)$.

Assume first that $p_m(i_1, sells) < p_m(i_3, buys)$. This implies

$$p_{i_1} + t_{i_1 m} < p_{i_3} - t_{m i_3},$$

which is a contradiction because (A.3) implies

$$p_{i_1} + t_{i_1 m} + t_{m i_3} \geq p_{i_3}.$$

In words, i_3 could instead buy the good from i_1 at a lower price. Now assume instead that $p_m(i_1, sells) > p_m(i_3, buys)$ and let $j_1 \in \mathcal{J}(m')$ be the agent that is buying the good from i_1 and let $j_3 \in \mathcal{J}(m'')$ be the agent that is selling to i_3, with markets m, m' and m'' possibly but not necessarily coinciding. We again reach a contradiction as j_1 could instead buy the good from j_3 at a lower price. To see this, note that

$$p_{j_1} = p_{i_1} + t_{i_1 m} + t_{m m'} + t_{m' j_1}$$

while

$$p_{j_3} = p_{i_3} - t_{j_3 m''} - t_{m'' m} - t_{m i_3}.$$

Combined with $p_m(i_1, sells) > p_m(i_3, buys)$, these two equations imply

$$p_{j_3} + t_{j_3 m''} + t_{m'' m} + t_{m m'} + t_{m' j_1} < p_{j1}.$$

The triangular inequality implies

$$p_{j_3} + t_{j_3 j_1} \leq p_{j_3} + t_{j_3 m''} + t_{m'' m} + t_{m m'} + t_{m' j_1} < p_{j_1},$$

which violates (A.3).

Appendix 4.G Model Extension with Seasonal Migration

In this appendix, we extend the model to allow for seasonal migration between rural markets as well as between rural and urban markets. As for trade in goods, labor can be traded between any two local labor markets subject to additive trade costs $t_{od,L}$ and/or iceberg trade costs $\tau_{od,L}$. We refer to this trade in labor as "seasonal migration", since we assume that migrants consume (and face prices) at their home location but earn wage $p_{i,L}$ on destination farm i, or $p_{h,L}$ when working for urban household h. We do not allow for international migration, i.e., $t_{od,L} = \tau_{od,L} = \infty$ for $o, d \in \{F\}$.

Our model exposition in Section 2 and the general functional forms in Appendix 4.B continue to apply to the model with migration. However, since labor supply is no longer perfectly inelastic in urban markets due to migration, we cannot treat output of manufacturing good $g(h)$ as an endowment. Instead, output of manufacturing variety $g(h)$ is given by:

$$q_h = a_h \sum_o x_{oh,L},$$

where $x_{oh,L}$ are flows of labor from any origin o to urban household h, and a_h is a productivity shifter. As for wages in rural markets, we need to account for wages $p_{h,L}$ that clear urban labor markets in equilibrium. Due to perfectly competitive labor markets, the urban wage follows $p_{h,L} = p_{h,g(h)} a_h$, where $p_{h,g(h)}$ is the price of manufacturing variety $g(h)$ for urban household (or city) h.

In equilibrium, rural and urban households maximize utility taking prices as given, prices respect no-arbitrage conditions given trade costs, and all markets clear. The equilibrium is a set of prices, $\{p_{j,g}\}$ and trade flows $\{x_{od,g}\}$ (measured in quantity at the destination). The equilibrium conditions (2)-(4), laid out in Section 2, apply to the model with migration as well. Based on the discussion above, only equilibrium condition (5) for urban income changes to:[6]

$$I_h = p_{h,L} L_h, \quad \forall h \in \mathcal{H}.$$

[6] Here, we now need to make the distinction between urban labor endowment L_h and urban employment $\sum_o x_{oh,L}$.