

Scaling Agricultural Policy Interventions*

Lauren Falcao Bergquist[†], Benjamin Faber[‡], Thibault Fally[§]

Matthias Hoelzlein[¶], Edward Miguel^{||}, Andres Rodriguez-Clare^{**}

May 2026

Policies to raise agricultural productivity have been central in the fight against global poverty. Their impacts are often measured in experiments that provide strong causal identification but may be too small-scale to capture general equilibrium effects that arise when policies are scaled up. We develop a quantitative model of agricultural trade, featuring a granular economic geography, that combines parameters estimated from small-scale experiments with rich administrative microdata to quantify how treatment effects change when policies are scaled up. Applying this framework to input subsidies in Uganda, we find that average welfare gains fall by about 20% when implemented at scale. However, gains increase among the poorest households, as returns shift from land to labor, reducing the regressivity of the intervention by more than half. We examine how these forces vary with household and market characteristics and the geographic scale of implementation, with implications for randomized saturation designs.

JEL classification: F15, F63, O13

*This research was supported by funding from the International Growth Centre (IGC). We thank Peter Flo Grinde-Hollevik, Liz Kairima, Felipe Prates Tavares, and Camilla Sarzo for excellent research assistance. We thank audiences at multiple seminars for comments and Rodrigo Adao, David Atkin, Wyatt Brooks, Arnaud Costinot, Dave Donaldson, Kevin Donovan, Pablo Fajgelbaum, Doug Gollin, Paul Niehaus 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](#)). 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. To inform these policies using rigorous evidence, much of the recent literature uses randomized control trials (RCTs) or natural experiments to generate exogenous variation in policy exposure across households or local markets. 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 to a broader segment of the population.¹

For example, a small-scale RCT of a fertilizer subsidy program may find that farmers use more fertilizer, see higher yields, and enjoy higher revenues as a result. However, when this program is implemented at scale, output prices may fall, input prices may rise, and consumption bundles may shift. Though it is understood in a qualitative sense that this will alter treatment effects, quantification and analysis of such differences have largely remained outside the scope of policy evaluations in this space.² How far off are small-scale experimental estimates from the ultimate full-scale policy impacts they are designed to inform? To what extent do GE effects alter the distributional implications at scale? For which types of markets and households are GE forces most pronounced in shifting impacts at scale?

To shed light on these questions, we develop a theoretical framework and practical toolkit for researchers conducting small-scale agricultural experiments who wish to quantify the implied impacts of these policies at scale. We develop a rich but tractable model for quantifying GE policy counterfactuals at the level of households in agricultural settings, which

¹There may be other important differences that emerge when policies are scaled, including variation in implementation details and questions of external validity ([Muralidharan and Niehaus, 2017](#)). While the primary focus of this paper is on differences driven by general equilibrium effects, we also explore extensions of our model to allow for differences in the intervention’s effectiveness at scale (Section 7). To the extent that treatment effects vary across populations due to fundamentals and forces present in our model, our approach can also speak to questions of external validity.

²An earlier literature in agriculture and development used computable general equilibrium (CGE) analysis to quantify GE implications. However, these studies often rely on less well-identified moments for parameter estimation and largely abstract from modeling the granular economic geography of farm production, consumption and trade costs that underlies the propagation of shocks and their incidence in GE. See e.g. [de Janvry and Sadoulet \(1995\)](#) for a review of this literature.

we combine with estimates from local experiments and rich administrative microdata. We then evaluate the local versus at-scale implications of one of the most widespread and costly agricultural policies in low-income countries: a subsidy to modern inputs, including chemical fertilizers and hybrid seed varieties. We do so in the context of Uganda, quantifying how the average treatment effect and distributional implications of the subsidy differ, for the same group of households, if the policy is scaled up from a small-scale pilot to the national level.

To achieve this, we must overcome two challenges that have limited the extent to which quantitative models in trade and spatial economics can speak to the results from small-scale agricultural experiments in low-income countries. The first is that models of a national or regional economy frequently feature a level of aggregation that is much higher than the individual household, the common unit of analysis in small-scale experiments. The second is that these models typically abstract from some features of the agricultural setting that are critical mediators for the propagation of shocks in GE. Motivated by cross-country trade in manufacturing, a common assumption is differentiation of goods by origin in a setting where all markets trade with each other, precluding a policy shock from changing which markets are connected through trade.³ Further, trade costs are typically assumed to be ad valorem (“iceberg”), implying complete pass-through of local price changes across trading pairs, an assumption that is not supported empirically in our setting. To bridge these gaps, our theory features individual households that populate the economy, where farmers trade in homogenous crops within and across markets subject to potentially both additive (per-unit) and multiplicative (ad valorem) trade frictions.⁴

These features are fundamental for modeling a granular and realistic economic geography that underlies the propagation of shocks across markets and households when an agricultural policy is scaled up to a national level. However, as we show, incorporating these features disrupts the convenient properties of “structural gravity,” including the use of “exact hat algebra” as the conventional solution method for counterfactual analysis in the liter-

³Notable exceptions are [Costinot and Donaldson \(2016\)](#) and [Sotelo \(2020\)](#). Since the additional data proposed to solve these models (on either production possibility frontiers or farm-gate prices) are rarely available, especially at the household-level, we propose a new solution method that unlocks the scope for counterfactual analysis in this environment.

⁴In addition to these features, we also allow for non-homothetic preferences, so that food price changes can have distributional implications beyond affecting household incomes, and technology choice in crop production similar to [Farrokhi and Pellegrina \(2022\)](#), such that the adoption of modern inputs can more flexibly affect the production function with respect to other inputs.

ature (e.g., [Costinot and Rodríguez-Clare \(2014\)](#)). After laying out the model, we propose a new quantitative approach that leverages rich but widely available microdata on household location, production and consumption. We use data on trade costs between and within markets, together with household expenditure shares and agricultural production quantities, to formulate a price discovery problem that recovers equilibrium farm-gate prices and trade flows consistent with observed behavior and the underlying trade-cost network. Given these objects, we then combine exact hat algebra with mixed-complementarity programming to solve for counterfactual equilibria. This approach has several advantages. First, it allows us to solve the model without imposing structural gravity, while avoiding strong data requirements such as observing the full set of initial farm-gate prices. It also allows the set of active trading pairs and the direction of trade to adjust in counterfactuals, capturing extensive-margin responses that are typically ruled out by assuming a fixed trade network. Second, from a computational perspective, the solution method is capable of handling high-dimensional GE counterfactuals at the level of individual households.

We then describe the use of small-scale experiments and administrative microdata to quantify the model. To estimate the key demand and supply elasticities, we use exogenous variation in consumer and producer prices from RCTs ([Bergquist and Dinerstein, 2020](#), [Carter et al., 2020](#)). On the supply side, we also make use of a natural experiment that exploits variation in how changes in crops' world market prices propagate to local markets as a function of trade costs to the nearest border crossing. To calibrate trade costs, we make use of estimates from [Bergquist et al. \(2024\)](#), using Ugandan market and trader survey microdata to provide information on market-to-market trade flows and crop prices at origin and destination. 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 use the quantified model to estimate the local and at-scale effects of a 75% cost subsidy for modern inputs in Uganda, including chemical fertilizer and hybrid seeds. For each of the roughly 4,500 rural parishes in Uganda, we randomly select 2.5 percent of the local population (a sample of roughly 100,000 households nationwide). We estimate two counterfactuals for these households. First, we solve for the household's "local effect:" the counterfactual changes in household welfare (real income in our model) due to an intervention that targets the subsidy only at the randomly selected households in a single parish,

keeping the rest of Uganda unexposed – akin to implementing roughly 4,500 separate RCTs. Second, we estimate the household’s “at-scale effect:” the welfare change experienced by the same household under an intervention that scales the subsidy to all rural households in Uganda. We find that the average effect of the subsidy at small scale, pooling all local randomized interventions, is a 4.3 percent increase in household real income. This is driven 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 changes by as much as + or -5 percentage points across households. 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 (a 3.5 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.

We document important distributional implications underlying these comparisons. In the local intervention, 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 the local subsidy increases returns to land without discernible effects on output or input prices, income gains are concentrated among land-rich households – who, in addition, also use modern inputs more intensively at baseline. At scale, however, GE effects on average decrease the local producer price index and increase the price of local labor. The resulting reduction in land revenues and increase in labor compensation benefit households with higher initial reliance on wage labor relative to the land-rich. GE forces therefore reduce the regressivity of the policy.

After documenting these changes in impact at scale, we use our framework to provide additional insights for research at the intersection of agriculture and development. A growing literature employs “randomized saturation designs” that randomize not only treatment across individuals, but also the saturation rate across geographic areas (“clusters”) (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.⁵ To identify the impact of policies at scale, one must thus typically extrapolate from these discrete points of saturation, subject to two assumptions: that GE forces are linear with respect to changes in the saturation rate; and 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 assess these assumptions by exploring how welfare implications evolve as a function of saturation rates, and at different geographical scales. At the nationwide level, we start with the local intervention that treats 2.5% of farmers in each parish, and then estimate how treatment effects among that original sample of farmers evolve as the program is sequentially scaled up in steps of 10% of the remaining rural Ugandan population. We find that the average gains decline close to linearly as a function of scale-up to the rest of the country, providing some reassurance about the lessons that can be drawn from designs relying on two discrete saturation rates.

However, our results also suggest some caution about these designs. Experiments typically randomize saturation 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, increases in saturation in steps of 10% of the population within subcounties (a larger than typical but feasible unit for randomized saturation) lead to almost no change in the average welfare gains, even at 100% saturation within the subcounty, and different distributional implications compared to the intervention at national scale. We document that this is not due to the absence of GE forces under subcounty saturation, but rather due to their different nature compared to at national scale. These findings suggest some caution when extrapolating from GE effects observed within smaller geographic units to the effects at a broader scale of program rollout.

We proceed to conduct variance decompositions to identify which household and market characteristics are the key drivers of changes in the policy's impact at scale relative to the local intervention. We find that initial household land income shares, which are at the center of the distributional changes discussed above, are also the single most important mediating factor explaining the extent of changes in the effect at scale, followed by initial revenue shares of maize – the main staple crop in our setting and which also has a relatively high average use of modern inputs at baseline – and the degree of market remoteness. We then use these insights to understand how well the current body of agricultural RCTs map to

⁵Variation in saturation can also stem from unit-level randomization and their proximity in space. See, e.g., [Egger et al. \(2022\)](#) who combine both types of variation.

their likely impacts at scale. We conduct a meta-analysis of registered agricultural RCTs in Sub-Saharan Africa over recent decades, documenting the prevalence of the key drivers of differences at scale across the study sites and overlaying the predicted changes in treatment effects at scale. We also report computations related to publication bias and show that findings from local experiments could affect learning about policy impacts disparately among land-poor vs. land-rich households: namely, under-stating effectiveness for the poor, but over-stating for the rich.

Finally, we present additional sensitivity and model validation exercises. These highlight the important role that RCTs and natural experiments play in identifying impacts in a given policy environment. They also point to some important limitations of our approach and possible extensions to address some of these in other policy contexts. For example, our baseline model does not allow the direct effect of the subsidy on yields to vary with the scale of implementation. However, there can be scenarios in which these effects either increase or decrease as a function of the scale of policy implementation.⁶ We follow recent advances from modeling agglomeration and congestion effects in spatial economics to extend our framework and then document the implications for changes in impact at scale in such scenarios.

We proceed as follows. Section 2 lays out the model and solution method. Section 3 describes the use of local experiments to estimate the key supply and demand parameters. Section 4 describes the use of administrative microdata to calibrate a granular economic geography. Section 5 presents the counterfactual analysis comparing local and at-scale policy impacts. Section 6 uses the framework to derive additional insights for research. Section 7 presents sensitivity and validation analyses and concludes with a discussion of the limitations and potential extensions of our approach.

2 Model and Solution Method

We present the model in three layers. First, we introduce the farmers' decisions about consumption and production, taking all prices as given. Second, we describe how farmers trade in village markets and how farm-gate prices are determined, taking village market prices as given. Lastly, we study how village prices are determined through trade with other village

⁶Complementarities at scale could arise from learning from others. Conversely, changes in the implementation protocols or oversight when scaling up could result in lower uptake or usage (Muralidharan and Niehaus, 2017).

markets, cities, and the rest of the world. After laying out the model, we develop the solution method to quantify policy counterfactuals given a rich but realistic data environment. Throughout, we provide intuition and discuss comparative statics with a focus on our application to policies subsidizing the use of modern inputs. We provide proofs and model extensions in [Appendix 5](#). To sharpen the exposition, we present the model here using the same functional forms for preferences and technology that we use in the counterfactual analysis, while the [Technical Appendix](#) provides a more general version that can be adapted to other contexts.

2.1 Farmer Decisions

Farmers consume a bundle of homogeneous crops and a bundle of differentiated manufacturing products. At the farm-gate, price p_k refers to good k indexing crops ($k \in \mathcal{K}^A$) or manufacturing products ($k \in \mathcal{K}^M$).⁷ Farmers have nested preferences. At the upper nest, farmers decide how much to consume of the crop and manufacturing bundles subject to non-homothetic Stone-Geary preferences, so that farmers must consume a minimum amount of the crop bundle \bar{C}_A for utility to be positive. At the lower nest, both the crop bundle and the manufacturing bundle are CES-aggregates with elasticities of substitution σ (crop bundle) and η (manufacturing bundle). The expenditure share on good k is then

$$\xi_k = \begin{cases} \frac{(b_k p_k)^{1-\sigma}}{P_A^{1-\sigma}} * \left(\zeta + (1 - \zeta) \frac{P_A \bar{C}_A}{I} \right) & \text{for } k \in \mathcal{K}^A \\ \frac{(b_k p_k)^{1-\eta}}{P_M^{1-\eta}} * (1 - \zeta) \left(1 - \frac{P_A \bar{C}_A}{I} \right) & \text{for } k \in \mathcal{K}^M \end{cases},$$

where I is income, $P_A = \left(\sum_{k \in \mathcal{K}^A} (b_k p_k)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$ is the price index of the crop bundle, and $P_M = \left(\sum_{k \in \mathcal{K}^M} (b_k p_k)^{1-\eta} \right)^{\frac{1}{1-\eta}}$ is the price index of the manufacturing bundle.⁸ The term $\zeta + (1 - \zeta) \frac{P_A \bar{C}_A}{I}$ is the share of expenditure allocated to the crop bundle, while $\frac{(b_k p_k)^{1-\sigma}}{P_A^{1-\sigma}}$ is the share of that expenditure allocated to crop k . Similarly, $(1 - \zeta) \left(1 - \frac{P_A \bar{C}_A}{I} \right)$ is the share of expenditure allocated to the manufacturing bundle and $\frac{(b_k p_k)^{1-\eta}}{P_M^{1-\eta}}$ is the share of that expenditure allocated to manufacturing product k . The term b_k is a demand shifter for good k . Holding prices constant, as income increases farmers reallocate their expenditure from the crop bundle to the manufacturing bundle. As income grows very large, preferences approach the Cobb-Douglas form with share $\zeta \in (0, 1)$ allocated to the crop bundle.

⁷We show below how farm-gate prices are determined in equilibrium but we can understand p_k as the shadow price of good k faced by a farmer even if it is not traded (i.e., under subsistence farming).

⁸We henceforth suppress the notation $k \in \mathcal{K}^A$ or \mathcal{K}^M unless needed to avoid confusion.

On the production side, farmers grow a bundle of crops using land, labor, and intermediate inputs. When producing crop k a farmer can choose between two techniques indexed by ω : a traditional technique using traditional inputs (land and labor), and a modern technique that additionally requires an intermediate input such as chemical fertilizer or hybrid seeds. The production function is Cobb-Douglas with crop-by-technique-specific cost shares $\alpha_{k,\omega}$ for intermediates, $\beta_{k,\omega}$ for labor, and $\delta_{k,\omega} = 1 - \alpha_{k,\omega} - \beta_{k,\omega}$ for land.

Letting p_n and p_l denote a farmer's shadow prices for the intermediate input and labor, respectively, the shadow return for an effective unit of land allocated to crop k with technique ω , denoted as $r_{k,\omega}$, is determined from the condition that price equals unit cost,

$$p_k = a_{k,\omega}^{-1} r_{k,\omega}^{\delta_{k,\omega}} p_n^{\alpha_{k,\omega}} p_l^{\beta_{k,\omega}}, \quad (1)$$

where $a_{k,\omega}$ is a productivity shifter. A farmer's optimization problem entails maximizing total land returns, $\sum_{k,\omega} r_{k,\omega} Z_{k,\omega}$, by choosing the effective units of land cultivated with crop k and technique ω , $Z_{k,\omega}$ for all crops and techniques. We assume that the marginal productivity of *physical* units of land allocated to a crop and technique is decreasing, so that the marginal *effective* unit of land falls with a larger physical plot allocated to a crop-technique. In particular, the farmer's optimization problem is constrained by the following set of feasible allocations of effective units of land:

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

where Z is the farmer's endowment of physical land.⁹ Farmers can substitute between crops with elasticity μ in the upper nest and between techniques within a given crop choice with elasticity κ in the lower nest.¹⁰ The solution to the farmer's optimization problem is characterized by the share of land income coming from crop k and technique ω for all crops and techniques:

$$\pi_{k,\omega} \equiv \frac{r_{k,\omega} Z_{k,\omega}}{\sum_{k,\omega} r_{k,\omega} Z_{k,\omega}} = \frac{r_{k,\omega}^\kappa}{\sum_\omega r_{k,\omega}^\kappa} * \frac{(\sum_\omega r_{k,\omega}^\kappa)^{\mu/\kappa}}{\sum_k (\sum_\omega r_{k,\omega}^\kappa)^{\mu/\kappa}}.$$

⁹Motivated by evidence that formal land markets are thin in many rural parts of sub-Saharan Africa (Acam-pora et al. (2022); Holden et al. (2010)), we abstract from trade in land. Having said this, informal land rental markets exist in many contexts, and incorporating such exchanges could be an interesting model extension depending on both the type of intervention and empirical setting.

¹⁰This is a nested constant elasticity of transformation production function as in e.g. Powell and Gruen (1968). The same expression can be obtained from an extension of the Roy-Frechet microfoundations in Costinot and Donaldson (2016) and Sotelo (2020), but now allowing for a nested Frechet structure, as in Farrokhi and Pellegrina (2022).

The first term on the right captures the share of land income coming from technique ω within total returns from crop k , while the second term captures the share coming from crop k in the bundle of all produced crops. The land return per unit of physical land is

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

and physical output of crop k and technique ω is

$$q_{k,\omega} = \frac{\pi_{k,\omega} r Z}{\delta_{k,\omega} p_k}.$$

The numerator is the value of the output of this crop-technique accruing to land, and the land cost share in the denominator scales land value to the total value of that crop-technique. Dividing by the farm-gate price of crop k transforms value into a physical quantity. Lastly, farmer income is $I = p_l L + rZ$, where $p_l L$ is income from the labor endowment, and rZ is income from the land endowment.

Considering the effects of a subsidy for modern intermediate inputs, we can characterize farmer adjustments through the implied changes in land returns given the wage and the output price. First, holding the land allocation fixed, the subsidy increases land returns as long as the modern technique is used. Second, within the land allocated to a given crop, the subsidy induces farmers to shift land to the modern technique. Finally, as different crops may require more or less use of modern intermediates relative to land, captured by $\alpha_{k,\omega}/\delta_{k,\omega}$, land returns increase more for some crops than others, inducing farmers to allocate more land to intermediate-intensive crops. These changes in land allocation further increase land returns and farmer income.

2.2 Local Markets, Trade Costs and Price Determination

After introducing how output prices and wages affect farmers' consumption and production decisions in the previous section, we now discuss how these farm-gate prices are determined by farmers' trading behavior in the village market. We define a village as a group of farmers $J(m)$ who have access to a local market m where crops, manufacturing products and labor (all referred to as "goods" of type g below) are traded. We assume that all local trade is conducted on the village market. We index a farmer by subscript i , such that farmer i 's land endowment is Z_i , for example. Trade in good g from farmer i to market m is subject to trade costs which stand in for various frictions to trade across space, including of course physical transportation costs. Trade costs can have both an additive component, $t_{im,g}, t_{mi,g} > 0$,

and an ad-valorem (iceberg) component, $\tau_{im,g}, \tau_{mi,g} > 1$: $t_{im,g}$ ($t_{mi,g}$) is the trade cost for farmer i when selling (buying) a unit of good g to (from) market m , while $\tau_{im,g}$ ($\tau_{mi,g}$) is the number of units of good g that farmer i must ship (purchase) for one unit to arrive at the destination. Additive trade costs are costs per unit shipped (e.g., per sack of maize), while iceberg trade costs scale with the value of the goods.¹¹

There are three cases for how farm-gate prices and market prices are related depending on the direction of trade in a good. First, if farmer i buys good g from the market, she incurs trade costs from market to farm-gate, and her farm-gate price is $p_{i,g} = \tau_{mi,g}p_{m,g} + t_{mi,g}$. Second, if farmer i sells good g at the market, her farm-gate price is $p_{i,g} = (p_{m,g} - t_{im,g})/\tau_{im,g}$. Lastly, farmer i can decide not to trade good g on the market, neither as seller nor as buyer. She is then in autarky (or subsistence) in this specific good and her farm-gate price must lie between the seller and buyer prices, $p_{i,g} \in ((p_{m,g} - t_{im,g})/\tau_{im,g}, \tau_{mi,g}p_{m,g} + t_{mi,g})$. The specific (shadow) price of good g in this interval is determined by equalizing own production and consumption of good g in equilibrium.

To make the notation more compact, we use i and j to index “agents,” which include farmers and markets. Denoting the trade flow in quantity of good g from agent i to agent j with $x_{ij,g}$, we can then capture the previous three cases by

$$\tau_{ij,g}p_{i,g} + t_{ij,g} \geq p_{j,g} \perp x_{ij,g}, \quad \forall i, j.$$

The symbol \perp between a weak inequality and a variable indicates that the weak inequality holds as equality if the variable is strictly positive. Thus, if agent j buys g from i then $x_{ij,g} > 0$, implying that the price j faces is given by agent i 's price plus trade costs.¹² If $x_{ij,g} = x_{ji,g} = 0$, the agents are not trading good g and we only know that i 's and j 's prices must be within intervals determined by trade costs and each others' prices.

Before introducing trade across markets to close the full setup of the model below, we can already see a rich set of farmer adjustments (within markets for now) to policy shocks that our environment is able to capture. Given homogeneous goods and trading frictions, a subsidy for intermediate inputs can lead to changes in farmers' trading behavior. Suppose the farmer initially buys a crop k . Now, the subsidy lowers production costs. The farmer

¹¹Additive trade costs are in units of a “transportation good” (e.g., fuel). We assume that this good is only imported and that there are no trade costs for this good, so that all agents can access it at the same price. Setting this price to one by choice of numeraire, $t_{ij,g}$ becomes the additive trade cost from i to m for good g . We thus model trade costs as exogenous (unaffected by GE price adjustments).

¹² $x_{ij,g} > 0$ implies $x_{ji,g} = 0$, since otherwise the expression above implies $\tau_{ji,g}p_{j,g} + t_{ji,g} = p_{i,g}$, which cannot simultaneously hold with $\tau_{ij,g}p_{i,g} + t_{ij,g} = p_{j,g}$ if any of the trade costs between i and j are positive.

might stop buying and even start selling the crop. Similarly, if the farmer is initially in autarky (subsistence) for crop k , then she might start selling it. In this sense, our model is able to capture the extent to which policies shift farmers out of subsistence farming and into market sales. Furthermore, additive trade costs also change the degree to which a price shock passes through to farmers. Since the proportion of the farm-gate price that covers additive trade costs does not scale with the price, the pass-through of price shocks will be incomplete, leading to weaker effects of the subsidy for remote farmers. These features we describe here within markets will also carry over to adjustments in cross-market trading connections, to which we turn below.

2.3 Urban Centers, Foreign Markets and the Trade Network

We now introduce urban centers and the rest of the world, as well as how trade patterns connecting these and rural villages determine market-level prices. Given our focus on the effects of agricultural policies, we follow a simple but standard setup for the urban sector. Each urban center (city) is populated by a representative household h . Preferences remain as described above. On the production side, each urban household is endowed with labor units L_h and produces a differentiated manufacturing product $g(h)$. The manufacturing technology is linear in labor, so that the quantity of manufacturing product $g(h)$ produced by urban household h is $a_h L_h$, where a_h is a positive constant. Given that urban labor supply is fixed, we can treat $a_h L_h$ as the urban household's endowment of manufacturing product $g(h)$. The urban household's income is $I_h = p_{h,g(h)} a_h L_h$, where $p_{h,g(h)}$ is the price for urban household h of manufactured good $g(h)$ (i.e., its own manufacturing product variety).

The final type of agent we need to introduce is what we call Foreign, indexed by F , which stands for the rest of the world. Foreign produces manufacturing product $g(F)$, as well as crops and the intermediate input. We assume that our model economy, Home, is "small" relative to Foreign. Thus, agents in Home can buy or sell any amount of crop k at exogenous price $p_{F,k}^*$ at a border crossing with Foreign. Similarly, agents in Home can buy the Foreign manufacturing product at world-market price $p_{F,g(F)}^*$ at a border crossing with Foreign.¹³ In turn, Foreign's demand for the manufacturing good produced by urban household h is such that the value of exports of good $g(h)$ is

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

¹³An agent buying or selling a good from Foreign would also pay the trade costs between their place of residence and the border.

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

For intermediate inputs, we assume that they are imported from Foreign at exogenous prices $p_{i,n}$ that farmer i pays for the input. This provides the 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

We allow for trade across markets on a fully connected graph, which includes border crossings with Foreign. There is an additive cost $t_{mm',g}$ and an ad-valorem cost $\tau_{mm',g}$ assigned to each direct road connection or border crossing between two markets m and m' . The trade costs between any arbitrary two markets can then be computed from the additive and ad-valorem trade costs along the lowest cost sequence of markets that are directly connected by a road or by a border crossing in the case of Foreign.

To illustrate how this works across rural markets, consider the case in which there are only additive-trade costs. In this case, if farmer i in market m imports crop k from market m' , her price must be

$$p_{i,k} = p_{m',k} + \sum_{m'',m''' \in \{m' \rightarrow m\}} t_{m''m''',k} + t_{mi,k}, \quad (2)$$

where $\{m \rightarrow m'\}$ indicates the sequence of markets that is the lowest cost route between market m' and farmer i 's market m , while $t_{mi,k}$ covers the trade cost from market m to i 's farm gate. Since agents can always export a crop to Foreign or import a crop from Foreign, their export price cannot fall below $p_{F,k}^*$ net of any costs between farm gate and border, while their import price cannot exceed $p_{F,k}^*$ plus trade costs from the border. In this sense, $p_{F,k}^*$ provides an upper bound for any agent's sale price of crop k and a lower bound for the purchase price. The market m price for crop k within this interval depends on whether the market as a whole imports/exports the crop from/to other domestic markets or Foreign, in the same way in which farm-gate prices depend on the direction of trade between farm and market. In equilibrium, this price is therefore a function of the difference in local consumption and production (excess demand), and the market's access to other markets for import/export (remoteness).

Following the convention in the trade literature, we assume that trade in manufacturing incurs iceberg trade costs only (i.e., $\tau_{ij,g} \geq 1$ and $t_{ij,g} = 0$ for all i, j and all manufacturing

¹⁴By assuming that agricultural inputs are all imported, we focus on the impact of input subsidies on farmers and ignore potential knock-on effects on domestic production of those inputs. In the Ugandan case, fertilizer is purely imported while hybrid seeds are mostly imported (IFDC, 2014).

goods g). Further, we assume that costs of selling or hiring labor across markets or with Foreign are prohibitively high. While we thus abstract from commuting and migration in our application, the model can readily incorporate it (see Appendix 5.C).¹⁵

2.4 Equilibrium

In equilibrium, farmers and urban households maximize utility taking prices as given, prices on the production side are subject to no-arbitrage conditions given trade costs within and between markets, and all markets clear. To formalize the definition of the equilibrium, we introduce excess demand function $\chi_{i,g}(\{p_{i,g}\}; W_{i,g})$ to capture the difference between the value of demand and supply of good g for agent i (see Appendix 5.A for detailed expressions). The function's arguments are the vector of endogenous prices $\{p_{i,g}\}$ and the vector of exogenous variables

$$W_{i,g} = \{L_i, Z_i, \{a_{i,g,\omega}\}, b_{i,g}, \{\alpha_{i,g,\omega}\}, \{\beta_{i,g,\omega}\}, p_{i,n}^*, a_h\}.$$

The equilibrium is a set of prices, $\{p_{i,g}\}$ and trade flows $\{x_{ij,g}\}$ measured in quantities at the destination, such that excess demand is equal to the difference in the value of imports ($\sum_j x_{ji,g}$) and exports ($\sum_j \tau_{ij,g} x_{ij,g}$) for each agent i and good g ,

$$\chi_{i,g}(\{p_{i,g}\}; W_{i,g}) = p_{i,g} \left(\sum_j x_{ji,g} - \sum_j \tau_{ij,g} x_{ij,g} \right) \quad \forall i, g, \quad (3)$$

and no-arbitrage conditions hold for all goods g ,

$$\tau_{ij,g} p_{i,g} + t_{ij,g} \geq p_{j,g} \perp x_{ij,g}, \quad \forall i, j. \quad (4)$$

Due to our assumption that crops are homogeneous goods, the direction of trade is determined in equilibrium, so that prices and trade flows must both be part of the equilibrium definition. This stands in contrast to work-horse models in international trade or economic geography, which make assumptions to ensure that the direction of trade is predetermined (e.g., by assuming goods are differentiated by origin of production), and a set of prices is sufficient to characterize an equilibrium. Here we need to keep track of a much larger set of equilibrium variables, as well as devise a new solution method to simulate policy interventions using our the model.

¹⁵Meaningful migration responses have not been found empirically in the context of the typical agricultural policies we consider here (e.g. [Huntington and Shenoy \(2021\)](#)), or in the context of broader shocks to agricultural productivity due to extreme weather events (e.g. [Emerick and Burke \(2016\)](#)).

2.5 Solution of Policy Counterfactuals

We are interested in solving for counterfactual equilibrium outcomes due to policy interventions. Policy interventions can be captured by changes in the vector of exogenous variables $W_{i,g}$ or trade costs $t_{ij,g}$ and $\tau_{ij,g}$ for some set of agents, chosen by the policymaker or researcher. Examples could be productivity shocks ($a_{i,g,\omega}$ or a_h), subsidies on modern intermediate input prices ($p_{i,n}^*$), land reform (Z_i), or infrastructure policies ($t_{ij,g}$ and/or $\tau_{ij,g}$). While our model thus allows for different types of interventions, we focus our exposition here on a change in the exogenous price of intermediates, $p_{i,n}^*$, consistent with the input subsidy we implement in the application below. The counterfactual equilibrium consists of a new set of prices $\{p'_{i,g}\}$ and trade flows $\{x'_{ij,g}\}$ which satisfy the equilibrium conditions above. Letting $\hat{x} = x'/x$ be the proportional change of any variable x from the initial to the counterfactual equilibrium denoted with x' (using ‘hat notation’), we can write the conditions for the counterfactual equilibrium as

$$\chi_{i,g} \left(\{\hat{p}_{i,g} p_{i,g}\}; \hat{W}_{i,g} W_{i,g} \right) = \hat{p}_{i,g} p_{i,g} \left(\sum_j x'_{ji,g} - \sum_j \tau_{ij,g} x'_{ij,g} \right) \quad \forall i, g, \quad (5)$$

$$\tau_{ij,g} \hat{p}_{i,g} p_{i,g} + t_{ij,g} \geq \hat{p}_{j,g} p_{j,g} \perp x'_{ij,g}, \quad \forall i, j. \quad (6)$$

The change in the price of the intermediate input, $\hat{p}_{i,n}^*$, leads to changes in land returns, $\hat{r}_{i,k,\omega}$, depending in part on the extent to which farmer i uses intermediates to produce crop k with technique ω . Changes in land returns feed into the excess demand functions $\chi_{i,g}(\bullet)$ through changes in farmers’ production choices, while the no-arbitrage conditions continue to hold in the counterfactual equilibrium. All prices, $\hat{p}_{i,g}$, adjust to bring the economy back to equilibrium.

To solve systems such as this, the literature in international trade typically applies a method called ‘exact hat algebra’ (Dekle et al., 2007). The idea is that once the endogenous variables in a system of equations can be written entirely in hat notation, the researcher can compute counterfactual equilibria by solving for a vector of endogenous *price changes* without having to recover the vector of initial price levels and unobserved fundamentals such as farmers’ productivities or preference shifters. The conventional assumptions about the manufacturing sector that we have invoked here, specifically ad-valorem trade costs and differentiated varieties across origins, imply that the counterfactual equilibrium conditions

for this sector satisfy the requirements for solving with exact hat algebra.¹⁶

In contrast, the agricultural sector features homogeneous goods such that prices are determined in combination with the direction of trade, which can change in the counterfactual equilibrium. Additionally, allowing for additive trade costs leads to a pass-through that is both imperfect and a function of the initial level of prices. Formally, we can directly observe that the right-hand side of equation (5) and the no-arbitrage conditions in (6) are in terms of initial *price levels* and counterfactual trade flows, thus precluding us from applying exact hat algebra for the agricultural sector. To overcome this complication, we develop a two-step procedure. In the first step (“Price Discovery”), we use the initial equilibrium conditions in combination with rich but available microdata to solve for the full vector of farm-gate prices in the observed (initial) equilibrium. In the second step, we then use the recovered price levels to obtain the vector of price changes $\{\hat{p}_{i,g}\}$ and counterfactual trade flows $\{x'_{ij,g}\}$ as a solution to the above system of equations.

As we discuss in Section 4, from the microdata we can either observe or directly infer the following variables: farmers’ and urban households’ expenditure shares on crops, $\xi_{i,k}$, $\xi_{h,k}$, Foreign crop prices, $p_{F,k}^*$, physical crop output, $q_{i,k,\omega}$, cost shares, $\alpha_{i,k,\omega}$, $\beta_{i,k,\omega}$, labor endowments, L_i , income of urban households, I_h , and trade costs for crops and labor: $t_{ij,k}$, $\tau_{ij,k}$, $t_{ij,l}$, and $\tau_{ij,l}$. We denote this set of observable variables used for the “Price Discovery” by

$$D_A = \left\{ \xi_{i,k}, \xi_{h,k}, p_{F,k}^*, q_{i,k,\omega}, \alpha_{i,k,\omega}, \beta_{i,k,\omega}, L_i, I_h, t_{ij,k}, \tau_{ij,k}, t_{ij,l}, \tau_{ij,l} \right\}.$$

With these variables from the initial equilibrium in hand, we can recast the excess demand functions for crops and labor of farmers, urban households and Foreign as functions of prices, $\{p_{i,g}\}$, and data D_A .¹⁷ We can then recover crop prices $\{p_{i,k}\}$ and wages $\{p_{i,l}\}$ in the initial equilibrium as a solution to the following system of equations for crops and labor only:

$$\tilde{\chi}_{i,g}(\{p_{i,g}\}; D_A) = p_{i,g} \left(\sum_j x_{ji,g} - \sum_j \tau_{ij,g} x_{ij,g} \right) \quad \forall i, \forall g \in \{k, l\} \quad (7)$$

$$\tau_{ij,g} p_{i,g} + t_{ij,g} \geq p_{j,g} \perp x_{ij,g}, \quad \forall i, j, \forall g \in \{k, l\}. \quad (8)$$

Using the recovered baseline prices and data we can in turn compute farmer income $\{I_i\}$.

In Appendix 5.B, we describe how to transform this price discovery step into an equiva-

¹⁶For more details on how we can use exact hat algebra to solve for the counterfactual equilibrium in the manufacturing sector see [Technical Appendix](#).

¹⁷We present these adjusted excess demand functions, $\tilde{\chi}_{i,g}(\bullet)$, in Appendix 5.A.

lent problem of finding the equilibrium of an exchange economy that is integrated as a small open economy with the rest of the world. Importantly, price discovery is more general than the full equilibrium formulation above, since we do not have to specify functional forms for crop supply and demand for this step but instead take production quantities and expenditure shares as given in the data. Consequently, the set of baseline prices is consistent with alternative formulations of the production process and demand system. We then provide a proof that, in the case of ad valorem trade costs, the goods in such an economy satisfy the connected substitutes property 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.¹⁸

3 Combining the Model with Local Experiments

While the interplay of forces governing how shocks propagate across households and markets described in the theory are difficult to capture in small-scale experiments alone, local experiments can play a critical role in informing policy impacts at scale when combined with the model. In this section we describe how we use treatment effects estimated in local RCTs to identify the key elasticities governing demand (σ, ζ, η) and supply (κ, μ). To do so, we use local experiments from a variety of East and Southern African countries: Uganda, Kenya, and Mozambique.¹⁹ [Appendix 1](#) provides additional details about the data used in the estimation.

Demand Estimation

To estimate the elasticity of substitution between crops in consumption, σ , we bring to bear a demand-side experiment conducted in [Bergquist and Dinerstein \(2020\)](#). This experiment was conducted in open-air maize markets in rural western Kenya, 30 km 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

¹⁸The sufficient conditions in our proof of uniqueness in [Appendix 5.B](#) no longer hold in the presence of additive trade costs because the demand for foreign goods is no longer strictly 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 vectors of farm-gate prices drawn randomly along the range of possible prices given exogenous international prices and trade costs. Reassuringly, we find the same equilibrium in all cases.

¹⁹Although these settings may differ in some dimensions, rural areas across the countries we include share many features, including crops grown, farming methods (mostly rain-fed agriculture), and overall levels of development.

amounts. The value of the discount ranged from roughly 0-15% of the baseline price. To estimate σ in the model, we regress log quantity purchased by individual i from seller s in market m on date d on log price,

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

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 $\sigma = 1$.²⁰

Table 1: Estimation of σ

	(1)	(2)	(3)	(4)
	OLS	IV	OLS	IV
Log P	-4.81*** (0.27)	-0.94 (0.62)	-5.02*** (0.34)	-1.00* (0.55)
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

Dependent Variable is Log Quantities (Instrument is Randomized Subsidy Amounts). Standard errors clustered at level of communities. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

To calibrate the demand parameter ζ , that governs the degree of 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 food consumption, and $\xi_{i,A}$ is the observed share spent on total food consumption, $\xi_{i,A} = \sum_k \xi_{i,k}$. Stone-Geary preferences imply that the share of income spent on subsistence approaches zero for the richest households, so we calibrate ξ_A to the average share of expenditure spent on total food consumption among the richest 5 percent of Ugandan households, which

²⁰A possible limitation of this RCT is that the subsidies were fairly short-run in their duration. [Bergquist and Dinerstein \(2020\)](#) (Appendix C) address this concern, exploiting the randomized order of their treatment periods to test for evidence of inter-temporal dynamics in demand. They find no evidence of stockpiling, in line with limited storage in this empirical setting ([Burke et al., 2019](#)). As part of the quantitative analysis, we also explore the sensitivity of the counterfactuals to a range of higher or lower values for σ in Section 7.

is close to 0.1 in the Uganda National Panel Survey (UNPS). 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 (see [Head and Mayer \(2014\)](#)).

Supply Estimation

To estimate the first key supply elasticity κ , that governs farmers' choice of land allocation across modern or traditional techniques within crops, we take advantage of the experiment implemented by [Carter et al. \(2020\)](#). Randomly selected farmers in Mozambique were offered fertilizer and improved seeds at a subsidized price. Data collected on farmers' use of modern inputs and output quantities by plot allow the estimation of the impact of changing input prices (instrumented by treatment) on land allocations across production techniques. We derive the following estimation equation from Section 2:

$$\log \left(\frac{\pi_{i,k,1}}{\pi_{i,k,0}} \right) = - \left(\kappa \frac{\alpha_{i,k,1}}{\delta_{i,k,1}} \right) \log p_{i,k}^{input} + \epsilon_{i,k},$$

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 ($p_{i,k}^{input}$) 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, $\alpha_{i,k,1}$ and $\delta_{i,k,1}$ respectively.

Table 2: Estimation of κ

	First Stage		Reduced Form		IV	
	(1) Cross-Section	(2) Panel	(3) Cross-Section	(4) Panel	(5) Cross-Section	(6) Panel
Treat	-0.76*** (0.05)	-0.76*** (0.05)	0.62* (0.36)	0.64* (0.36)		
Log Input Price					-0.82* (0.48)	-0.83 (0.49)
Observations	63	127	63	127	63	127
Community FX		Yes		Yes		Yes
Round FX		Yes		Yes		Yes
F-Stat					205.92	207.53

Dependent Variable is $\log \frac{\pi_{i1|kt}}{\pi_{i0|kt}}$ (Instrument is RCT Treat Indicator). Standard errors clustered at level of communities. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

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.²¹ 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 survey-round and village fixed effects.²² The IV point estimates in columns 5 and 6 are 0.82 and 0.83. Using the ratio of cost shares of land over fertilizer and hybrid seeds, this would imply that $\kappa = 2.5$. We use this estimate of the lower-nest (within-crop) elasticity as our baseline. To address the somewhat wide confidence intervals we find here, we explore the sensitivity of the counterfactual analysis across a range of higher or lower values for κ in Section 7.

We complement this RCT with a natural experiment that we use in Uganda to estimate the upper-tier supply elasticity in our model for substitution of land allocations across crops, μ .²³ The estimation equation derived from the model in Section 2 is:

$$\log \left(\sum_{\omega'} \pi_{i,\omega',t|k}^{-1} \left(\frac{q_{i,k,\omega',t}}{l_{i,k,\omega',t}^{\beta_{i,k,\omega'}} m_{i,k,\omega',t}^{\alpha_{i,k,\omega'}}} \right)^{\frac{1}{\delta_{i,k,\omega'}} \frac{\kappa}{\kappa-1}} \right)^{\frac{\kappa-1}{\kappa}} = \left(\frac{\mu-1}{\mu} \right) \log \pi_{i,k,t} + \log Z_{i,t} + \log \tilde{b}_{i,k,t} \quad (9)$$

The left-hand side of (9) is farmer i 's harvest quantities ($q_{i,k,t}$) for crop k aggregated across both techniques in survey year t , adjusted in the denominator for the reported quantities of labor ($l_{i,k,\omega,t}$), modern intermediates ($m_{i,k,\omega,t}$) 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 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

²¹Given 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.

²²Carter 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.

²³The experiment in Carter et al. (2020) did not induce changes in the allocation of land across crops that one could use for estimating μ .

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 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 $\frac{\Delta p_{world}}{p_{world,t_0} + bordercost_i}$. 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) and likely exogenous to domestic production (Uganda accounts for less than 2% of world coffee sales). We 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.²⁴

Panel A of Table 3 presents the first-stage regression both before and after including region-by-crop-by-technology-by-time fixed effects, and using all years of data (2005, 2009, 2010, 2011 and 2013) or just using long changes 2005-2013. 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. 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. Panel C presents the second-stage estimation of equation (9). We find significant point estimates in the range of 0.45-0.75. Recall that 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 reported in Sotelo (2020) ($\mu = 1.7$). To be conservative, we pick the low estimate of $\mu = 1.8$ as our baseline calibration.²⁵

²⁴Appendix Figure A.2 documents that 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.

²⁵As we show in the sensitivity analysis in Section 7, this is conservative in terms of welfare impacts, and in terms of the difference between local-vs-at-scale effects.

Table 3: Estimation of μ

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
	All Years	All Years	All Years	All Years	2005-13	2005-13	2005-13	2005-13
Panel A: First Stage								
Instrument		-0.467*** (0.121)		-0.367** (0.172)		-0.925*** (0.161)		-1.132*** (0.267)
Panel B: Dependent Variable is Log Harvest ($\log(q_{i,k,t})$)								
$\log\pi_{i,k,t}$	0.357*** (0.016)	0.793* (0.420)	0.357*** (0.016)	0.631 (0.494)	0.415*** (0.034)	0.925*** (0.310)	0.425*** (0.032)	0.904*** (0.216)
Panel C: Dependent Variable is Log Adjusted Output								
$\log\pi_{i,k,t}$	0.411*** (0.036)	0.394 (0.539)	0.406*** (0.036)	0.785 (0.723)	0.441*** (0.060)	0.553 (0.421)	0.438*** (0.062)	0.754** (0.315)
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
1st Stage F-Stat		14.80		4.57		32.82		17.93

Standard errors clustered at level of counties. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

4 Calibration of a Granular Economic Geography

The structure of the model allows us to combine the key elasticities that we estimate from local experiments with rich household microdata to capture a granular and realistic economic geography when evaluating effects at scale. In this section, we describe how we populate the vector of observable data used in the price discovery and counterfactual solution laid out in Section 2: $D_A = \{\xi_{i,k}, \xi_{h,k}, p_{F,k}^*, q_{i,k,\omega}, \alpha_{i,k,\omega}, \beta_{i,k,\omega}, L_i, I_h, t_{ij,k}, \tau_{ij,k}, t_{ij,l}, \tau_{ij,l}\}$.

Appendix 1 provides a summary and additional discussion of the administrative data used below in the calibration. These data are increasingly available in many low and middle-income countries. We consider 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. The intermediate input used in production under the modern technique encompasses chemical fertilizer and hybrid seed varieties in our empirical context. To estimate the cost shares of intermediates

($\alpha_{i,k,\omega}$), labor ($\beta_{i,k,\omega}$) and land ($\delta_{i,k,\omega}$) in the production function of each crop-technique-location combination, we take the average of the cost shares observed across households in the Uganda National Panel Survey (UNPS) microdata for each of the 4 regions of the country (Appendix Table A.4).

To calibrate the model to the full set of local markets and households populating Uganda, we require household-level information on pre-existing production quantities ($q_{i,k,\omega}$) and expenditure shares across crops and sectors ($\xi_{i,k}$, $\xi_{h,k}$) for the full population of households observed in the census microdata, which is generally not available as part of census data.²⁶ 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 below). For each of these outcomes from the UNPS on the left-hand side, we project them (using UNPS survey weights) 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 household head, number of dependents, number of household members, an asset ownership index (using the same assets reported in both datasets), potential yield given a farmer’s location from the FAO/GAEZ database, dummies for subsistence farming and urban households, district dummies and survey year fixed effects. For this estimation, we employ Poisson pseudo-maximum likelihood, which has the nice property of preserving aggregates in the predicted population data.

Geography and Trading Frictions

Households are located in roughly 4,500 rural parishes, which we treat as local markets (“villages” in the model), and 70 urban centers (“cities” in the model). Figure 1 presents a map of this setting.

²⁶Household 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 are from the FAO database we describe in Appendix 1.

Figure 1: Ugandan Markets and Transportation Network

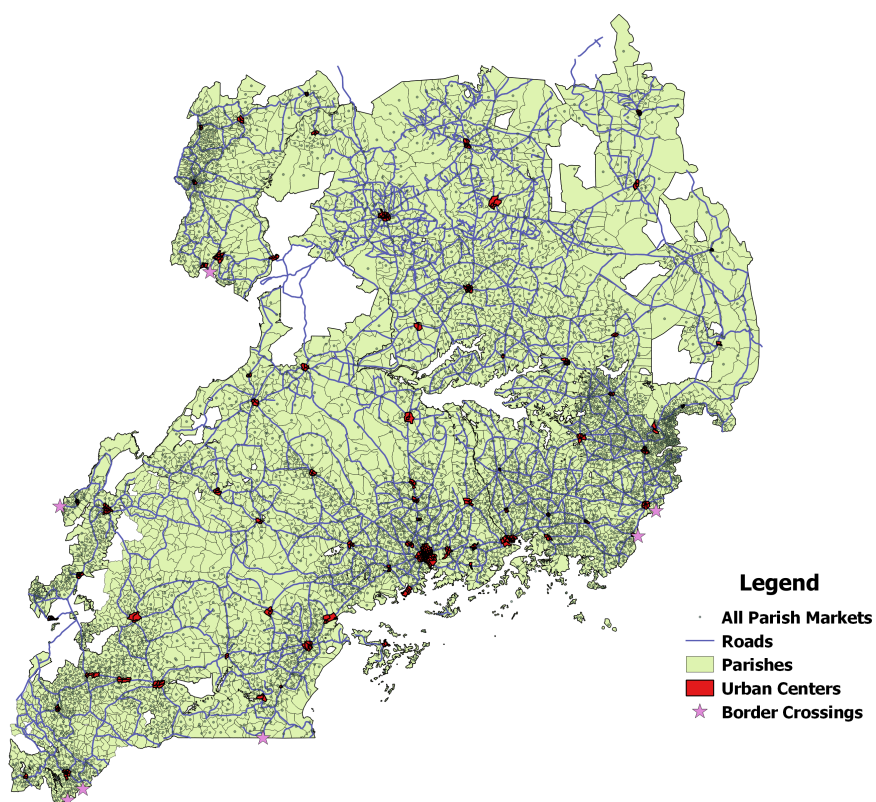


Figure displays local parish markets, urban markets, border crossings and the road network in Uganda.

To calibrate trade frictions across local markets, we use survey microdata collected by [Bergquist et al. \(2024\)](#) 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. 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. [Appendix 2](#) uses this database to explore the nature of trade costs between rural markets. The estimates suggest that trade costs are predominantly per-unit (additive) instead of being charged ad valorem in our context. Using only bilateral price gaps from market pairs during months in which we observe positive trade flows between the pair (following spatial arbitrage in the model), with information on the road distance traveled between the markets from the transportation

network database in Figure 1, we estimate the following specification:

$$t_{ij,g,t} = (p_{j,g,t} - p_{i,g,t}) = \alpha + \beta (\text{RoadDistance}_{ij}) + \epsilon_{ij,g,t},$$

where t indexes survey rounds and the error term $\epsilon_{ij,g,t}$ is clustered at the level of bilateral pairs (ij) . RoadDistance_{ij} is measured in road kilometers traveled along the transportation network. We estimate a single function of per-unit trade costs with respect to road distances across all crops, so that $t_{ij,g} = t_{ij}$.²⁷ The estimated trade cost for an additional road kilometer traveled between two markets is 1.2 Ugandan shillings (standard error 0.3), which implies a cost of about US\$0.5 per kilometer for one ton of shipments.²⁸ If we replace the specification above to be in logs on both left and right-hand sides, the distance elasticity is 0.0258 (standard error 0.0057), which is close to existing evidence for within-country African trade flows (e.g., [Atkin and Donaldson \(2015\)](#)). We use this distance elasticity to calibrate ad valorem trade costs τ_{ij} 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. 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 m at year-month t and $p_{m,g,t}$ is the local market price that we do not directly observe and capture with market-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. To measure farm-gate trade costs in the full-count census, we then project these residuals on the same household characteristics that we use in the prediction of consumption and production in the previous section.²⁹ 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 64 Ugandan shilling per kilogram, which amounts to roughly 8 percent of the average crop price. Finally, we use the UNPS microdata to estimate the trading frictions farmers face when hiring or

²⁷We do so for reasons of statistical power. The dataset covers two crops, maize and beans. Including a crop-month FE in the regression above yields very similar results.

²⁸[Bergquist et al. \(2024\)](#) document 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 \(2007\)](#)).

²⁹Since the distribution of trade costs is mechanically centered at zero, 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).

selling labor in the local market with the same strategy 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). On average, hiring farmers is subject to labor trading frictions of 241 shilling (or 10 US cents) per day for hiring a worker, or around 4% of the daily wage. Viewed through the lens of the “value of travel time” (Becker, 1965), this would be consistent with average forgone wages due to roughly 25 minutes of transit for hired labor on an 8-hour workday.

5 Counterfactual Analysis: Local vs. At-Scale 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. In a survey of 10 African countries, Jayne and Rashid (2013) find that input subsidy programs account for on average 30% of total public expenditure on agriculture, and estimate that over 60% of Sub-Saharan Africa’s population lives in a country with a major input subsidy program.³⁰ In this section, we investigate how the welfare impacts of such a policy differ between a local intervention and one at scale – among the same sample of farmers – and quantify the underlying mechanisms. In doing so, we also explore how GE forces shape the distributional effects of the policy.

Local Effects vs Scaling Up: Average and Distributional Effects

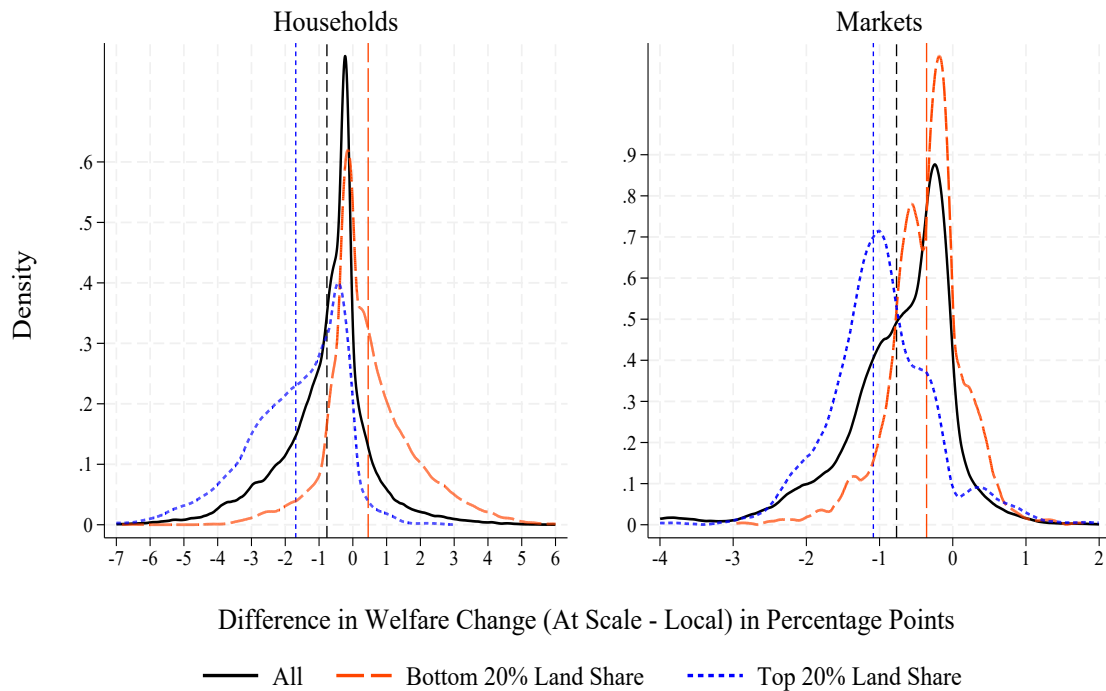
We focus on the effects of a 75 percent cost subsidy for modern inputs (chemical fertilizers and hybrid seed varieties in the data).³¹ We run two types of counterfactuals in the calibrated model. In the *local intervention*, we randomly select a 2.5 percent sample in each of the 4500 rural parishes (roughly 100,000 households nationwide). For each of the markets, we then treat 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

³⁰In Section 7, we discuss other types of policy counterfactuals that can be studied using our approach, or to which it could be extended.

³¹The subsidy uniformly lowers the price for intermediates $p_{i,n}$ by 75% across treated farmers i . To simplify the exercise, we leave aside for the moment the public finance dimension of the subsidy (akin to financing by international organizations or foreign donors). It would be straightforward to, e.g., have this financed by a lump-sum tax in the model.

separate small-scale RCTs. For the *at-scale intervention*, we offer the subsidy to all farming households in the Ugandan 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 scaled to the rest of the Ugandan countryside.

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



Notes: Figure plots distribution of percentage point difference in welfare (real income) changes from at-scale minus local interventions for sample of $\sim 100,000$ rural households (left panel), and their averages across parishes (right panel). Vertical bars indicate mean differences.

Figure 2 presents the difference in welfare effects between the at-scale and local interventions across all $\sim 100,000$ national sample households. Changes in welfare in the model are changes in real incomes, with the price index defined as the ideal price index given by the nested Stone-Geary preferences from Section 2. The left panel shows the at-scale impact minus the local intervention's 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. 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 ($\frac{r_i Z_i}{I_i}$ in the model).³² Those in the bottom quintile – whom we refer 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.³³

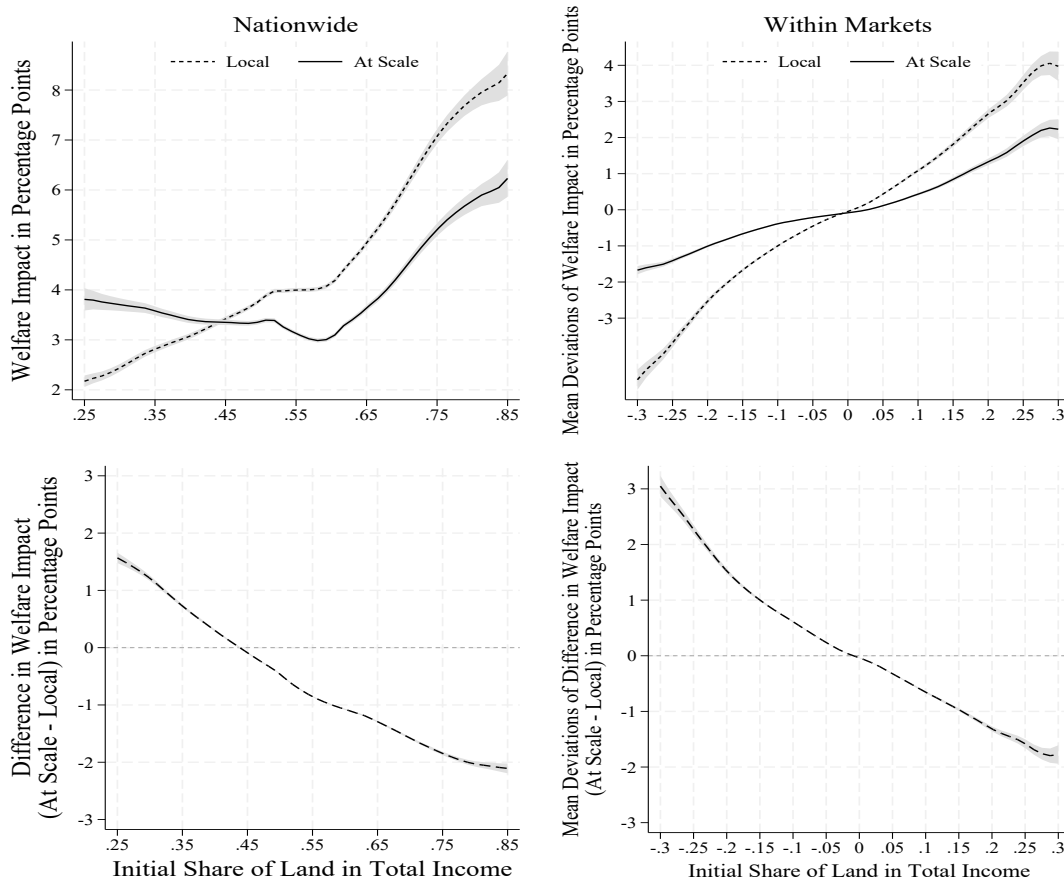
Two main insights emerge. First, the distribution is wide, with households experiencing more than +/- 5 percentage point changes in their welfare impact when the intervention is scaled-up, with the average household experiencing a decrease of about 1 percentage point, or about 20% of the average local welfare effect (Table A.5 shows the point estimates of local and at-scale effects across the different groups in Figure 2). 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.

To further investigate the distributional implications of scaling in this context, Figure 3 presents the non-parametric estimates of the local and at-scale welfare effect as a function of initial land income shares. We see in the top left panel that the local intervention is regressive, benefiting land-rich households more than the land-poor by roughly 5.5 percentage points moving from left to right in land income shares. However, the at-scale intervention is substantially less regressive, reducing this gap by more than half, to 2 percentage points. Driving this compression is the fact that land-poor households experience larger gains at scale than under the local intervention, with the poorest households experiencing welfare gains that are 1.5 percentage points larger at scale, as shown in the bottom left figure. In contrast, land-rich households fare worse at scale, with the richest experiencing a 2 percent-

³²Total income is the sum of returns from land (selling crops net of input costs) and labor income.

³³Appendix Figure A.3 plots flexibly estimated relationships between our measure of land income shares and households’ land ownership in acres or households’ calibrated total incomes. Fink et al. (2020) document similar patterns in another African context (Zambia).

Figure 3: Distributional Implications

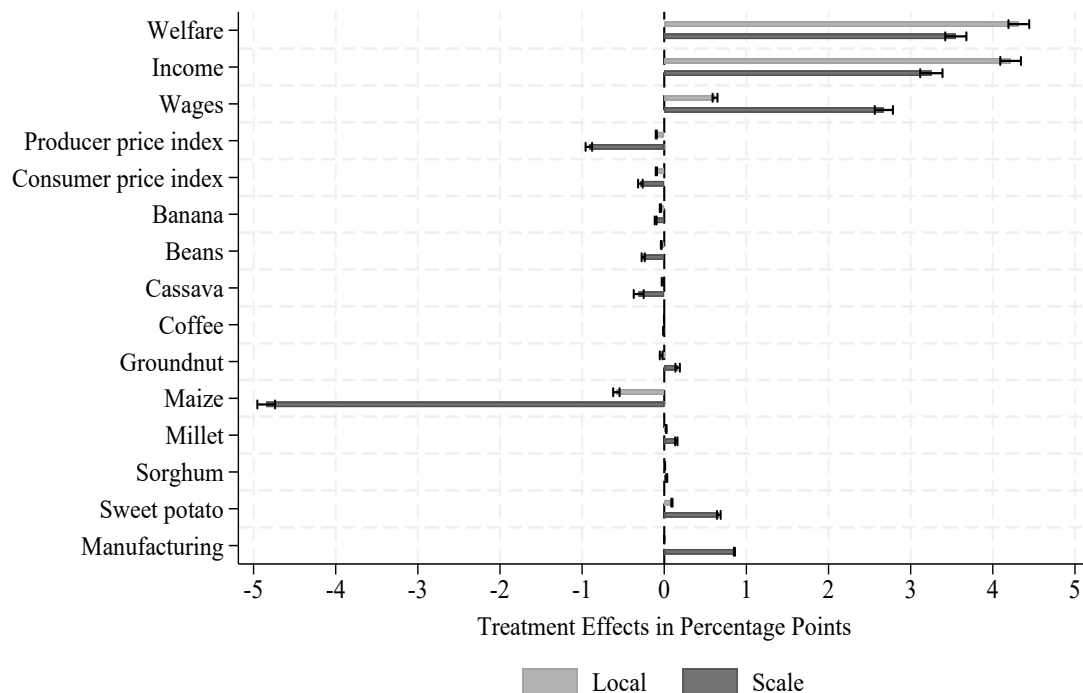


Notes: Figure plots local polynomial regressions of percentage point welfare changes of the at-scale and local interventions (in the top row) and the difference in impact at scale (at-scale impact - local impact) in the bottom row, both against initial household land income shares. Right column uses deviations from the parish means on both axes to study within-market variation. Shaded areas indicate 95% confidence intervals.

age point drop in their welfare gains relative to the local intervention. Qualitatively similar differences are present in the right column when comparing land-rich and -poor households within markets, after conditioning on parish market fixed effects, suggesting that these effects are not predominantly driven by differences across locations.

What drives differences in these impacts at scale, both the decrease in average effects and the attenuation of the policy's regressivity? Figure 4 presents mechanisms, displaying average treatment effects on welfare, income, and prices under the local intervention (in light gray) and the at-scale intervention (in dark gray). We see that the local intervention drives little movement in local market prices, with on average small changes in crop prices, minor increases in local wages, and no effect on manufacturing prices. Cost savings on modern

Figure 4: GE Mechanisms



Notes: Figure plots average impacts under the local intervention (in light gray) and under an at-scale intervention (in dark gray) on the following outcomes: welfare, income, and prices, including wages, manufacturing prices, and prices for each of the nine study crops. 95% confidence intervals are in black whiskers.

inputs and greater usage of these techniques, without meaningful changes to output or other input prices, translate to meaningful increases in farmer nominal incomes, which rise on average by 4.2 percentage points, and almost no change in the consumer price index. As a result, average welfare goes up by 4.3 percentage points. As these gains are concentrated in land income, the local effect favors land-rich households. They also benefit more from the local intervention in part because of higher pre-existing usage of modern technology (Figure A.4).

When the intervention is scaled up to the national level, including to 100% of the local farmers, prices move substantially more. Crop prices mostly fall due to increased supply; the drop is especially notable for maize, the most commonly grown crop in Uganda that also has one of the highest pre-existing production shares using the modern technology (Figure A.5). In contrast, prices for a few relatively less modern input-intensive crops rise as farmers shift out of these crops. Local wages rise, due to increased demand for local labor, a force that is amplified by the (non-Hicks neutral) nature of the increase in modern input usage. Finally,

manufacturing prices rise as demand for these goods increases, a channel that propagates the gains of the policy from the countryside to urban centers. Average consumer prices on net fall, though these effects are more muted, with the consumer price index decreasing by only 0.29 percentage points. Together, this brings down average welfare gains by 18% relative to the local intervention, from 4.3 to 3.5 percentage points. The distributional changes at scale are driven by the fall in producer prices and the rise in local wages, both favoring land-poor households who receive a larger share of their total income from labor relative to land.³⁴

6 Insights for Research in Agricultural Development

In this section, we use the framework to provide additional insights for research at the intersection of agriculture and development. First, we investigate how the sign and extent of GE forces differ as a function of saturation rates at different geographical scales, with implications for randomized saturation designs in the RCT literature. Second, we investigate the types of households and markets in which GE forces are most critical in changing a policy's impact at scale. Third, we use these insights to discuss implications for the large number of agricultural experiments in Sub-Saharan Africa conducted over the last 20 years, exploring how consideration of GE forces interacts with study site selection, and implications for publication bias.

GE Forces as a Function of the Intervention's Scale

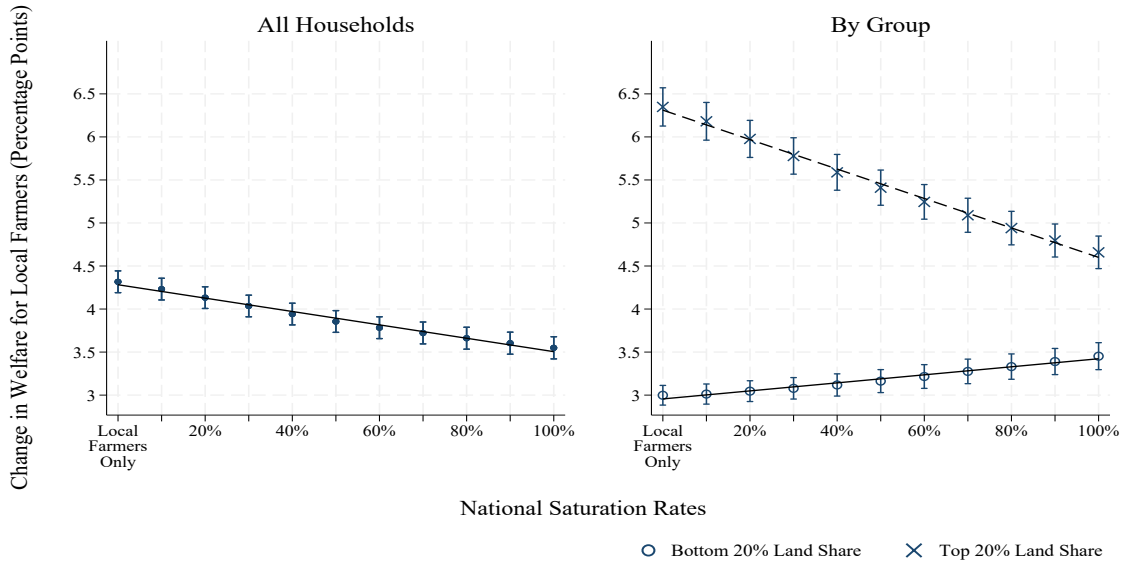
How do GE effects evolve as the intervention is scaled up to an increasingly large fraction of households nationwide, and as the geographic scale of the cluster is varied? Both have implications for the optimal design and lessons that can be learned from randomized saturation designs, an experimental approach to capturing GE effects. In these designs, the fraction of individuals treated (the "saturation") is randomized across geographic areas or "clusters" to study the market-level outcomes that emerge (see e.g., Baird et al. (2011); Burke et al. (2019); Egger et al. (2022)). It is most common to use only two levels of saturation, due to budget and logistical constraints, and clusters are often defined at the level of villages or groups of villages. Extrapolating from results from the two saturation points in this design to the impacts that would emerge under a full-scale, national policy would therefore require two assumptions: i) that GE forces are linear with respect to changes in the saturation rate;

³⁴Appendix Figure A.6 further documents the impacts of these GE forces on different components of incomes and consumer price indices as a function of initial land income shares.

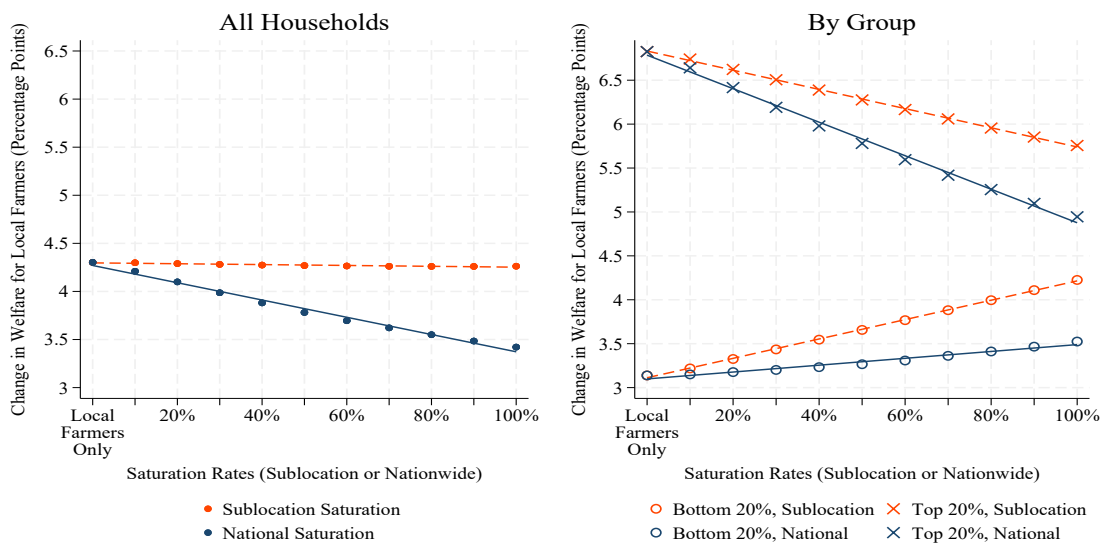
and ii) that the GE forces experienced at the level of local clusters are representative of the effects of saturation at a broader geographical scale (e.g. nationwide). We can use our approach to assess these two assumptions by computing counterfactuals to estimate how impacts evolve as a function of saturation rates at different geographical scales.

Figure 5: GE Forces as a Function of Saturation

Panel A: National Saturation



Panel B: National vs. Sublocation Saturation



Notes: Panel A plots average welfare effects for ~100,000 sample against national saturation rates with 95% confidence intervals. Panel B plots effects for subsample of ~6500 located in 51 subcounties against either national or subcounty saturation rates.

We start with the first question of linearity, presenting in Panel A of Figure 5 the 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 in all parishes. The right-most point on the x-axis corresponds to the at-scale intervention above where 100% of rural Ugandan households receive the subsidy 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.³⁵ 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. We find 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 Uganda 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.³⁶ Consider 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) and plot results for this random subset of subcounties (which includes roughly 6500 households of the same national 2.5% sample); the blue line in Panel B shows results that closely mirror those in Panel A. In the orange line, we consider the more feasible randomized saturation design in which we vary the saturation

³⁵We 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.

³⁶For 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.

rate within the 51 subcounties, but not the rest of Uganda.

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). 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 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 randomizes 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 (Figure 4), whereas changes in the saturation within subcounty populations have much more muted implications on output prices on average. 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 – 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. At the same time, the right graph of Panel B in Figure 5 shows that randomized saturation designs do capture an important subset of GE forces that would be absent from small-scale field experiments, mainly related to knock-on effects of increased production on local factor prices and the labor market. Those insights could then be combined with approaches such as the one laid out here to assess impacts at a broader

geographical scale.

Where do GE Forces Drive the Biggest Changes in Impact at Scale?

Because of the granularity of our approach, which allows us to run counterfactuals akin to 4,500 RCTs across each parish of Uganda, we can identify for which types of households and markets GE forces most alter policy impacts at scale. This is relevant for understanding where researchers and policymakers should exercise the most caution in extrapolating from small-scale results and where approaches like ours could be most complementary.

In the previous section, we document the role of households' pre-existing land income shares as a significant driver of changes in the policy's impact at scale, relative to the local intervention. But, of course, other factors, such as types of crops, modern input usage, market remoteness or expenditure shares, could act as additional mediators. To get a more complete picture of the drivers of small-scale welfare impacts, impacts at scale, and their difference, Table 4 presents variance decompositions. We include all farmer and market characteristics that the model uses as information in the baseline equilibrium: a household's initial production by crop and technology regime (18 variables), trading frictions to other markets and within the local market, and expenditure shares across the 9 crops and manufacturing. We group these household and market characteristics into seven bins, as shown in Table 4. We then report the percentage of the total explained variation (R-squared) in three outcomes – the local effect, the at-scale effect, and the difference – that can be attributed to the group of variables (via Shapley decompositions, regressing each outcome on all variables described above for the roughly 100,000 household sample).³⁷

Table 4: Variance decomposition

	Land share	Modern share	Maize share	Coffee share	Oth. prod.	Consumption	Trade Costs
Local effect	4.15	39.98	0.15	13.01	28.37	10.13	4.22
At-scale effect	0.28	41.47	3.02	11.96	28.77	11.21	3.28
Scale-local effect	38.60	1.84	29.85	3.86	11.52	6.12	8.20

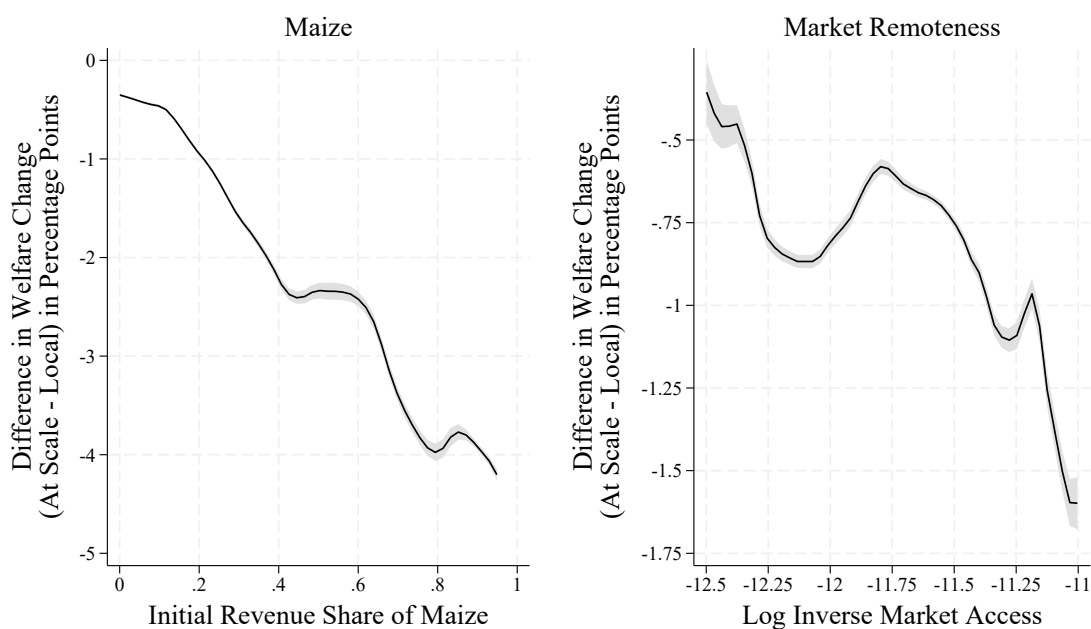
Notes: Table presents results of a Shapley decomposition, based on regressions of one of the three outcomes listed in rows on the left of the table. The regressions have the roughly 100,000 sample households as unit of observation. Each column corresponds to a group of right-hand-side variables included in the regression. The listed numbers indicated percentages of the total explained variation, summing to 100 in each of the rows.

While initial production quantities across crops and modern technology usage are the most predictive of the variation in the policy's local effect and the effect at scale across

³⁷The R-squared exceeds 0.95 for both local and at-scale interventions and 0.7 for the difference.

households, three characteristics stand out for explaining the *difference* between the local and at-scale effects: the household's initial land income share accounts for 39% of the total explained variation, the initial revenue share of maize explains another 30% and trade costs explain another 8%. In the previous section, we documented the mechanisms underlying the role of land income shares, which the variance decomposition here now shows quantitatively accounts for a large part of the overall variation in this difference of impacts at scale. Turning to the role of maize shares, Figure 4 above showed that maize, the most frequently grown staple crop, experiences the largest average price decline in the at-scale intervention, in part because it has among the highest pre-existing average usage of modern inputs (Figure A.5). As a result, producer prices change more in maize-producing regions, driving a larger wedge between the local and at-scale effects. The left panel of Figure 6 documents this mechanism, showing that the average difference in impact at scale (relative to local) goes from a minimal change for those with low initial revenues shares of maize to -4 percentage points among those that mainly grow maize at baseline.

Figure 6: Drivers of Average Changes at Scale

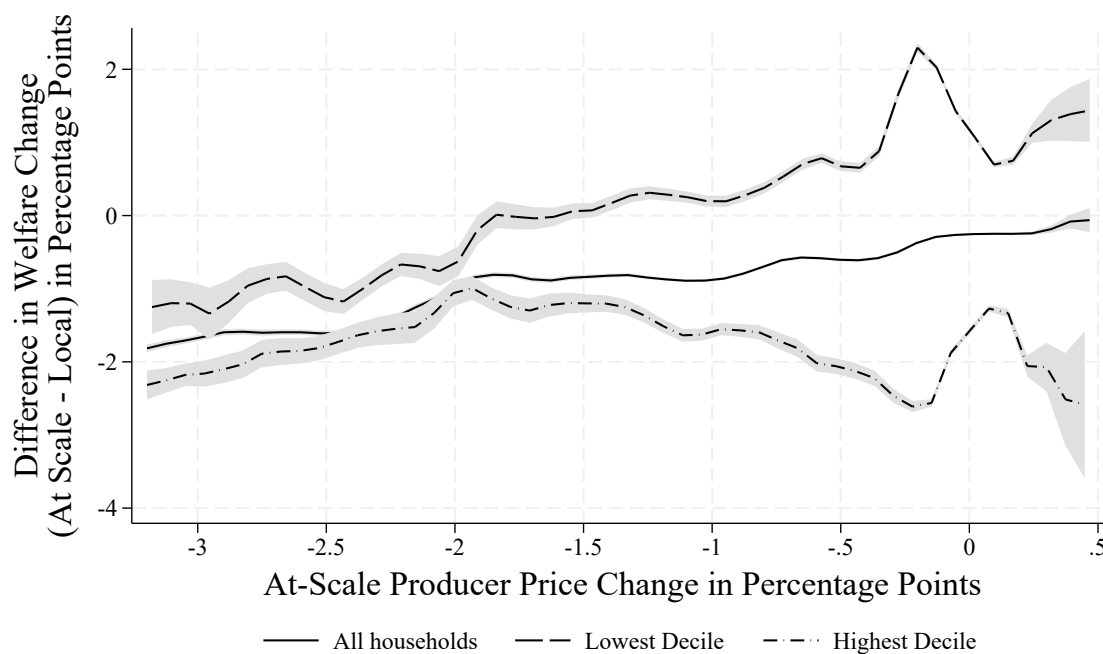


Notes: Figures are based on the roughly 100,000 household sample. Panels plot the difference of the effect at scale relative to the local intervention as a function of initial maize revenue shares (left) and log inverse market access (right). Shaded areas indicate 95% confidence intervals.

The right panel documents the role of market remoteness. Because sub-Saharan Africa has some of the highest transport costs in the world (e.g., [Porteous, 2019](#)), we expect re-

moteness to be a relevant mediator in our setting. Theory would predict that GE forces are stronger in poorly-integrated markets, where local prices are less pinned-down by market prices at the border or in larger cities due to higher trade cost wedges. Figure 6 confirms this hypothesis, showing that more remote markets – measured by the log inverse market access of local parish markets ($1 / \sum_{d \neq o} \frac{Pop_d}{Distance_{od}}$) – experience more pronounced declines in gains at scale compared to the local intervention. The reason is that agricultural producer prices are on average more negatively affected by the at-scale intervention in more remote markets, where higher trade cost wedges both weaken the disciplining role of external prices – allowing local supply expansions to have larger price effects – and limit the transmission of shocks from other regions.

Figure 7: Distributional Changes at Scale



Notes: Figure is based on the roughly 100,000 household sample. The graph plots the difference in effects at scale vs. the local intervention as a function of average producer price index changes due to the at-scale intervention in the parish, separately for all sample households and the top or bottom deciles of initial land income shares. Shaded areas indicate 95% confidence intervals.

Figure 7 turns attention to drivers of the changes in distributional implications at scale. First, in the solid line, the figure plots the average difference in impact at scale for all households as a function of the change experienced in the producer price index in their market when the policy is implemented at scale. As discussed above, places that see more pronounced drops in producer prices at scale experience larger reductions in the average wel-

fare gains of the policy. The two dashed lines then present this same relationship, but for the bottom and top deciles of households in terms of initial land income shares (“land poor” in long-dashed and “land rich” in short-dashed). On the left-most side of the figure, we see that both groups are negatively impacted in locations with large declines in output prices, as the returns to both land and labor go down with the declining overall returns in agriculture. However, in the right-most side of the figure, we see that in places where output price changes are more muted – perhaps in better connected markets – effects diverge for the land-rich and land-poor. In these locations, crop production rises without the offsetting reductions in prices. This puts upward pressure on wages in local labor markets, shifting gains from land to labor, which benefits the land-poor and reduces the gains to the land-rich. In contrast to the changes in average effect at scale, we therefore see that the most pronounced redistribution of gains from the local intervention toward households at the bottom of the distribution occur in markets where producer price changes are muted.

Implications for Study Selection and Publication Bias

We now put these insights into context with the existing body of closely related local experiments. We begin by identifying all agricultural RCTs that have been conducted in Sub-Saharan Africa and listed on the American Economic Association (AEA) Registry since its launch in 2013. This search yielded 143 RCTs conducted across 23 countries. Using location data, we merge in additional information about the three dominant drivers of changes in impact at scale that we identified above: households’ average land income shares, shares of maize in farmer revenues, and market remoteness.³⁸

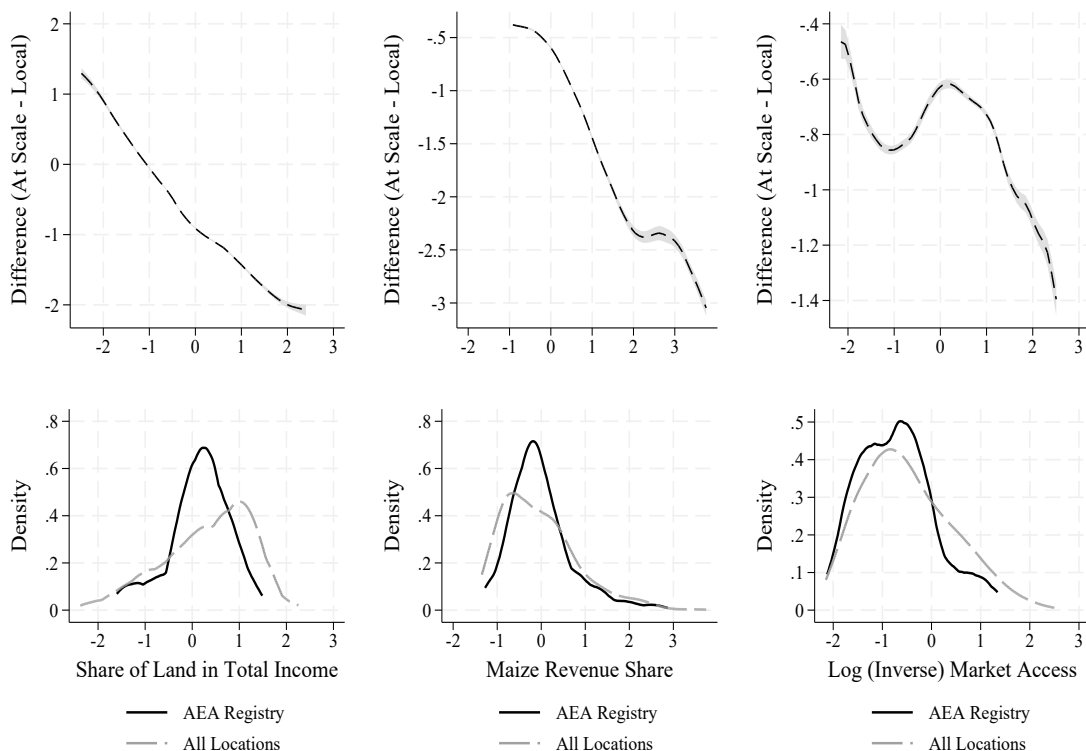
The three bottom panels in Figure 8 plot the distribution of these characteristics across study sites (solid line) and across the entire population of possible study sites, weighted by population (dashed line) in all rural regions of the same countries as a comparison to indicate propensities relative to random sampling among the rural population. In the top panels, we overlay the implications for how local effects change at scale, as estimated in the Ugandan policy setting above.

We find the existing body of local field experiments tend to be located in rural regions that have lower land income shares (p-value 0.12), lower maize revenue shares (p-value 0.32) and which are less remote (p-value 0.00) relative to the rural population of these countries as a whole. A comparison to the top panels suggests that these are the study sites and

³⁸ [Appendix 4](#) describes the data sources and provides details on the computations for each study location.

populations for which GE forces are less pronounced, driving weaker differences between the local and at-scale intervention on average. On the other hand, we found in Figure 6 that those locations also experience the most meaningful changes in distributional implications at scale relative to findings from local interventions.

Figure 8: Study Site Selection of Agricultural RCTs in Sub-Saharan Africa



Notes: Three bottom panels plot in the black solid line the distribution of the 143 agricultural RCTs conducted in Sub-Saharan Africa listed on the AEA RCT Registry since its launch in 2013 (see Appendix 4 for more detail). To facilitate comparison to a representative set of possible rural study populations, the dotted gray line presents the distribution of all rural grids or enumeration areas across the 23 relevant countries, weighted by population (with the 2.5% of grids or enumeration areas representing urban areas (defined as populations of over 300 people/sq km removed). From left to right, the plots present the distribution of the three characteristics found to be most predictive of the difference in impacts between the local and at-scale intervention in Table 4: households' average land income shares, shares of maize in farmer revenues and market remoteness (as measured by log inverse market access). All variables are presented as z-scores, normalized by the mean and standard deviation with the country in which the study was conducted. The top panel overlays the difference between the at-scale - local treatment effect with respect to these three characteristics, as estimated in our model.

It is possible that these study populations are in part intentionally chosen to avoid GE effects, which may be seen as identification threats given their potential to generate spillovers.³⁹

³⁹In practice, the lack of discernible price or wage effects due to the local intervention from our model

Or, more plausibly, they may be chosen simply for convenience of access or other unrelated reasons. Regardless, the end result is that the evidence we have on the impact of agricultural policies on locations and populations likely to be strongly affected by GE forces is “doubly thin” – not only are they less studied, but when local studies are conducted among these populations, measured effects diverge to a greater extent from those that would be seen at scale.

Table 5: Implications for Publication Bias

	Share detectable	At-scale effect (non-detected)	Share at-scale < local (detected)	% change at scale (detected)
All households	0.32	1.34	0.80	-0.12
Bottom decile	0.20	2.36	0.07	0.29
Top decile	0.49	0.88	0.98	-0.25
Market level	0.34	1.30	0.85	-0.09

Notes: Table presents results based on the roughly 100,000 household sample. “Share detectable” in the first column presents the fraction of households with welfare changes due to the local intervention ≥ 5 percentage points. The bottom row is based on market-level average treatment effects (fractions of market-level RCTs rather than households). The second column shows the average at-scale treatment effects among the “non-detected” households or markets. The third column shows the fraction of households or markets with “detectable” local effects that have smaller effects at scale. The final column shows the average percentage change in the effect at scale (relative to local) among the “detectable” cases.

Finally, we explore some potential implications for publication bias. In the first column of Table 5, we compute the fraction of household- or average market-level local treatment effects in our counterfactual analysis with (absolute) effect sizes greater or equal to 5 percentage points. One of the advantages of the model-based approach is that our counterfactuals isolate the pure effects of the single policy intervention, while keeping all other shocks affecting households and markets constant. Field experiments estimate these effects in a real-world scenario, and with frequently noisy survey data, such that defining “detectable” effect size at a cutoff of 5 percentage points should likely be viewed as a lower-bound.⁴⁰ We find only 32% of the roughly 100,000 households experience real income gains under the local intervention that would be detectable above that threshold and 34% when looking at market-level average treatment effects.

estimates in Figure 4 suggest that concerns of GE spillovers on control units may be overstated, at least at the scale at which most local agricultural interventions are run.

⁴⁰E.g., based on estimated standard errors on treatment effects observed in the AEA RCT database we built, the median study would be powered to detect increases in harvests of about 13% (with few studies reporting estimates on household welfare).

However, column 2 suggests that one would be wrong to conclude that the remaining two-thirds of potential study sites would experience zero effects if the policy were to be scaled up: the average at-scale household benefit computed among this group in column 2 is 1.3 percentage points, with roughly one third of those households experiencing gains of more than 2 percentage points, suggesting meaningful impacts at-scale. In a setting where statistical zeros have lower propensities for eventual publication, columns 1 and 2 would thus indicate that at-scale benefits could be obscured due to the combination of imprecise estimation and publication bias against insignificant results.

Next, we turn attention to the “detectable” effects that exceed real income changes of 5 percentage points in the local intervention. According to our analysis in column 3, 80% of the estimated effects in a local field intervention would over-estimate household benefits compared to the same intervention at scale, and 85% when using market-level average treatment effects. The average estimated local effect among this subset of markets and households is roughly 8.5%, so that the change in impact at-scale displayed in column 4 implies that the average results from the detectable local interventions would be over-estimated by slightly more than 1 percentage point.

The impact of publication bias is not symmetric across land-poor vs land-rich households, as documented in the second and third rows of Table 5. Because local interventions have on average more positive effects among the land-rich, almost 50% of the land-rich households would have their effects detected in local interventions, but only 20% of the land-poor. However, in line with the asymmetric changes in impacts at scale we documented above, land-poor households with non-detectable local effects would experience on average more than 2.3 percentage point welfare gains from the at-scale intervention. These findings imply that publication bias could hinder learning about policy effectiveness disproportionately among land-poor households. In columns 3 and 4, we then focus on households and markets that would be less likely subject to publication bias. Here, the estimated local effects would under-state the benefits of input subsidies among the land-poor, and they would over-state the benefits among the land-rich, compared to the intervention at scale. They do so by large magnitudes: column 4 suggests that on average local interventions would under-state the benefits among the land-poor by roughly 30% and over-estimate the gains among the land-rich by on average 25%.

Our review of experiments listed in the AEA registry suggests that concerns around

publication bias are likely empirically relevant for the existing body of agricultural RCTs. Even when we restrict attention to studies launched over five years ago, only 57% of studies are published. Moreover, we find that the factor identified in Table 4 as most predictive of larger local effects – the share of modern inputs used at baseline – is also predictive of publication status, suggesting that experiments with larger local effects may indeed be more likely to be published (see Appendix Table A.7).

7 Sensitivity, Validation, Extensions and Limitations

In this final section, we present additional sensitivity analysis and model validation. We conclude with a discussion of some important limitations of our approach as laid out above and model extensions that can help address some of these in different research contexts.

Sensitivity Analysis and Validation

In Section 3, we emphasized the role that local experiments can play in rigorous identification of key elasticities. The more sensitive 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 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 important to identify accurately with credible approaches.

We quantify 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 Appendix Figure A.7. We find that while results are less sensitive to the value of the demand elasticity σ and the upper-tier supply elasticity μ (across crops), the magnitude of the lower-tier supply elasticity (across traditional vs. modern technology), κ , is quite important for our estimates.⁴¹ 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 in-

⁴¹In our setting, estimates may be less sensitive to σ and μ because cost shares of modern inputs do not differ substantially across crops, with the exception of maize. How households trade off these crops is therefore less critical for the changes in the policy's impact locally vs at scale. However, in other contexts (e.g. with an intervention targeted at one particular crop), both σ and μ could play more important role at scale.

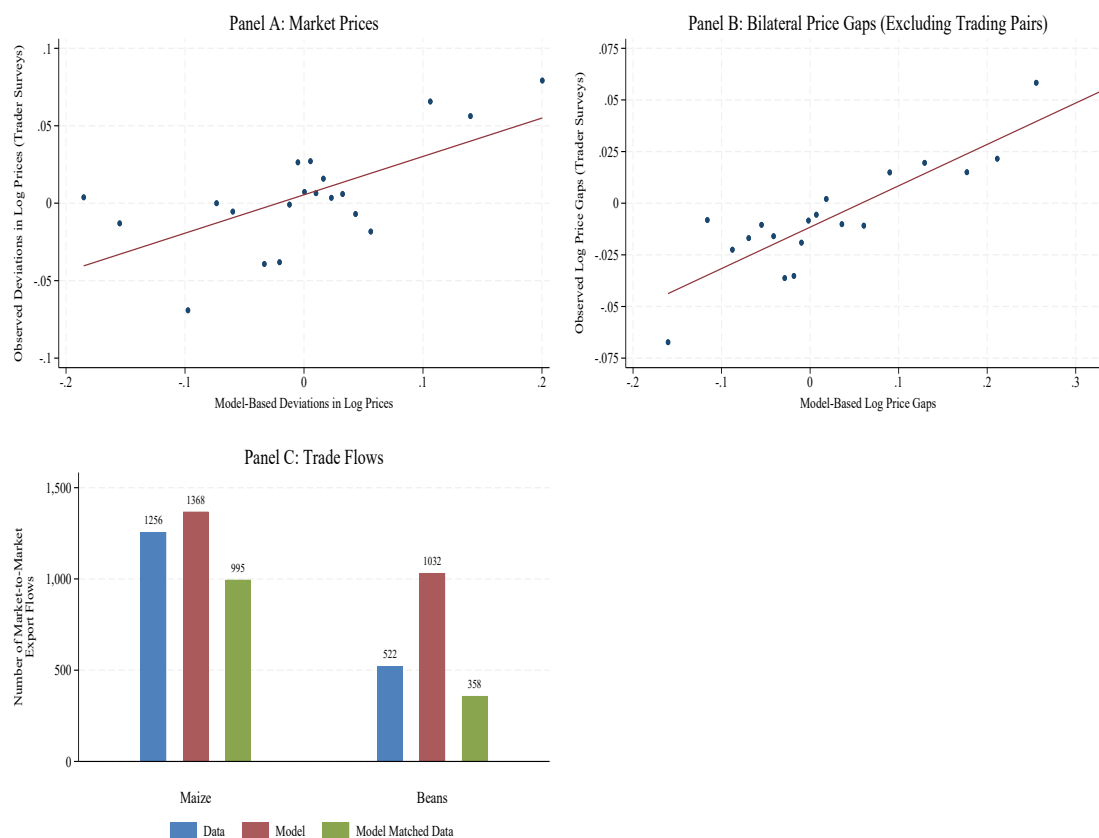
intervention in GE, as greater responsiveness on the part of others leads to larger output and factor price changes at scale compared to local intervention (at original prices). This highlights the importance of careful identification of this parameter. Using exogenous variation in prices coming from 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.

One additional sensitivity check concerns our assumptions about the tradability of certain crops. For example, matooke, a local banana variety, is not commonly traded across the Ugandan border with other countries. Looking at Uganda's trade statistics, sorghum, sweet potatoes and cassava have similar or lower shares of exports to domestic production. A natural question is how important the assumption of cross-border trading is for the counterfactual analysis. In Appendix Figure A.8 we assess this sensitivity, presenting side-by-side results from our baseline model and an alternative version in which we set cross-border trading frictions for these four crops to prohibitively high levels, such that Uganda remains effectively in autarky for them. Reassuringly, we find very similar results, with the gap in at-scale effects now slightly stronger than in our baseline calibration. This makes sense following the discussion of the drivers above, as higher trade costs (more remoteness) increase the scope for GE price effects at scale, since market prices are less constrained (pinned-down) by world market prices across the border.

Beyond sensitivity analyses, 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 to use data on crop prices and trade flows between 260 Ugandan markets in the trader surveys collected by Bergquist et al. (2024). Comparing these marketplaces in our 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 9 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

Figure 9: Model Validation Using Price Data and Trade Flows from Trader Surveys



Notes: Panel A presents a binned scatter plot of residualized log median crop-by-market prices (y-axis) from the trader survey data on residualized crop-by-market log prices in the model. Panels B and C make use of additional trade flow data between the 260 markets for maize and beans. In Panel B, we convert the market price dataset from Panel A into bilateral price pairs (counting each pair only once per month and crop), excluding active trading pairs. Panel C compares the trade flows reported across the markets in the trader surveys for each month of data to the bilateral trade flows from the models' price discovery algorithm.

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.⁴²

⁴²There are, of course, many reasons why price deviations can differ in data vs. model. On the survey data side, there could be measurement error, unobserved variation in crop quality, or temporary shocks on the day that information was collected across different markets. On the model side, household locations, expenditure

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 predicted by the 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.

Panel C of Figure 9 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 995 active trading relationships (79%). For beans, the model predicts 69% of the observed bilateral trading relationships (358/522). 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 (73% for maize and 35% for beans). One explanation for this discrepancy is that the trader surveys are based on a sample of traders at a given point in time (i.e., we might miss individual traders serving specific routes in reality, especially for thinner markets such as for beans), whereas the model captures aggregate trade flows between markets. Another possibility is that trade costs in our model are too stylized (i.e., linear in distance) and do not account for various real-life features in transportation such as fixed costs, empty back-haul, or market sequencing. Hence, our model predicts many small trade flows that might not be profitable in reality because of these factors. As we discuss below, we regard the introduction of such features in the model as interesting avenues for future research. That being said, we view the high proportion of observed bilateral trading relationships that the model correctly predicts as a further reassuring piece of evidence that the model-based price discovery algorithm reveals meaningful economic variation across

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 marketplaces 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.

markets.

Finally, we can also use the AEA RCT database that we built to conduct some (arguably ambitious) out-of-sample validation exercises. For each study, we collect information about the intervention type, location, and treatment effects on harvests, agricultural revenues, and total revenues when available. As discussed above and in [Appendix 4](#), we add information from a variety of other sources, including data from the FAO GAEZ database about local agricultural practices. Using this information, we can compare the drivers of variation in the local intervention’s treatment effect that we estimate in our quantitative model of the Ugandan input subsidy application to those we observe across the realized field experiments. We note upfront that this exercise suffers from two major caveats: (1) substantial heterogeneity in setting and intervention details and (2) small sample sizes.⁴³

With these caveats in mind, in [Appendix Figure A.9](#) we overlay the estimated treatment effects in the AEA registry studies and those estimated in our model, as a function of the factor found to be the main driver of the local effect size in our variance decomposition in [Table 4](#): the initial farmer revenue share under the modern production technique. We do this for all registered agricultural interventions with reported treatment effects for each outcome, as well as restricted to experiments studying the same specific intervention we study (input subsidies). Given the small sample sizes, the confidence intervals around the AEA registry data are large (even when using point estimates of treatment effects as data points, ignoring their own confidence intervals). Having said this, reassuringly we find that our model’s estimates fall within these confidence intervals. Moreover, in both our model and in the AEA registry data, we find a positive relationship between treatment effects and the share of modern production techniques used. Such out-of-sample validation is, of course, quite tentative. Nevertheless, this meta-analysis adds some further indication that the economic forces captured in our model appear to be present across multiple different empirical contexts.

Limitations, Extensions, and Avenues for Future Research

We now turn to some important limitations of our approach, as well as possible extensions to address some of these depending on the policy context. We conclude by highlighting

⁴³Of the 143 studies identified in the review, only 115 have working papers and among these, only 53 report harvest impacts, 37 agricultural revenue impacts, and 30 total household revenue impacts; further, only one-third of experiments specifically study input subsidies, with only 21 reporting harvest impacts, 16 agricultural revenue impacts, and 11 total household revenue impacts.

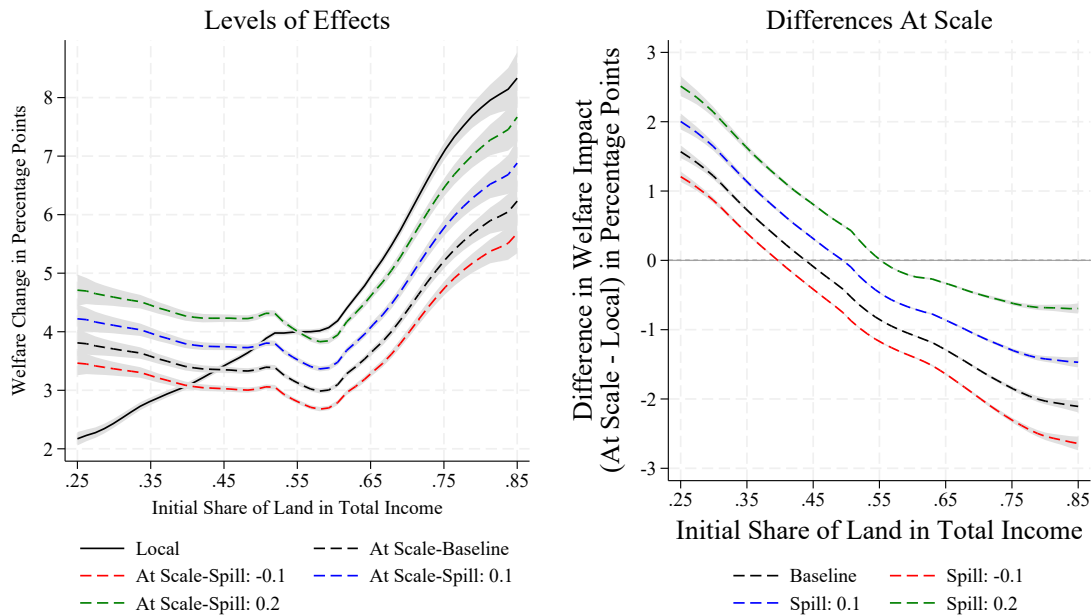
avenues for future research.

One important limitation of our baseline model is that it does not allow the effect of the subsidy on take-up or land productivity to vary with the scale of implementation. However, there can be scenarios in which the effectiveness of treatments either increases or decreases as a function of the scale of policy implementation: for example, complementarities at scale could arise due to learning from others in the village (Beaman et al. (2021), Chandrasekhar et al. (2022)) or shifting norms; conversely, changes in the implementation protocols or oversight when scaling up could imply lower take-up or impact (Muralidharan and Niehaus, 2017). We follow recent advances from modeling agglomeration and congestion effects in spatial economics (e.g., Redding (2022)) to extend our framework in the following way. We modify the production function of the modern technique by introducing external returns to scale of modern technology adoption within the local market (parish). Specifically, we posit that the productivity shifter $a_{i,k,\omega}$ for farmer i producing crop k using technique ω , introduced in equation (1), follows $\bar{a}_{i,k,\omega} Q_{m,k,\omega}^\varsigma$ for $\omega = \{\text{modern}\}$ and $\bar{a}_{i,k,\omega}$ for $\omega = \{\text{traditional}\}$, where $Q_{m,k,\omega} = \sum_{i \in \mathcal{J}(m)} q_{i,k,\omega}$ is the total output of crop k produced by all farmers in village m using the modern technique, and $\bar{a}_{i,k,\omega}$ is an exogenous productivity shifter. The parameter $\varsigma > 0$ (< 0) is the agglomeration (congestion) elasticity. When $\varsigma > 0$, then the productivity of a farmer using the modern technique increases with the scale of local technology adoption, capturing potential at-scale complementarities, for example, through learning or shifting norms. When $\varsigma < 0$, on the other hand, the productivity of the farmer using modern inputs decreases as the scale of local technology use rises, allowing, for example, for decreased implementation efficiency at-scale.

Figure 10 presents the levels of the local and at-scale effects (left panel) and the differences in at-scale vs. local effects (right panel) across alternative assumptions about the extent of congestion or agglomeration that arise at scale, ranging from a congestion force of $\varsigma = -0.1$ to complementarities within the range of existing estimates of agglomeration forces, $\varsigma = 0.1$ and $\varsigma = 0.2$. Given the small fraction of farmers affected, the level of the local intervention effects (plotted in the solid line in the left panel) is not affected by the presence of congestion or agglomeration forces. At scale, we see that the effects are roughly shifted in parallel across the land income distribution, negatively compared to our baseline estimates in case of congestion and positively in presence of agglomeration. As a result, the differences in the effect at-scale are shifted in the same directions in the right panel. For

sufficiently large complementarities at scale, this can push average impacts at-scale to be greater than those in the local intervention.

Figure 10: Model Extension with At-Scale Complementarities



Notes: Figure plots local polynomial regressions of the welfare effect in percentage points (left panel) or the difference the effect (at scale - local) (right panel) as a function of initial household land income shares. Results from the baseline model are in black. The remaining graphs plot the same relationships in model extensions that allow for local congestion externalities at scale (red), and different degrees of positive agglomeration effects at scale (blue and green). Shaded areas indicate 95% confidence intervals.

Model extensions like the one above can serve as a reference for policy settings in which implementation at scale alters treatment effectiveness. Another future extension would be to relax some of the simplifying assumptions we make when taking the model to the data: for instance, the assumption that bilateral trade frictions are a linear function of distance on the existing road network in Section 4 could explain some of the unmatched parts of the validation exercises on trade flows that we discuss above. It clearly leads to an approximation of trade costs that we impose across bilateral connections in Uganda, relative to actual trading frictions that may involve route-specific fixed costs and considerations about route sequencing and dynamic costs due to empty backhauling. The increasing availability of more granular and high-frequency transportation data in LMICs provides promising avenues for further improving the calibration of a granular economic geography in different

contexts.⁴⁴

Lastly, in terms of types of policies, our approach is most directly tailored to three common types of interventions: (1) shocks to agricultural productivity ($\hat{b}_{j,k,\omega}$) due to, e.g., climate change, new seed varieties, irrigation technology or input subsidies as we lay out above; (2) demand-side shocks, including cash transfers, other income shocks (including those in cities or other regions) or changes in preferences due to nutritional information campaigns ($\hat{a}_{j,k}$); and (3) policies affecting trade costs ($\hat{\tau}_{od,k}$ and $\hat{t}_{od,k}$ in our model), such as road building or trade reforms. Of course, there is a range of agricultural policies for which our current model would need to be tailored or extended to speak to effects at-scale. For example, an important set of policy counterfactuals that our current model would need to be extended for is land market reforms such as those that title land rights (e.g. [Ali et al. \(2015\)](#)). We abstract from trade in land in Section 2, in part because land markets have been found to be generally thin in our empirical setting ([Acampora et al., 2022](#), [Holden et al., 2010](#)). But this market would be straight-forward to incorporate as an additional module to analyze policy counterfactuals aimed at land reform. Our current (static) model would also need to be modified to be applied to dynamic policies that aim to reduce risk (e.g. [Donovan \(2021\)](#)) or alleviate the impact of inter-temporal shifts in preferences ([Duflo et al., 2011](#)) or prices ([Burke et al., 2019](#)).

We consider the approach that we develop here as a first step to unlock quantitative analysis of the at-scale impacts of agricultural policies evaluated under local interventions, with these and other extensions as promising avenues for future research in this area.

References

- Acampora, Michelle, Lorenzo Casaburi, and Jack Willis**, “Land Rental Markets: Experimental Evidence from Kenya,” *Working Paper*, 2022.
- Ali, D., K. Deininger, M. Goldstein, and E. La Ferrara**, “Impact of the land tenure regularization program in Rwanda,” *World Bank Development Research Group Case Study*, 2015.
- Atkin, David and Dave Donaldson**, “Who’s Getting Globalized? The Size and Implications of Intra-national Trade Costs,” *NBER Working Paper*, 2015, (21439).

⁴⁴E.g., recent work on India’s road transportation by [Barnwal et al. \(2024\)](#) make use of high-frequency GPS records, trucking logs, and a survey of truck drivers.

- Baird, S., C. McIntosh, and B. Ozler**, “Cash or Condition? Evidence from a Cash Transfer Experiment,” *Quarterly Journal of Economics*, 2011, 126, 1709–1753.
- Barnwal, Prabhat, Jonathan I Dingel, Daniil Iurchenko, Pravin Krishna, and Eva Van Leemput**, “Internal Trade Barriers in India,” *Manuscript, Columbia University*, 2024.
- Beaman, Lori, Ariel BenYishay, Jeremy Magruder, and Mushfiq Mobarak**, “Can Network Theory-based Targeting Increase Technology Adoption?,” *American Economic Review*, 2021, 111 (6), 1918–1943.
- Becker, Gary S.**, “A Theory of the Allocation of Time,” *The economic journal*, 1965, 75 (299), 493–517.
- Bergquist, Lauren Falcao and Michael Dinerstein**, “Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya,” *American Economic Review*, 2020, 110 (12), 3705–47.
- , **Craig McIntosh, and Meredith Startz**, “Search Costs, Intermediation, and Trade: Experimental Evidence from Ugandan Agricultural Markets,” *Working Paper*, 2024.
- Berry, Steven, Amit Gandhi, and Philip Haile**, “Connected substitutes and invertibility of demand,” *Econometrica*, 2013, 81 (5), 2087–2111.
- Bouguen, Adrien, Michael Yue Huang, Michael Kremer, and Edward Miguel**, “Using RCTs to Estimate Long-Run Impacts in Development Economics,” *Annual Review of Economics*, 2019, 11, 523–561.
- Burke, Marshall, Lauren Falcao Bergquist, and Edward Miguel**, “Sell Low and Buy High: Arbitrage and Local Price Effects in Kenyan Markets,” *Quarterly Journal of Economics*, 2019.
- Carter, Michael, Rachid Laajaj, and Dean Yang**, “Subsidies and the African Green Revolution: Direct Effects and Social Network Spillovers of Randomized Input Subsidies in Mozambique,” *American Economic Journal: Applied Economics*, 2020.
- Castaneda, Andres, Dung Doan, David Newhouse, Ming Cong Nguyen, Hiroki Uematsu, and Joao Pedro Azevedo**, “Who are the poor in the developing world?,” *World Bank Policy Research Working Paper*, 2016, (7844).
- Chandrasekhar, Arun G, Esther Duflo, Michael Kremer, João F Pugliese, Jonathan Robinson, and Frank Schilbach**, “Blue spoons: Sparking communication about appropriate technology use,” Technical Report, National Bureau of Economic Research 2022.
- Costinot, Arnaud and Andrés Rodríguez-Clare**, “Trade theory with numbers: Quantify-

- ing the consequences of globalization,” 2014, 4, 197–261.
- **and Dave Donaldson**, “How Large are the Gains from Economic Integration? Theory and Evidence from US Agriculture 1880-1997,” *NBER Working Paper*, 2016, (22946).
- de Janvry, Alain and Elisabeth Sadoulet**, *Quantitative Development Policy Analysis*, The John Hopkins University Press, 1995.
- Dekle, Robert, Jonathan Eaton, and Samuel Kortum**, “Unbalanced Trade,” *The American Economic Review*, 2007, pp. 351–355.
- Donovan, Kevin**, “The Equilibrium Impact of Agricultural Risk on Intermediate Inputs and Aggregate Productivity,” *Review of Economic Studies*, 2021.
- Duflo, Esther, Michael Kremer, and Jonathan Robinson**, “Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya,” *American Economic Review*, 2011, 101 (6), 2350–90.
- Egger, Dennis, Johannes Haushofer, Edward Miguel, Paul Neihaus, and Michael Walker**, “General Equilibrium Effects of Cash Transfers: Experimental Evidence from Kenya,” *Econometrica*, 2022.
- Emerick, Kyle and Marshall Burke**, “Adaptation to climate change: Evidence from US Agriculture,” *American Economic Journal: Economic Policy*, 2016.
- Farrokhi, Farid and Heitor S Pellegrina**, “Trade, technology, and agricultural productivity,” *Forthcoming, Journal of Political Economy*, 2022.
- Fink, Gunther, B Kelsey Jack, and Felix Masiye**, “Seasonal liquidity, rural labor markets, and agricultural production,” *American Economic Review*, 2020, 110 (11), 3351–92.
- Head, Keith and Thierry Mayer**, “Gravity equations: Workhorse, toolkit, and cookbook,” in “Handbook of international economics,” Vol. 4, Elsevier, 2014, pp. 131–195.
- Holden, Stein T, Keijiro Otsuka, and Frank M Place**, “The emergence of land markets in Africa: Impacts on poverty, equity, and efficiency,” 2010.
- Hummels, David**, “Transportation costs and international trade in the second era of globalization,” *Journal of Economic perspectives*, 2007, 21 (3), 131–154.
- Huntington, Heather and Ajay Shenoy**, “Does Insecure Land Tenure Deter Investment? Evidence from a Randomized Controlled Trial,” *Journal of Development Economics*, 2021.
- IFDC**, “Uganda Fertilizer Assessment,” *Alabama: International Fertilizer Development Center*, 2014.

- Jayne, T.S. and Shahidur Rashid**, “Input subsidy programs in sub-Saharan Africa: a synthesis of recent evidence,” *Agricultural Economics*, 2013, 44, 547–562.
- Muralidharan, Karthik and Paul Niehaus**, “Experimentation at Scale,” *Journal of Economic Perspectives*, 2017, 31 (4), 103–24.
- Porteous, Obie**, “High trade costs and their consequences: An estimated dynamic model of African agricultural storage and trade,” *American Economic Journal: Applied Economics*, 2019, 11 (4), 327–366.
- Powell, Alan A and FHG Gruen**, “The constant elasticity of transformation production frontier and linear supply system,” *International economic review*, 1968, 9 (3), 315–328.
- Redding, Stephen J**, “Trade and geography,” *Handbook of International Economics*, 2022, 5, 147–217.
- Sotelo, Sebastian**, “Domestic trade frictions and agriculture,” *Journal of Political Economy*, 2020, 128 (7), 2690–2738.

Scaling Agricultural Policy Interventions: Online Appendix

[Appendix 1](#) describes the datasets 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 referenced in the main text and in the stylized facts below. [Appendix 4](#) describes the data collection and construction of the AEA Registry data on agricultural RCTs, as well as additional results using the data. [Appendix 5](#) lays out additional model details and proofs. Finally, the [Technical Appendix](#) provides a more general version of the model that can be adapted to other contexts.

Appendix 1 Data

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.

The UNPS farmer-level data is organized in plots. Crops produced (either single stand or intercropped), seed variety (quantity and cost), fertilizer (quantity and cost), pesticides (quantity and cost), labor (days of HH labor and hired labor, and cost of hired labor) and land (size and value) are reported for each plot and season. The crops we consider are matooke (banana), beans, cassava, coffee, groundnuts, maize, millet, sorghum, and sweet potatoes. When a single crop is grown on a plot, we assign the plot and all inputs directly to a crop. In the case of more than one crop on a given plot, UNPS reports the share of the plot allocated to each crop. We allocate all inputs to each crop according to this share. However, we estimate input shares in production using only single stand crops (see below). We consider a plot-crop as modern if at least one of the modern inputs is used: modern seed varieties, chemical fertilizers, or pesticides.

From 2009 forward, the survey reports quantities of harvest and sales as well as the value

and point of sale for each plot and crop. Harvest and sales quantities are reported in different units (i.e., sack, kg, container, bunch) and different states/conditions (i.e., dried, wet, in the shell). In many cases, the survey reports conversion factors for various units to kg. For each unit, we assign the mode of these conversion factors to all observations. Unfortunately, no conversion is provided for states/conditions by the UNPS. To bring quantities to a common state for each crop, we manually construct conversion factors using scientific publications from the FAO, USDA and HarvestPlus. The exact sources and conversion factors are available in the replication package and are described in the Readme file.

For our analysis, we construct the following variables from the UNPS data:

- Land allocation $\pi_{i,k,\omega,t}$: We aggregate the land sizes of all plots (or shares of plots for intercropped plots) which are used for a given crop-technique k, ω , across seasons in a year t . We divide by the total amount of land under all crops and techniques cultivated by farmer i .
- Harvest $q_{i,k,\omega,t}$: We sum output in kg and a common state over all plots, which farmer i uses to produce a crop-technique pair k, ω , across seasons in a year t .
- Intermediates $m_{i,k,\omega,t}$: We add expenditures on modern seed varieties, chemical fertilizers, and pesticides for each crop-technique pair k, ω , across seasons in a year t .
- Labor input $l_{i,k,\omega,t}$: We add days of labor of household members and hired labor for each crop-technique pair k, ω , across seasons in a year t .
- Expenditure Shares $\xi_{i,k}$: UNPS includes consumption expenditures and quantities for a large set of goods. Non-durable consumption including food is reported for the past week, semi-durables for the past month, and durables for the past year. We adjust these categories to represent annual expenditures. We manually link food expenditure to the 9 main crops. The expenditure share on food $\xi_{i,A}$ is constructed as the sum of all food expenditure over total expenditures. The expenditure shares of each crop are constructed for all expenditures related to a crop k (i.e., maize flour, maize kernels) over total expenditure on the 9 main crops.
- Shares of intermediates $\alpha_{i,k,\omega}$ and labor $\beta_{i,k,\omega}$ in production: We construct costs of land, labor and intermediates for each farmer, crop and technology using only plots that are exclusively occupied by a single crop. Land costs are constructed using rents for a parcel reported by the farmer and when not available, we assign the average land

rent per acre in the district in a given season multiplied by the parcel size. Labor costs are constructed from days worked by household members and hired workers. When a farmer hires workers to work on a plot, the total cost of hired labor is reported. We construct daily wages from the average daily wage of hired workers in the district and season. To calculate total labor costs, we assign this daily wage to all days worked by household members in a season and add the costs of hired labor if applicable. Costs of intermediates (modern seed varieties, chemical fertilizer, pesticides) are directly reported. We then calculate cost shares by dividing the cost of each category by the sum of all categories. Finally, we take averages for each crop and technique in the four regions of Uganda.

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 at 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.

The following variables are constructed from the census data:

- Asset index: We add dummies for ownership of a bicycle, motorcycle, other transport equipment, and mobile phone.
- Labor endowment L_i : We count all household members aged 15-65 as the labor endowment of the farm household.
- Urban income I_h : We calculate urban income as the average total consumption expenditure of urban households in the UNPS. We then multiply this average with the number of households in each city in Uganda.

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 comple-

ment this database with geo-referenced information on crop suitability from the Food and Agricultural Organization (FAO) Global Agro-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. \(2024\)](#) 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. \(2024\)](#).

Demand Estimation To estimate the slope of the demand curve for crops, we bring to bear transaction-level microdata from maize markets in rural Kenya that was collected as part of an experiment in [Bergquist and 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 and 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 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. [Figure 1](#) provides a map of the Ugandan geography we use in our counterfactual analysis.

Product Differentiation Across Farmers [Appendix Table A.1](#) 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. \(2024\)](#) 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.1](#) 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. \(2024\)](#). Specifically, we estimate:

$$t_{ijkt} = (p_{jkt} - p_{ikt}) = \alpha + \beta p_{ikt} + \theta_{ij} + \phi_t + \epsilon_{ijkt}$$

where t_{ijkt} are per-unit trade costs between origin i and destination j for crop k (maize or beans) observed in month t , p_{ikt} are origin unit prices, θ_{ij} are origin-by-destination fixed effects, and ϕ_t are month fixed effects. Alternatively, origin-by-destination-by-month fixed effects (θ_{ijt}) can be included. Following [Bergquist et al. \(2024\)](#), 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_{ikt} 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.2](#), 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 and 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.1](#) reports a non-parametric estimate of the household Engel curve for food consumption. We estimate flexible functional forms of the

following specification:

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

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_m , 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.3, 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

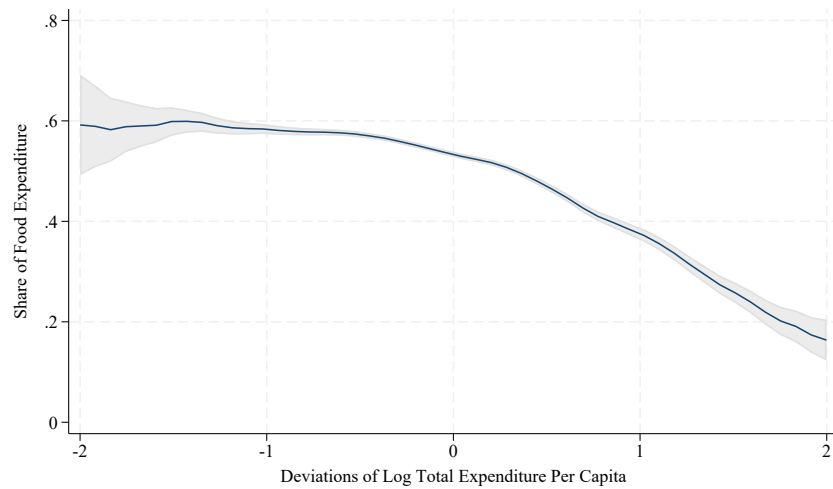
Table A.1: Product Differentiation – Missing Trade Flows

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

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

See Appendix 2 for discussion and Appendix 1 for description of the data. *** p<0.01, ** p<0.05, * p<0.1

Figure A.1: Household Preferences (Non-Homotheticity)



See Appendix 2 for discussion and Appendix 1 for description of the data.

Table A.2: 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*** (0.0188)	-0.0419** (0.0206)	-0.0082 (0.0256)	-0.0002 (0.0274)
Observations	8,522	8,428	7,151	7,077
Pair FX	yes	.	yes	.
Month FX	yes	.	yes	.
Pair-by-Month FX	no	yes	no	yes

Standard errors clustered at level of bilateral pairs.

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

See Appendix 2 for discussion and Appendix 1 for description of the data.

Table A.3: Technology Adoption and Production Cost Shares

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

Standard errors clustered at level of farmers.

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

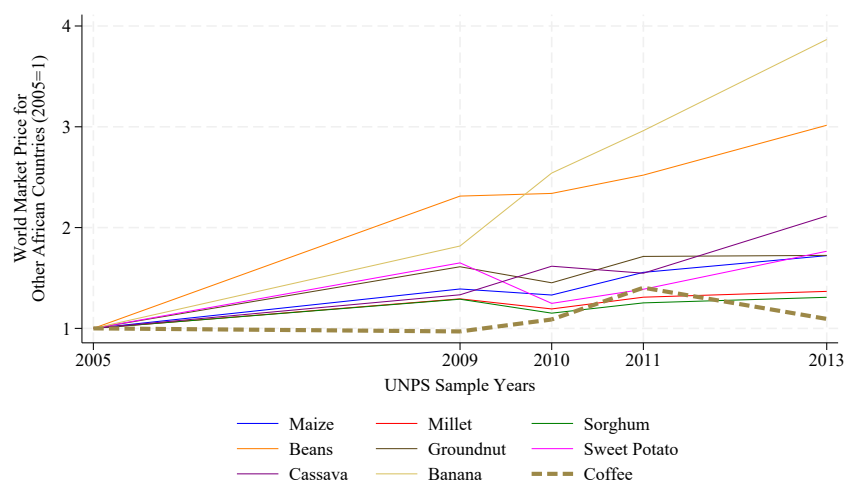
See Appendix 2 for discussion and Appendix 1 for description of the data.

Table A.4: Calibrated Cost Shares in Production

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

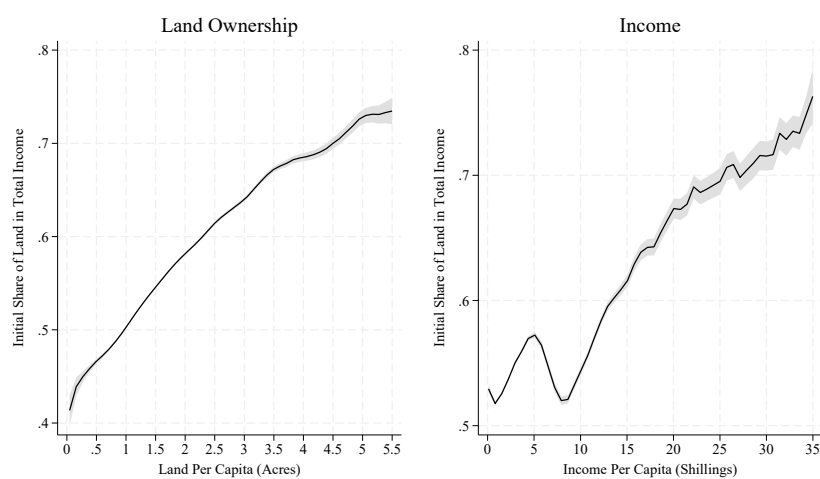
See Section 4 for discussion and Appendix 1 for description of the data.

Figure A.2: Relative World Price Changes Over the Sample Period



See Section 4 for discussion.

Figure A.3: Land Income Shares, Land Ownership and Household Incomes



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

Table A.5: 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	4.32*** (0.06)	3.55*** (0.07)	3.00*** (0.06)	3.45*** (0.08)	3.92*** (0.07)	2.94*** (0.08)	6.35*** (0.11)	4.66*** (0.10)
Observations	104,361	104,361	19,829	19,829	19,828	19,828	20,872	20,872
No Clusters	4502	4502	3578	3578	4133	4133	4088	4088

Standard errors clustered at market-level.

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

See Section 5 for discussion.

Table A.6: Channels

Panel A: Local Effects

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Income Local	Wage Local	P_manu Local	P_banana Local	P_bean Local	P_cassava Local	P_coffee Local	P_groundnut Local	P_maize Local	P_millet Local	P_sorghum Local	P_sweetpot Local
Effect	4.2165*** (0.0642)	0.6182*** (0.0152)	0.0000 (0.0000)	-0.0467*** (0.0038)	-0.0359*** (0.0026)	-0.0188*** (0.0062)	-0.0022*** (0.0003)	-0.0375*** (0.0085)	-0.5831*** (0.0198)	0.0247*** (0.0021)	0.0088*** (0.0009)	0.0931*** (0.0036)
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

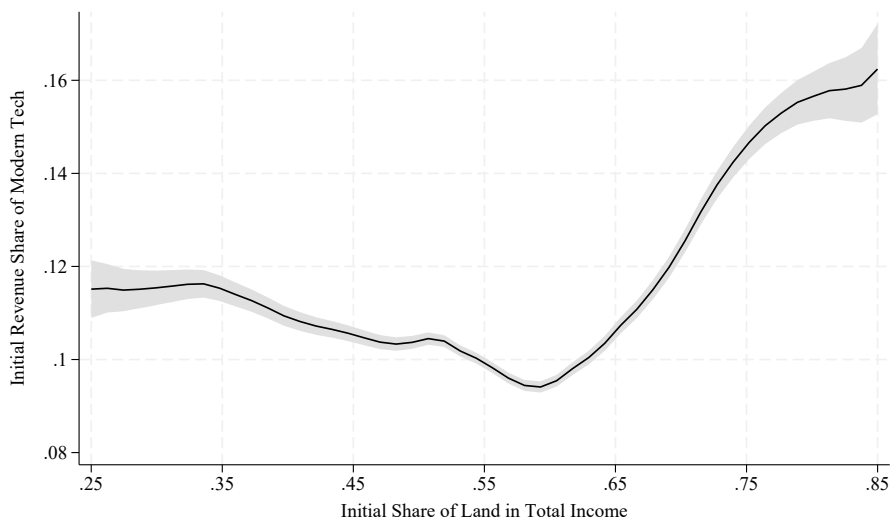
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Income At Scale	Wage At Scale	P_manu At Scale	P_banana At Scale	P_bean At Scale	P_cassava At Scale	P_coffee At Scale	P_groundnut At Scale	P_maize At Scale	P_millet At Scale	P_sorghum At Scale	P_sweetpot At Scale
Effect	3.2518*** (0.0694)	2.6734*** (0.0563)	0.8544*** (0.0032)	-0.1049*** (0.0055)	-0.2559*** (0.0097)	-0.3109*** (0.0309)	-0.0114*** (0.0005)	0.1626*** (0.0137)	-4.8461*** (0.0549)	0.1466*** (0.0071)	0.0297*** (0.0034)	0.6640*** (0.0114)
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

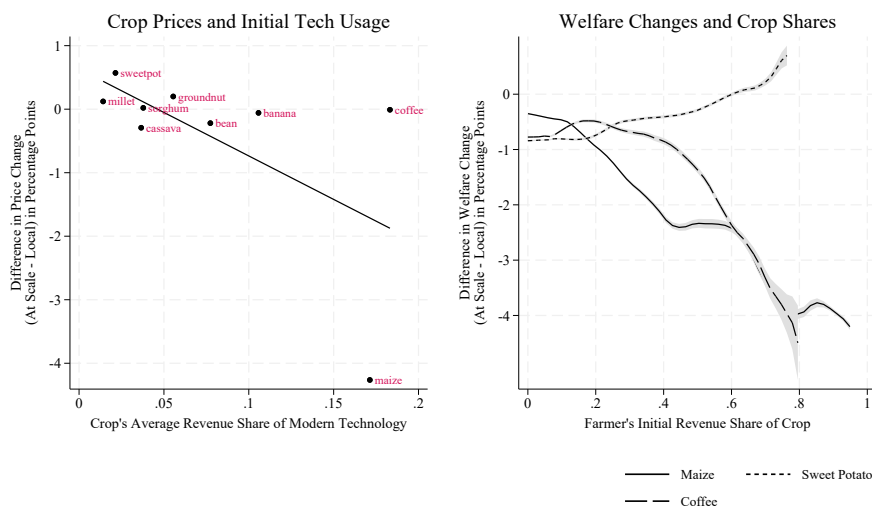
See Section 5 for discussion.

Figure A.4: Initial Usage of Modern Inputs Across Land-Poor vs Land-Rich Households



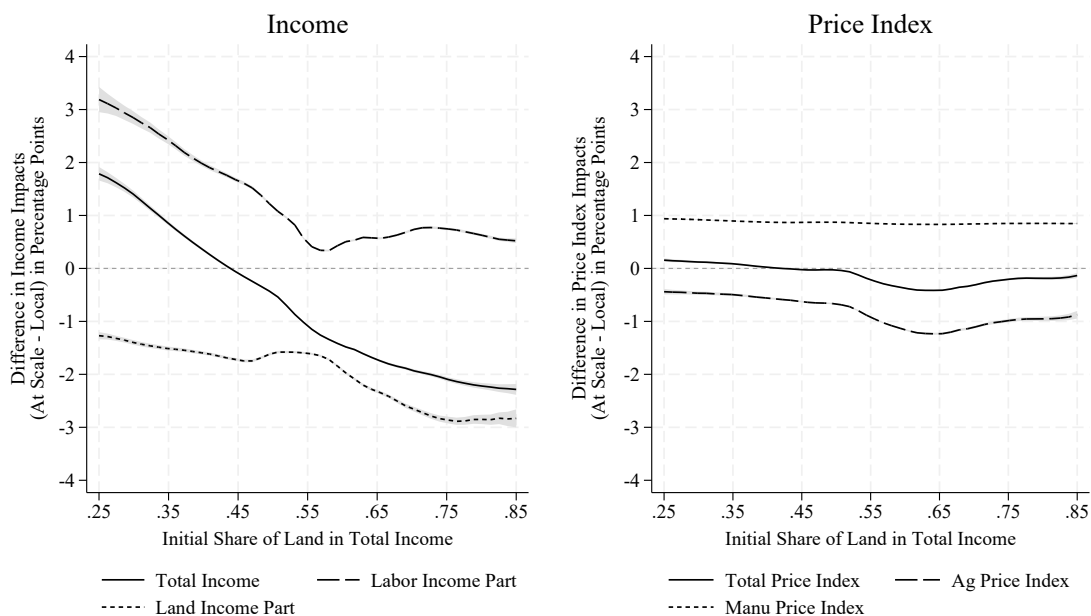
The figure plots estimates from local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 5 for discussion.

Figure A.5: Effects as a Function of Initial Crop Shares



Right panel plots local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 5 for discussion.

Figure A.6: 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 100,000 rural Ugandan households. Shaded areas indicate 95 percent confidence intervals. See Section 5 for discussion.

Figure A.7: Sensitivity to Alternative Parameters

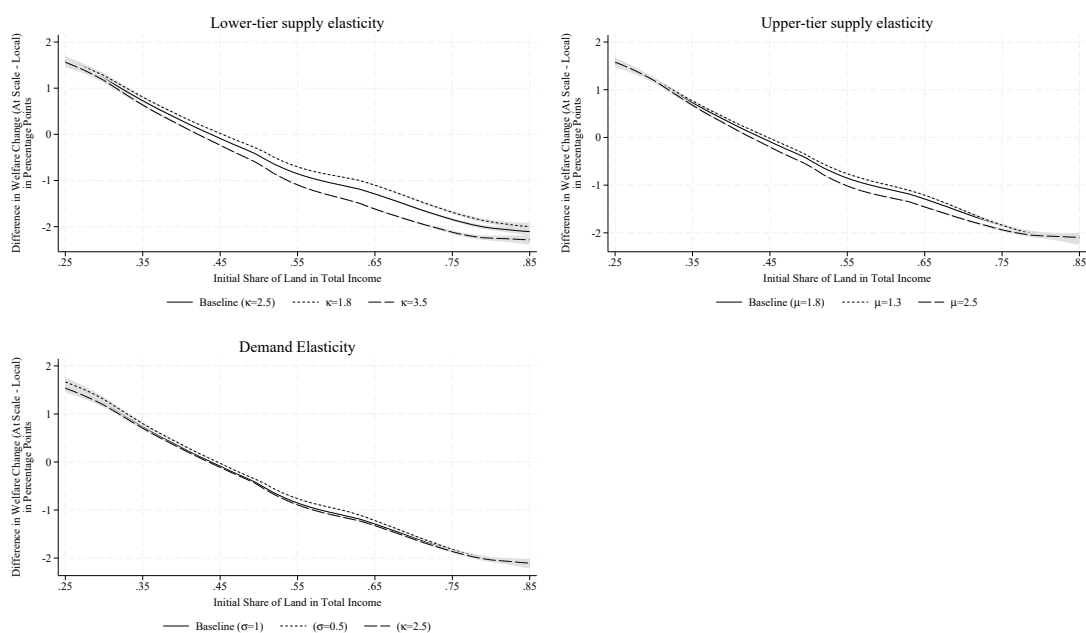


Figure plots local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 7 for discussion.

Figure A.8: Alternative Assumptions on Cross-Border Trading

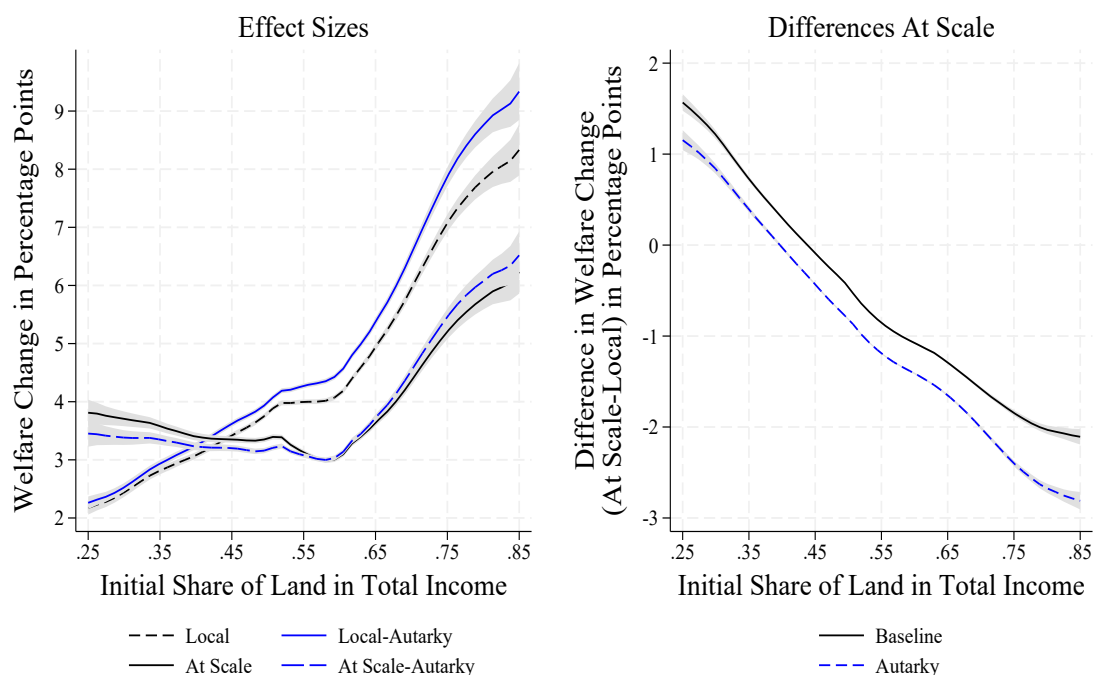


Figure plots local polynomial regressions of the welfare effect in percentage points (left panel) or the difference the effect (at scale - local) (right panel) as a function of initial household land income shares. Results from the baseline model are in black and from the alternative parametrization are in blue. Shaded areas indicate 95% confidence intervals. See Section 7 for discussion.

Appendix 4 AEA Registry of Agricultural RCTs

4.A Database Construction

To construct a database of agricultural RCTs conducted in Sub-Saharan Africa, we turn to the American Economic Association (AEA) RCT registry (www.socialscienceregistry.org/). We collect all studies that have been conducted in Sub-Saharan Africa and listed on the registry from the registry's launch in 2013 through 2024. To do so, we search the database for keyword "agricultur*" for each country in Sub-Saharan Africa. This yields a set of 181 trials. Trials are then manually reviewed by two independent reviewers to ensure they are related to agriculture. After discarding 38 trials that feature neither an agricultural intervention nor an agricultural outcome, we are left with 143 trials conducted across 23 countries. These comprise the studies in our dataset.

With the help of three RAs, we then manually code information for each study, drawing data from the information listed on the registry, attached pre-analysis plans, and working

papers or published papers identified in a web search including Google Scholar, authors' websites, and EconLit. When available, information collected on each study included: (1) intervention type; (2) study location; (3) working paper and publication status; (4) sample size; (5) randomization details; (6) coefficient and standard error of the treatment effect on harvests (in kgs), agricultural revenues, and total household revenues, as well as the mean of the variable in the population; (7) mean and standard deviation of demographic variables of the study population (land size, education, household expenditure) and (8) minimum detectable effects (MDEs). For the full codebook, please contact the authors. Note that several studies do not have working papers and, when they do, not all reported all outcomes above (more on this below).

The location data is collected at the finest level described in the text, which is almost always a town, district or other subnational administrative unit. For studies that only listed country or broad region within the country (e.g. "Eastern Uganda"), we contacted the authors to get a more precise location. We thank the many authors who replied to this request with detailed study location information.

Using the location data, we merge in several auxiliary data sources that provide data on market access, land share, maize share, and modern share. These data are drawn from four additional sources, merged with our RCT data based on the nearest grid cell or enumeration area centroid:

1) Market access data is constructed using population counts on 15-arc-minute grids from the Socioeconomic Data and Applications Center at NASA ([CIESIN, 2018](#)) and distance between each grid cell from Harvard WorldMap ([WorldMap, 2021](#)).

2) Land shares are constructed at the lowest possible enumeration area level using data on household consumption and non-farm income collected in nationally representative household surveys (Living Standards Measurement Surveys, for countries that conduct them, and other nationally representative households surveys manually collected from national statistics agencies, for countries that do not run LSMS surveys).

3) Data on maize shares and modern shares are constructed using spatial data at a 5 arc-minutes grid level from the Food and Agriculture Organization's Global Agro-Ecological Zones (GAEZ) v4 dataset. Data on maize production value and total production value allowed for the calculation of maize shares. Data on "achievement ratios", which GAEZ calculates by "comparing downscaled actual yields with agro-ecological attainable yield

simulated under high input/advanced management assumptions” across 22 major crops, are used as a proxy for modern share.

4) Additional demographic data is added on asset holdings, household size, and education and age of household head from Demographic and Health Surveys.

In addition, to facilitate comparison to the full population of all possible rural sites across the 23 countries studied in the AEA registry studies, a database is constructed with the same four variables collected for all grid cells (or enumeration areas) across all rural areas in the 23 countries. Grid cells with a population of greater than 300 people/square km are identified as “urban” (following the definition of the UN Statistical Commission (Commission, 2020)) and are dropped. This results in a drop of 2.5% of all grid cells.

4.B Summary statistics

First, we explore where the RCTs listed on the AEA Registry were conducted. Uganda, our country of focus, is the second most commonly studied country, accounting for 25 studies, second only to Kenya at 27 studies. In descending order, the next most common are: Ethiopia (14 studies), Ghana (13 studies), Malawi (12 studies), Mozambique (7 studies), Zambia (6 studies), Liberia (5 studies), and Rwanda (5 studies). All other countries had fewer than five registered studies.

Second, we explore intervention type, coding the interventions evaluated in the 143 studies into (non-exclusive) categories. We find that 42% study extension services or farmer training, 38% financial vehicles or services (including credit, savings, insurance, and cash transfers), and 36% free or subsidized inputs. While our paper shows that even these types of productivity-enhancing interventions can have strong GE effects at scale, we note it is much less common that experiments study interventions with more obvious GE implications, such as those that directly target output or input market access (including land and labor markets); these total only 17% of interventions in the registry. This may in part be driven by concern about GE effects arising even in the pilot and contaminating the control group.¹

4.C Publication bias

We use this database to study the role of publication bias among agricultural RCTs listed on the AEA registry. As discussed above, we collect data on whether each study has been published. Restricting attention to experiments that were launched more than five years

¹Only 4% of the studies reviewed does not fall into one of the above categories. These are mostly studies measuring the impact of nutritional information, food safety, or health campaigns.

ago (before 2020) to allow enough time to produce results, we find that 57% of studies are published.²

We then regress a dummy for being published on the same set of covariates explored in the variance decomposition in Table 4: land shares in income, modern shares of production, maize revenue shares, and remoteness. Results are presented below in Appendix Table A.7.

Table A.7: Predictors of Publication Status among RCTs in the AEA Registry

	Published
Land share	-0.09 (0.07)
Modern share	0.16* (0.08)
Maize share	0.01 (0.06)
Log (inverse) market access	-0.03 (0.08)
Observations	105
Mean of DV	0.59
R squared	0.08

Notes: Regression of indicator for whether the study has been published on land share, modern share, maize share, and remoteness (as proxied by log inverse market access). See text above for details and sources for each variable construction. Regression sample is restricted to experiments launched more than five years ago (before 2020) to allow time for publication.

Consistent with the results from the variance decomposition, which finds that the share of revenues produced under the modern production technique is the strongest predictor of the local effect, we find studies conducted among populations with greater modern shares are significantly more likely to be published, possibly because they have stronger local treatment effects. However, these populations have lower treatment effects at scale, suggesting again that publication bias may lead to an over-reporting of results in settings where treatment effects are likely to be greater in the local intervention than at scale – and an under-reporting of the opposite.

4.D Validation among AEA registered experiments

Finally, we can also use our AEA RCT database to conduct some (arguably ambitious) out-of-sample validation exercises. As discussed above, for each study, we collect information about treatment effects on harvests, agricultural revenues, and total revenues when available.

²84% have a working paper, consistent with a broader meta analysis conducted by [Hoces de la Guardia \(2025\)](#) which finds that 81% of all AEA registry entries have a working papers within eight years.

We can then compare the drivers of variation in the local intervention’s treatment effect that we estimate in our quantitative model of the Ugandan input subsidy application to those we observe across the realized field experiments.

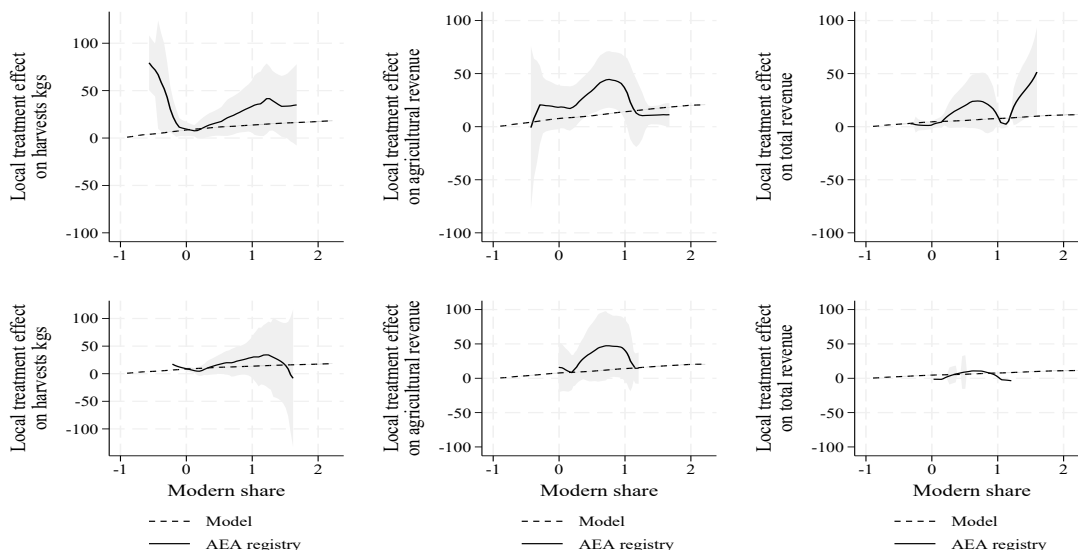
We note upfront that this exercise suffers from two major caveats: (1) substantial heterogeneity in setting and intervention details (only one-third of experiments specifically study input subsidies) and (2) small-sample size (of the 143 studies identified in the review, only 115 have working papers and, among these, only 53 report harvest impacts, 37 agricultural revenue impacts, and 30 total household revenue impacts).³

With these caveats in mind, in Appendix Figure A.9 we overlay the estimated treatment effects from the AEA registry studies and those estimated in our model, as a function of the factor found to be the main driver of the local effect size in our variance decomposition in Table 4: the initial farmer revenue share under the modern production technique. We do this for all registered agricultural interventions with reported treatment effects for each outcome as shown in the top row of Figure A.9, while the bottom row plots a version restricted to those experiments that specifically study input subsidies. Given the small sample sizes, the confidence intervals around the AEA registry data are large (even when using point estimates of treatment effects as data points, ignoring confidence intervals).

Having said this, reassuringly we find that our model’s estimates fall within these confidence intervals. Moreover, in both our model and in the AEA registry data we find a positive relationship between treatment effects and the share of modern production techniques used. Such out-of-sample validation is, of course, quite tentative – as there could be many meaningful differences between the studies’ settings and our calibrated application using the Uganda data. Nevertheless, this meta analysis adds some further reassurance that the economic forces captured in our model appear to be present across multiple different empirical contexts.

³Among the input subsidy experiments, which match more closely the intervention modeled here, sample size is even more constrained, with only 21 reporting harvest impacts, 16 agricultural revenue impacts, and 11 total household revenue impacts.

Figure A.9: Model Validation using Results from AEA Registry Trials



Presented is a local polynomial fit between treatment effects – in the AEA registry studies in black and in the model in dotted gray – as a function of the modern share (standardized z-score, normalized by the within country mean and standard deviation). Gray areas present the 95% confidence intervals around each line (which are more visible in the AEA registry due to the much smaller sample size than in the granular model data). Results are presented for treatment effects on harvests (in kgs), agriculture revenues, and total household revenues, as a percentage of the baseline mean. The top row presents results for all studies in the AEA registry with the given outcome reported (53 studies for harvest treatment effects, 37 for agricultural revenues, and 30 for total household revenue impacts). The bottom row presents results restricted to the experiments studying the same intervention we model: a subsidy for modern inputs. This results in a smaller sample size, with only 21 studies reporting harvest impacts, 16 agricultural revenue impacts, and 11 total household revenue impacts.

Appendix 5 Model and Solution Method

In Appendix 5.A, we first present the excess demand functions $\chi_{i,g}(\bullet)$ used in the text to define the equilibrium, and we then present the excess demand functions for the “price discovery” step. In Appendix 5.B, we develop the proof for uniqueness in price discovery for the special case with iceberg trade costs. Finally, in Appendix 5.C, we extend the model to allow for seasonal migration between rural markets and between rural and urban markets. For readers interested in more technical detail on the model and solution method we refer to the [Technical Appendix](#).

5.A Excess Demand Functions

The excess demand function for farmers are given by

$$\chi_{i,g}(\{p_{i,g}\}; W_{i,g}) =$$

$$= \begin{cases} \xi_{g,i}(\{p_{i,g}\}; W_{i,g}) I_i - p_{i,g} \sum_{\omega} q_{i,g,\omega}(\{p_{i,g}\}; W_{i,g}) & \text{for crop } g, \\ \xi_{g,i}(\{p_{i,g}\}; W_{i,g}) I_i & \text{for manufacturing product } g, \\ \sum_{k,\omega} \beta_{i,k,\omega} p_{i,k} q_{i,k,\omega}(\{p_{i,g}\}; W_{i,g}) - p_{i,g} L_i & \text{for labor } g = l. \end{cases}$$

The excess demand functions for urban household h for crops and manufacturing products g are given by:

$$\chi_{h,g}(\{p_{h,g}\}; W_{h,g}) = [\xi_{h,g}(\{p_{h,g}\}; W_{h,g}) - \mathbb{1}(g = g(h))] I_h$$

where expenditure share function $\xi_{i,g}(\bullet)$ and crop output function $q_{i,g,\omega}(\bullet)$ are defined in the main text where we skip the functional dependencies to avoid excessive notation. Indicator function $\mathbb{1}(g = g(h))$ is equal to one only if manufacturing product g belongs to urban household h and zero otherwise.

For Foreign and for crop g , $\chi_{F,g}$ equals $-\infty$ if $p_{F,g} < p_{F,g}^*$, equals $+\infty$ in the reversed case, and has a finite value if $p_{F,g} = p_{F,g}^*$. For manufacturing products g produced in Home: $\chi_{F,g} = X_{F,g}(p_{F,g})$.

For the price discovery step excess demand as functions of data D_A and prices $\{p_{i,g}\}$ of crops and labor for farmers and urban households are given by

$$\begin{aligned} \tilde{\chi}_{i,g}(\{p_{i,g}\}; D_A) &= \xi_{i,g} I_i(\{p_{i,g}\}; D_A) - \sum_{\omega} p_{i,g} q_{i,g,\omega}, & \text{for crop } g, \\ \tilde{\chi}_{i,g}(\{p_{i,g}\}; D_A) &= \sum_{k,\omega} \beta_{i,k,\omega} p_{i,k} q_{i,k,\omega} - p_{i,g} L_i & \text{for labor } g = l, \\ \tilde{\chi}_{h,g}(\{p_{h,g}\}; D_A) &= \xi_{h,g} I_h, & \text{for crop } g, \\ \tilde{\chi}_{F,g}(\{p_{F,g}\}; D_A) &= \begin{cases} -\infty & \text{if } p_{F,g} < p_{F,g}^* \\] -\infty, \infty[& \text{if } p_{F,g} = p_{F,g}^* \\ \infty & \text{if } p_{F,g} > p_{F,g}^* \end{cases} & \text{for crop } g, \end{aligned}$$

where

$$I_i(\{p_{i,g}\}; D_A) = \sum_{k,\omega} (1 - \alpha_{i,k,\omega}) p_{i,k} q_{i,k,\omega} + p_{i,L} L_i.$$

5.B Price Discovery

In this appendix, we show that, in the case with only iceberg trade costs (i.e., $t_{i,j} = 0$ for all i, j, 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 (7) and (8) (for a given set of Foreign prices) and excess demand functions in 5.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 $(o, g) \in \mathcal{J} \times \mathcal{K}_A \cup \{l\}$ where \mathcal{K}_A is the set of crops. Agent o 's endowment of (o, g) is $q_{o,g}$, which 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 $\mathbf{p} \equiv \{p_{o,g}\}_{(o,g) \in \mathcal{V}}$, the excess demand function (in value) for a variety $(o, g) \in \mathcal{V}$ is given by

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

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 ,⁴ we have $I_d = \sum_g p_{d,g}q_{d,g}$ and:

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

In turn, for $d = F$ we have $X_{F,o,g}(\mathbf{p}) = \infty$ if $p_{o,g} < p_{F,g}^*$, zero in the flipped case, and finite if $p_{o,g} = p_{F,g}^*$. We henceforth follow the convention that $q_{o,g} = 0 \implies p_{o,g} = \infty$ and $X_{d,o,g}(\mathbf{p}) = 0$, and also let $X_F(\mathbf{p}) \equiv \sum_{d \in \mathcal{J},g} X_{d,F,g}(\mathbf{p})$ denote the aggregate expenditure on goods from Foreign (imports).

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

$$\chi_{o,g}(\mathbf{p}) = 0, \quad \forall (o, g) \in \mathcal{V}. \tag{A.1}$$

We also assume that each agent $j \in \mathcal{J}$ produces at least one good (to ensure positive income)

⁴Recall 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.

and has a positive expenditure share on each good that it produces:

- Assumption A1:** 1) Endowments: $\sum_{g \in \mathcal{K}} q_{o,g} > 0, \forall o \in \mathcal{J}$.
 2) Demand: $q_{o,g} > 0 \implies \xi_{o,g} > 0, \forall o \in \mathcal{J}, g \in \mathcal{K}_A \cup \{L\}$.

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 $\chi_{o,g}(\mathbf{p})$ for any variety (o, g) with $\xi_{o,g} > 0$.

We say that a set of prices \mathbf{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) $X_{d,o,g}(\mathbf{p}) = X_{o,d,g}(\mathbf{p}) = 0, \forall o \in \mathcal{J}_1, d \in \mathcal{J}_2$ (i.e., no trade between the two blocks) and (ii) $X_{F,o,g}(\mathbf{p}) = 0, \forall o \in \mathcal{J}_1$ or $X_{F,o,g}(\mathbf{p}) = 0, \forall o \in \mathcal{J}_2$ (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 \mathbf{p} that solves the system of equations A.1. We do so by appealing to the result in Corollary 1 of [Berry, Gandhi and Haile \(2013\)](#) – henceforth BGH – which states sufficient conditions under which a function is injective on a set. We apply this result to our excess demand function $\{\chi_{o,g}(\mathbf{p})\}_{o,g}$.

Following 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, $X_F(\mathbf{p})$, as the “demand for good 0.” Trade balance then implies that

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

as in equation (2) of BGH.⁵ 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, which is satisfied here.⁶

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, $X_F(\mathbf{p})$, increases weakly with the price of any domestic variety in \mathcal{V} while demand for variety (o, g) , $\chi_{o,g}(\mathbf{p})$, 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.

⁵BGH add +1 to demand for good “0,” but this does not affect any results nor assumptions on monotonicity.

⁶Here 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.

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 $X_F(\mathbf{p})$ increases strictly in $p_{o,g}$ or (ii) there is a variety $(o', g') \in \mathcal{V} \setminus \mathcal{V}_0$ such that $\chi_{o',g'}(\mathbf{p})$ 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 $(o, g') \in \mathcal{V} \setminus \mathcal{V}_0$ then an increase in $p_{o,g}$ leads to an increase in revenues for agent o and an increase in demand for (o, g') through an income effect (under 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 $(o', g') \in \mathcal{V} \setminus \mathcal{V}_0$ such that agent o purchases good g' from o' , i.e. such that $X_{o,o',g}(\mathbf{p}) > 0$. 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.

The third case is one where, for any agent o , either all or none of the varieties are in \mathcal{V}_0 , 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 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 $p_{o,g}$ leads to greater demand for Foreign goods $X_F(\mathbf{p})$.

5.C 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 $t_{ij,l}$ and/or iceberg trade costs $\tau_{ij,l}$. 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., $t_{ij,l} = \tau_{ij,l} = \infty$ for $i, j \in \{F\}$.

Our model exposition in Section 2 continues 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 $a_h \sum_i x_{ih,l}$, where $x_{ih,l}$ are flows of labor from any origin i 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 $\{x_{ij,g}\}$ (measured in quantity at the destination). The equilibrium conditions (3) and (4), laid out in Section 2, apply to the model with migration as well. Based on the discussion above, only the equilibrium condition for urban income changes to:⁷

$$I_h = p_{h,l}L_h, \quad \forall h.$$

To implement this model extension in counterfactuals, one needs to calibrate the size of regional migration costs. One way of doing so is guided by the model's implication that two markets connected through migration have equal wages net of migration costs – similar to the estimation of trading frictions using bilateral price gaps. But of course, this implication is conditional on (partly unobserved) worker characteristics, as discussed at length in the literature on rural-urban wage gaps in developing countries (Gollin et al. (2014); Young (2013); Herrendorf and Schoellman (2018)).

⁷Here, we now need to distinguish between urban labor endowment L_h and urban employment $\sum_i x_{ih,l}$.

References for Online Appendix

- Bergquist, Lauren Falcao and Michael Dinerstein**, “Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya,” *American Economic Review*, 2020, 110 (12), 3705–47.
- , **Craig McIntosh, and Meredith Startz**, “Search Costs, Intermediation, and Trade: Experimental Evidence from Ugandan Agricultural Markets,” *Working Paper*, 2024.
- Berry, Steven, Amit Gandhi, and Philip Haile**, “Connected substitutes and invertibility of demand,” *Econometrica*, 2013, 81 (5), 2087–2111.
- Carter, Michael, Rachid Laajaj, and Dean Yang**, “Subsidies and the African Green Revolution: Direct Effects and Social Network Spillovers of Randomized Input Subsidies in Mozambique,” *American Economic Journal: Applied Economics*, 2020.
- Commission, UN Statistical**, “A recommendation on the method to delineate cities, urban and rural areas for international statistical comparisons,” *UN Statistical Commission*, 2020, *Fifty-first session*.
- de la Guardia Fernando, Edward Miguel Gufran Pathan Viviane Rocha Erik Ø. Sørensen Bertil Tungodden Hoces**, “Pre-registration, Reporting Guidelines, and Publication Patterns in Economics,” *Working paper*, 2025.
- Gollin, Douglas, David Lagakos, and Michael E Waugh**, “The agricultural productivity gap,” *The Quarterly Journal of Economics*, 2014, 129 (2), 939–993.
- Herrendorf, Berthold and Todd Schoellman**, “Wages, human capital, and barriers to structural transformation,” *American Economic Journal: Macroeconomics*, 2018, 10 (2), 1–23.
- Network), Columbia University CIESIN (Center For International Earth Science Information**, “Gridded Population of the World, Version 4 (GPWv4): National Identifier Grid, Revision 11 (Version 4.11) [Data set],” 2018. Palisades, NY: NASA Socioeconomic Data and Applications Center (SEDAC). <https://doi.org/10.7927/H4TD9VDP>.
- WorldMap, Harvard**, 2021. Available at: <https://worldmap.maps.arcgis.com/>.
- Young, Alwyn**, “Inequality, the urban-rural gap, and migration,” *The Quarterly Journal of Economics*, 2013, 128 (4), 1727–1785.

Technical Appendix More General Model and Exact Hat Algebra

In Appendix A, we present the model allowing for general functional forms on preferences and technology, for which exact hat algebra is feasible. In Appendix B, we formally describe this class of functions. In Appendix C, we provide additional details on recovering trade shares in manufacturing to apply exact hat algebra for that sector.

A 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 Technical Appendix B. The model equilibrium and solution to counterfactuals in the main text and excess demand functions in 5.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 $V(\{b_{j,k}p_{j,k}\}_j, I_j)$, 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 . Roy's identity implies that

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

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 good'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

¹With 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 $\alpha_{i,g,k,\omega}(\{p_{i,n}\}_i, r_{i,k,\omega})$.

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

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

Letting $\pi_{i,k,\omega}$ denote the share of land returns of farmer i coming from production of crop k with technique ω , an envelope result implies that

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

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

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

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.

B 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 $g(\bullet)$ 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 the following proposition 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:

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

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

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 $p_0, p_1, \hat{p} \in \mathbb{R}_+^n$, $f(p_0) = f(p_1) \implies f(p_0 \cdot \hat{p}) = f(p_1 \cdot \hat{p})$
- 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 $\pi : \mathbb{R}^n \rightarrow E$, 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.

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:

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

Proof of ii) implies i) For the converse, let's construct a function $K : Im(f) \rightarrow \mathbb{R}^n$ 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 $x, \delta \in \mathbb{R}^n$. With $p = \exp(x)$ and $\hat{p} = \exp(\delta)$ this implies: $f(p \cdot \hat{p}) = g(f(p), \hat{p})$.

Proof of iii) implies ii) With this projection, $F(x_0) = F(x_1)$ implies $\pi(x_0) = \pi(x_1)$, and: $F(x_0 + \delta) = F(\pi(x_0 + \delta)) = F(\pi(x_0) + \pi(\delta)) = F(\pi(x_1) + \pi(\delta)) = F(\pi(x_1 + \delta)) = F(x_1 + \delta)$

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. 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 $\pi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ 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: $F(x_0 + \pi(x_1)) = F(x_0 + x_1)$, and thus:

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

Similarly, as it implies $F(2\pi(x)) = F(\pi(2x))$, we get: $F\left(\frac{\pi(x_0) + \pi(x_1)}{2}\right) = F\left(\pi\left(\frac{x_0 + x_1}{2}\right)\right)$

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

$$\pi\left(\frac{x_0 + x_1}{2}\right) = \frac{\pi(x_0) + \pi(x_1)}{2} \quad (\text{T.1})$$

for all x_0, x_1 . For any F , we can construct such a projection π by choosing 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 T.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 $\pi : \mathbb{R}^n \rightarrow E$ 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 in iii).

Examples Cobb-Douglas production functions provide an extreme example where we just need to know the functional form and the relative change in inputs. Level sets (in log) are planes and are thus affine as described above.

Next, consider expenditure shares 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, as it is well documented in the literature. In this case, level sets (in log) are all the lines parallel to the $(1, \dots, 1)$ 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 . Such f is however not conducive to exact hat algebra.³ 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. $\phi_1 = 0$. Function f is then invertible up to p_1/w , noticing that p_1/w does not influence any expenditure share, and is now conducive to exact hat algebra.⁴

C 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 $j \in \mathcal{I} \cup \mathcal{H}$ where we denote the set

²Note 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.

³For instance, if $n = 2$, $\phi_i = 1$ and $\alpha_i = 1/2$, we have: $f_1(p_1/w, p_2/w) = \frac{1}{2} [1 + p_1/w - p_2/w]$ 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 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 observables.

⁴Note that other counter-examples can be found for homogeneous (homothetic) functions.

of urban households with \mathcal{H} and the set of farmers with \mathcal{I} . 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 e.g. [Donaldson and Hornbeck \(2016\)](#). We assume that we have some data on the 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 (e.g. 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 $\chi_{j,g(h)}$). 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 manufacturing variety k .

References for Technical Appendix

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