UC Berkeley
UC Berkeley Electronic Theses and Dissertations

Title
Essays in International Trade and Spatial Economics

Permalink
https://escholarship.org/uc/item/2x53s5q9

Author
Hoelzlein, Matthias M

Publication Date
2020

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

Essays in International Trade and Spatial Economics

by

Matthias M Hoelzlein

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Benjamin Faber, Co-chair
Professor Andres Rodriguez-Clare, Co-chair
Professor Cecile Gaubert
Professor Victor Couture

Spring 2020

Essays in International Trade and Spatial Economics

Copyright 2020
by
Matthias M Hoelzlein

Abstract

Essays in International Trade and Spatial Economics

by

Matthias M Hoelzlein

Doctor of Philosophy in Economics

University of California, Berkeley

Professor Benjamin Faber, Co-chair

Professor Andres Rodriguez-Clare, Co-chair

This dissertation studies how heterogeneous agents' exposure to local markets determines welfare and distributional effects of shocks such as policy interventions. In both chapters recent modeling tools in the international trade literature are applied to the study of spatial inequality in cities and its sources in chapter 2 or scaling up agricultural policy interventions in chapter 3. These tools are used to answer research questions in the context of quantitative general equilibrium models that feature rich heterogeneity but remain tractable. Moreover, the models in this dissertation predict clear relationships between the choices of households, firms and aggregate variables that are disciplined with detailed microdata.

The first chapter *Two-Sided Sorting and Spatial Inequality in Cities* studies a new economic force underlying the spatial sorting of rich and poor households in cities. On the demand side, households with different incomes choose neighborhoods and differ in their expenditures across various local services. On the supply side, service establishments sort into neighborhoods while taking into account proximity to their consumers. This two-sided sorting leads to endogenous differences in the local price index that amplify the concentration of household groups. A recent literature in urban economics has rationalized spatial sorting of households that is left unexplained by local incomes or housing costs by modeling pure amenity spillovers. In this chapter, I quantify the contribution of endogenous price indices to spatial sorting that is usually projected onto such reduced-form spillovers, and study the implications of two-sided sorting for urban policy. To do so, I develop a quantitative equilibrium model of the city that features two-sided sorting and nests many urban models. I estimate the key parameters of the model using detailed microdata for Los Angeles from 1990-2014. I find that spatial variation in local price indices decreases the estimates of reduced-form spillovers by about 30-50 percent. To shed light on the policy implications, I simulate policy counterfactuals, and compare the effects to the existing framework with only

reduced-form amenity spillovers. By studying a number of prominent place-based policies in Los Angeles, I find substantially different effects on neighborhood composition and welfare between both models.

The second chapter *Scaling Agricultural Policy Interventions: Theory and Evidence from Uganda* studies the welfare and distributional effects of scaling agricultural interventions in general equilibrium.[1] Interventions aimed at raising agricultural productivity in developing countries have been a centerpiece in the global fight against poverty. These policies are increasingly informed by evidence from field experiments and natural experiments, with the well-known limitation that findings based on local variation generally do not speak to the general equilibrium (GE) effects if the intervention were to be scaled up to the national level. In this chapter, we develop a new framework to quantify these forces based on a combination of theory and rich but widely available microdata. We build a quantitative GE model of farm production and trade, and propose a new solution method in this environment for studying high-dimensional counterfactuals at the level of individual households in the macroeconomy. We then bring to bear microdata from Uganda to calibrate the model to all households populating the country. We use these building blocks to explore the average and distributional implications of local shocks compared to policies at scale, and quantify the underlying mechanisms.

[1] This chapter is based on a working paper with Lauren Falcao Bergquist, Benjamin Faber, Thibault Fally, Edward Miguel and Andres Rodriguez-Clare (Bergquist *et al.* (2019))

To my wife Heather

Contents

List of Figures

List of Tables

Acknowledgments

This dissertation has been supported by many people in my life - advisors, family, colleagues and administrators - and would not have been possible without them. I would like to emphasize the contributions of a few in these acknowledgments but I am indebted to many more who remain unnamed and I am grateful for their support as well.

First, I would like to thank my advisors Ben Faber, Cecile Gaubert, Andres Rodriguez-Clare and Victor Couture for their excellent guidance and the many hours they put in to help me achieve my goals. I am particularly grateful to my main advisor Ben whose encouragement and advice has been invaluable for this dissertation and shaped my thinking as an academic economist. I also want give a special thanks to Andres who always had an open door for me and who has supported my carrier from the very beginning when I first came to Berkeley as an visiting graduate student. I also want to thank the other faculty members at Berkeley who supported me over the years, in particular, Thibault Fally, Galina Hale, Amir Kermani, Pat Kline, Enrico Moretti, Martha Olney, Christina Romer, David Romer, Benjamin Schoeffer and Nick Tsivanides. Moreover, I would like to thank my coauthors Lauren Falcao Bergquist and Edward Miguel. My gratitude goes to the many seminar participants at the UC Berkeley Economics Department and Haas School of Business for many helpful discussions. Navigating the PhD studies and ultimately finding my way into academia would not have been possible without the great administrators at the Economics Department, in particular, Patrick Allen.

In my time at UC Berkeley I have made many friends that supported my dissertation by providing advice and many kind conversations, in particular, Marc Dordal i Carreras, Sergii Meleshchuk and Nick Sander who have been excellent office mates and made our tiny office special. I am also grateful to Dennis Egger, Ezequiel Garcia-Lembergman, Johannes Hermle, Carla Johnston, Isabella Manelici, Evan Rose, Jose Vasquez and Roman Zarate. I also would like to thank the numerous undergraduate research assistants that contributed to my work over the years.

My dissertation would not have been possible without the tremendous support of my wonderful wife Heather who always had my back, especially, after our sweet daughter Vera was born. Gladly, I would like to dedicate my dissertation to Heather. I also would like to send my gratitude to my parents who supported my decision to pursue a doctorate in Berkeley in many ways and have been the best teachers in my life.

Lastly, I would also like to thank the Fisher Center for Real Estate and Urban Economics at UC Berkeley for financial support. Researcher(s) own analyses calculated (or derived) based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researcher(s) and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. All errors are mine.

Chapter 1

Dissertation Introduction

In this dissertation, I study the impact of different policies on heterogeneous agents and how the average and distributional welfare effects of these policies are determined by an agent's exposure to her local market. In particular, the location of an agent has important implications for the transmission of policy shocks from other agents in the economy through market prices for goods and services or wages in the local labor market. To quantify such general equilibrium effects empirically or experimentally has been very difficult or even impossible in the literature. Therefore, my dissertation relies on modern modeling techniques from the international trade literature that allows me study such policy shocks in general equilibrium with rich heterogeneity and complex geography to capture such indirect effects. Moreover, the models in this dissertation are being estimated and quantified with detailed microdata on household choices, firms and geography.
In Chapter 2 of my dissertation *Two-Sided Sorting ans Spatial Inequality in Cities* I evaluate how prominent urban policies such as local tax incentives to firms and social housing interact with a novel economic force that endogenizes the spatial sorting of rich and poor households in cities. If households value proximity to consumption amenities like restaurants or retail, and firm profits depend on local demand, the joint location choice of both sides of the market result in pecuniary externalities that amplify segregation and spatial inequality. The strength and direction of these pecuniary externalities vary by the degree to which different residents value certain amenities in their consumption bundles. A rich consumer's cost of living may be low in a rich neighborhood while costs might be high for a poor consumer because the neighborhood's amenities are less suited for her tastes. In my theory, such externalities operate on the local price index faced by households of different skill and amplify segregation of skill groups in cities. I develop a quantitative general equilibrium model of the city that incorporates three crucial ingredients: first, skill groups and firms in various sectors jointly choose where to locate. Second, demand by skill groups and profits by sectors are linked through non-homothetic preferences. Third, spatial frictions limit consumers' access to firms and vice versa. Using rich microdata on Los Angeles, I estimate the model and show that relative price indices of consumption are quantitatively important drivers of spatial inequality. In particular, I find that ignoring this channel leads to a sizeable upward bias of 30-50% in the degree of amenity spillovers, a concept used commonly in the literature to capture the endogenous nature of local amenities. With the estimated model, I assess the counterfactual

predictions of two place-based policies, a local tax incentive to firms and social housing, in the presence of two-sided sorting. I find that allowing for relative price indices changes the incidence of place-based policies on different groups in the city as compared to a model that only relies on amenity spillovers to amplify spatial segregation.

Chapter 3 *Scaling Agricultural Policy Interventions: Theory and Evidence for Uganda* studies how the average and distributional effects of local policy interventions can change when such an intervention is scaled to the regional or national level. Over the past decades an expanding literature in development economics has used local Randomized Controlled Trials (RCT) to inform the effectiveness of policies aimed at improving agricultural productivity in developing countries such as input subsidies or training programs. Such field experiments are designed to pick up local differences between a treated and control population and, hence, cannot speak to general equilibrium effects of the policy in question that affects both groups. Furthermore, RCTs require considerable resources limiting studies to small-scale interventions. Hence, in their limited scope they cannot fully inform policy makers who are interested in rolling out policies at a larger scale. This chapter develops a quantitative spatial model of farm production and trade that incorporates many linkages between individual farmers, in particular, how farmers interact on local markets and how local markets are integrated in the national economy through rich trade linkages. The model is then quantified with rich but widely available microdata from Uganda. In a series of policy counterfactuals the effects of a local subsidy for modern inputs such as chemical fertilizer on a small subset of Ugandan farmers are simulated and then compared to the results of a national program treating all approximately four million farm households in Uganda. The degree to which farmers are integrated in the local and national economy has important implications for the difference in effects between local and scaled intervention. For example, isolated markets have strong local price adjustments in product and labor markets in the scaled intervention shifting the surplus from the policy intervention from larger, richer producers towards poorer, less educated households. Due to such general equilibrium effects the scaled intervention benefits poorer households more than the local intervention within the same set of households. The chapter then proceeds to answer several other questions regarding the differences between local and scaled interventions, in particular, how the differences behave with varying levels of local saturation rates, the nature of trading frictions and differences in the level of aggregation.

Chapter 2

Two-Sided Sorting and Spatial Inequality in Cities

2.1 Introduction

Spatial inequality in cities and the clustering of rich and poor households across neighborhoods have attracted widespread attention by policy makers and given rise to public debate. These sorting patterns are usually rationalized by initial differences in access to employment, housing costs or natural amenities. Over time, differences in the social composition of neighborhoods are then reinforced by endogenous changes in the urban landscape that reflect the preferences and resources of each population. For instance, richer neighborhoods provide residents with more and different consumption options, greater safety, or access to better schools.

In the recent literature, these effects are typically summarized and modeled as reduced-form amenity spillovers[1]; utility of different groups is assumed to depend directly on the composition of residents, a reduced-form way to capture many channels through which neighborhood composition feeds back into utility of its residents. Since such amenity spillovers abstract from household or firm behavior, researchers and policy makers cannot directly observe them in the data. However, many urban policies target the behavior of households and firms to affect inequality or efficiency outcomes of a city. Hence, studying the microfoundations underlying the spatial sorting of household groups is crucial for the qualitative and quantitative implications of urban policies.

In this chapter, I set out to open up the black box of amenity spillovers. On the supply side, business establishments in retail and services decide where to locate while taking into account proximity to their consumers. On the demand side, households in different skill (or income) groups vary in their expenditure shares across these local services due to non-homothetic preferences. These basic forces give rise to two-sided sorting in cities whereby the spatial sorting of high and low skilled households[2] is a function of endogenous differences in access to consumption services

[1]See, for example, Diamond (2016), Su (2018b), Tsivanidis (2018), Guerrieri *et al.* (2013), Brueckner *et al.* (1999), Fajgelbaum & Gaubert (2019)

[2]For the remainder of the chapter I refer to "high skilled" as someone with at least a bachelor's degree and everyone

that enter the local price index.

In this setting, the chapter aims to answer two main research questions. First, to what extent does accounting for differences in the local price index across neighborhoods affect existing estimates of reduced-form amenity spillovers within and across household groups? Second, what are the implications of allowing for two-sided sorting for widely used urban policies that address spatial inequality in cities? In answering these questions, the chapter makes three main contributions. First, I propose a quantitative general equilibrium model of the city that introduces two-sided sorting of skill groups and firms in various local service sectors but nests other forces present in workhorse urban models. Second, I combine my theory with rich microdata from Los Angeles in order to quantify the fraction of spatial clustering that is due to endogenous price indices arising from two-sided sorting. Third, I simulate counterfactuals in the estimated model to study the implications of two-sided sorting for a number of prominent urban policies.

The analysis proceeds in four steps. In the first step, I document a number of motivating stylized facts. First, expenditure shares across retail and service sectors vary considerably with household income or skill. For example, high-income (or high skilled) households spend a much larger fraction of income on recreation providers such as gyms, education services, and restaurants than households with lower income. Second, the spatial distribution of establishments by sector is systematically correlated with the local skill or income composition. Some local service sectors concentrate in rich neighborhoods; whereas others are more equally spread out. Third, income elasticities of demand predict which sectors collocate with skill groups. Establishments in income-elastic sectors, such as gyms, are more likely to be found in neighborhoods with many high-skilled residents.

In the second step, I develop a quantitative spatial model that captures these moments in the data by incorporating three main ingredients. First, households and firms in various sectors simultaneously choose where to locate. Second, demand by skill groups and profits by sectors are linked through non-homothetic preferences. Third, spatial frictions limit consumers' access to firms and vice versa.

In the model, high skilled and low skilled households are identical except for incomes. They choose where to live based on residential rents, the local price index of consumption amentities and non-pecuniary amenities. Households have non-homothetic CES preferences with sector-specific income elasticities of demand. Following Hanoch (1975), this non-homothetic demand system has recently been applied by Comin *et al.* (2018), Borusyak & Jaravel (2018) and Matsuyama (2019); however, my paper is the first to leverage its convenient properties in the context of spatial sorting. The non-homotheticity at the sector level implies that high skilled households spend relatively more on income-elastic goods. Hence, profits in income-elastic sectors rise disproportionally in neighborhoods with many high skilled residents. More firms in income-elastic sectors locate in rich neighborhoods. Since high skilled households value firms in such sectors more than those who are low skilled, everything else being equal, the price index of the high skilled is lower in rich neighborhoods than in poor ones.[3] In equilibrium, high skilled and low skilled households face

else I classify as "low skilled". When I use the term "high skilled household" I refer to a household with a high skilled head.

[3]In more formal terms, expenditure shares are log-supermodular in income and the sector income elasticity of

different price indices in the same neighborhood and this difference is a function of the local skill composition. Two-sided sorting of firms and households leads to skill-location-specific pecuniary externalities that result in segregated neighborhoods. Since my model allows for reduced-form amenity spillovers, as in Diamond (2016), Su (2018b) and Tsivanidis (2018), I can trace out the qualitative and quantitative contributions of pecuniary externalities generated by two-sided sorting and reduced-form spillovers in explaining observed spatial inequality in Los Angeles.

I begin this task by characterizing how these pecuniary externalities operate between different skill groups and neighborhoods through the lens of the model. The key to understanding how different populations affect each other through price indices lies in the covariance of their expenditure shares across goods sectors. First, the impact of the externality is stronger within than across skill groups because expenditure shares are more correlated within groups. Second, the impact is stronger for any two neighborhoods that are geographically located close to one another, since the residents in both locations buy goods from similar shopping destinations. Furthermore, I show that the strength of these pecuniary externalities across skill groups varies with the initial income inequality, precisely because expenditure shares are not constant when preferences are non-homothetic. Hence, this dependence on initial conditions suggests that the pecuniary externality generated by two-sided sorting cannot be captured with constant exogenous amenity spillovers.

In the next step of the paper, I apply the model to detailed microdata from Los Angeles. To capture the spatial distribution of households by skill, I combine tract-level data from the National Historical Geographic Information System (NHGIS) and household-level microdata from IPUMS covering the years 1990-2014. I use this dataset to estimate key elasticities on the household side of the model. On the firm side, the National Establishment Time-Series Database (NETS) provides detailed, geo-coded information on the near-universe of establishments that allows me to estimate the spatial supply elasticity of firms from exogenous shocks to firm density. To discipline the strength of non-homotheticities in the model, I estimate income elasticities of demand for 28 local service and retail sectors[4] with household-level expenditure data from the Consumer Expenditure Survey (CEX) and Nielsen Consumer Panel. Taken together, the estimated elasticities and spatial distributions of households and firms characterize the key ingredients of the model.

In the main empirical part of this paper, I assess the extent to which allowing for two-sided sorting affects previous estimates of within and cross-group amenity spillovers in accounting for spatial clustering in the city. To this end, I estimate two different models. First, I relate population changes by skill to a tract's exposure to changes in the surrounding skill-mix, but assume that price indices of goods consumption do not vary by skill group and location. A similar relationship has been used to infer the strength of within and cross-group spillovers in Diamond (2016) and Su (2018b). Using this model, I find large reduced-form spillover elasticities similar to previous

demand. Therefore, firm profits are log-supermodular in the share of high skilled residents and the income elasticity. Both taken together, implies that the price index of goods consumption is log-supermodular in income and the share of high-skilled residents.

[4] I categorize sectors as local if consumers physically go to an establishment to purchase a good or service. In my sector definitions, I try to account for quality differences as much as possible given the constraint that sectors need to match to industry codes in the establishment microdata and expenditure categories in the expenditure data. For example, I differentiate fast food restaurants and full service restaurants or department and dollar stores.

estimates. Second, I estimate the same relationship but this time accounting for variation in price indices of goods. This allows me to jointly recover the key supply elasticity on the household side of the model and a set of unbiased reduced-form spillover elasticities. The comparison of the two sets of spillover elasticities shows that accounting for variation in price indices reduces the relative importance of spillovers for low and high skilled households by 30-50%. This result indicates that two-sided sorting is a quantitatively important driver of spatial sorting by skill groups.

My estimation also makes a methodological contribution. Since data on household expenditures and establishment-level prices is not available for most services at the level of skill groups and tracts, I cannot directly construct price indices at this level of disaggregation. Instead, I rely on the demand structure of the model to overcome this issue. Conditional on observed changes in income and residential rents, variation in the expenditure share on goods that I observe by skill groups at the tract-level, provides a sufficient statistic for changes in the price index of goods.

For identification, I use plausibly exogenous variation in changes in access to service establishments across census tracts, since changes in population and price indices may be correlated with unobserved shocks to the attractiveness of a location. In particular, I construct a shift-share instrument for changes in the local retail environment by interacting the initial sector shares of establishments in a location with sector growth rates in the total citywide number of establishments from other large urban centers in California. Pre-existing local sector shares are able to capture that locations provide different sector-specific supply-side advantages such as access to distribution networks, worker pools, or natural characteristics. These initial differences lead to variation in the exposure of locations to overall differences in plausibly exogenous growth across sectors. With this instrument, I exploit exogenous shifts in the availability of local consumption varieties leading to changes in the relative price index of goods that inform changes in real income of a neighborhood. To estimate reduced-form spillover elasticities in the same regression, I require a second source of exogenous variation. I construct a relative shift-share instrument that uses the same sector growth rates, but I interact the initial shares of establishments with the difference in citywide expenditure shares between high and low skilled households for each sector. Relative sector-level expenditure shares inform how growth across sectors differently enters into price indices of high and skilled households surrounding a tract.

After recovering the resident supply (mobility) elasticity, I proceed by estimating the spatial supply elasticity of firms, which is identified from the sector-specific relationship of individual establishment profits and the number of establishments in a location. The estimation poses two challenges to identification. First, local profits and the number of establishments are correlated with unobserved sector-location-specific productivity. Therefore, I exploit the differential exposure of sectors to plausibly exogenous variation in local demand that is driven by households' preference for the steepness of a location, a natural amenity highly valued by households. Second, I address selection bias due to sorting of firms on idiosyncratic productivity differences across locations by comparing establishments that belong to the same multi-establishment firm.

In the final step of this paper, I assess the implications of two-sided sorting and relative price indices for our understanding of urban policies. To this end, I simulate two place-based policies in Los Angeles, a new place-based tax incentive to firms and social housing. In both counterfactual exercises, I compare outcomes from two versions of the model. In the baseline model, skill groups sort on relative price indices and my unbiased estimate of reduced-form spillovers. In the model without price index effects, sorting is a result of only reduced-form amenity spillovers (which are biased upwards in the calibration when omitting the price index channel from the model).[5]

In my first counterfactual, I shock the firm distribution by simulating a new tax incentive to invest in economically disadvantaged areas, so-called Opportunity Zones (OZ).[6] I implement this policy by subsidizing profits of firms in the 257 OZs in Los Angeles and assume that the city's government finances the subsidy with lump-sum taxes on households. Firms respond strongly to the subsidy by moving operations into these zones. In the baseline model, the increased supply of consumption varieties induces households to locate in or close to OZs; however, high skilled households respond more to the now lower price index of consumption. As a result, the policy leads to gentrification of these initially disadvantaged areas. In the model without price index effects, the policy does not trigger any sizable mobility response of households since the location of firms has no bearing on the price index of consumption. Although the policy leads to modest average welfare losses for both skill groups of around .1-.2% of consumption, I find that welfare losses are smallest in the baseline model. With price index effects, the policy benefits local OZ residents, a population with high marginal utility and lack of access to firms in the initial equilibrium.

In the second exercise, I assess the effects of social housing on the spatial distribution of households and firms in LA. Using a newly collected dataset on address-level rent savings from social housing, I assume that the benefits of Social Housing accrue to low skilled households in the form of a rent subsidy financed by the city government. In both model versions, social housing leads to an inflow of low skilled households due to the subsidy and a corresponding decrease in skilled residents because of higher market rents. In the baseline model with price index effects, firms leave neighborhoods that have a large presence of social housing but more so in income-elastic sectors amplifying the clustering of the low skilled in treated neighborhoods due to an increase in relative price indices. In the model without price index effects, firms in all sectors do not respond thereby muting the effect of the policy on households.

The rest of the chapter is organized as follows. Section 2.2 discusses the paper's contribution to the existing literature. Section 2.3 describes the data. Section 2.4 presents stylized evidence on the joint location of firms and skill groups. Section 2.5 introduces the model. Section 2.6 builds intuition for the model. Section 2.7 takes the model to the data. Section 2.8 presents policy counterfactuals. Section 2.9 concludes.

[5] I can turn off price effects without changing the preference structure of the model by removing spatial frictions. Without such frictions the location choices of households and firms are unrelated.

[6] Opportunity Zones were implemented as part of the 2017 Tax Cuts and Jobs Act. The policy offers generous tax benefits to investors if they invest capital gains from previous investments in businesses located in roughly 8,700 designated tracts. According to U.S. Treasury Secretary Steven Mnuchin the total investment in Opportunity Zones will exceed $100B in 2019.

2.2 Related Literature

In addition to the work discussed above, this paper relates to several strands of the literature. First, there is a growth in literature that studies how heterogeneous preferences for consumption amenities and differential access to services, such as restaurants and retail, lead to sorting of households within cities.[7] In contemporaneous work, Couture *et al.* (2019) model competitive neighborhood developers who choose the local supply of a representative service sector and non-homothetic demand for housing and services. Their framework generates endogenous differences in access to services that induce sorting of households in different income groups. My approach nests their mechanism, but I extend it by modeling how non-homothetic preferences across many types of services reinforce sorting patterns. This additional layer of heterogeneity allows me to jointly capture the spatial distributions of heterogeneous firms and households in the data, and to study how the response of specific types of firms amplifies the effects of urban policies. Hence, my paper adds a new dimension to recent work on place-based policies in cities, for example Busso *et al.* (2013), Diamond & McQuade (2019), Diamond *et al.* (2018), and Davis *et al.* (2018).

Second, extensive literature documents large spatial differences in the availability and variety of goods and services associated with the size and social composition of a local population.[8] In particular, Handbury (2013) shows that when accounting for non-homothetic preferences across food items, income-specific price indices across cities are systemically correlated with local income. Specifically, poor households face higher food price indices in rich relative to poor cities and vice versa for rich households. My paper makes the analogous argument for the relative price indices of local services across neighborhoods. To formally account for these findings, I provide a general equilibrium framework that features a market for consumption amenities where heterogeneous service firms cater to local residents with different incomes and non-homothetic preferences across services.

Third, my paper contributes to a smaller literature that studies the spatial sorting of heterogeneous firms. Motivated by the uneven distribution of productivity across space, work in this area aims to separate local agglomeration externalities in production from the sorting of firms that are ex-ante heterogeneous in productivity. For example, Behrens *et al.* (2014) and Gaubert (2018) find that firm sorting explains a sizable share of the productivity premium of large cities. Brinkman *et al.* (2015) and Ziv (n.d.) study how agglomeration forces and firm sorting interact within cities. Different from these contributions, my paper focuses on demand-side complementarities between local resident composition and the determinants of firm demand, such as income elasticities. Hence, I add to this literature by evaluating how firm sorting contributes to the uneven distribution of household groups within cities.

Lastly, my model builds on the quantitative spatial economics literature that studies the rich structure of cities (Ahlfeldt *et al.* (2015); Allen *et al.* (2015)). The focus of this literature is pri-

[7]For example, Couture & Handbury (2017) and Baum-Snow & Hartley (2016) document that changing tastes for services over the last couple of decades are important drivers of the observed movement of college graduates into downtown neighborhoods.

[8]Waldfogel (2008), Schiff (2014), Couture (2016), and Davis *et al.* (2019) study variety and density of restaurants. Glaeser *et al.* (2018) look at several categories of local services.

marily on the trade-off between job location and residence. Moreover, it features homogeneous households with homothetic preferences.[9] I complement these papers by modeling spatial linkages within the city that are driven by consumption patterns of heterogeneous households with common non-homothetic preferences and the endogenous location choices of firms.

2.3 Data

In this section, I provide an overview of data sets I use to characterize the Los Angeles Metropolitan Area and to estimate the model.[10]

Throughout my analysis, I focus on outcomes for high-skilled and low-skilled households. A high-skilled household is defined as having a household head with at least a bachelor's degree. In 2014, Los Angeles consisted of approximately 1.1 million high-skilled and 2 million low-skilled households.

I use the 2010 Census tracts as the geographic definition of neighborhoods. The urban part of Los Angeles county, which is the basis of my analysis and what I refer to as Los Angeles from this point forward, consists of 2235 tracts with a total population just under 10M in 2014. The National Historical Geographic Information System (NHGIS) provides data on tracts for the Census 1990, 2000 and American Community Survey (ACS) 2012-2016[11]. All census tract data are interpolated to constant 2010 census tract boundaries using the Longitudinal Tract Data Base (LTDB). The primary information I extract from NHGIS are income distributions and distributions of expenditure shares on housing by income at the tract-level. The US Census and the ACS specifically provide household counts within defined income bins and household counts within income-expenditure share of housing, rent, and owner-cost bins. For each year, I combine this tract-level information with sample microdata from IPUMS at the level of Public Use Microdata Areas (Puma) to impute counts of households by skill, household income by skill and expenditure shares on housing/rent/owner cost by skill for each census tract and year. Since IPUMS microdata reports only pre-tax income of households, I compute the tax liability for each household using NBER's TAXSIM software and adjust tract income and housing expenditure share by group accordingly.

To capture the location and size of firms, I use the National Establishment Time-Series Database (NETS), collected by Duns and Bradstreet (D&B). This dataset provides annual information on exact geographic location, employment, and sales, as well as NAICS six-digit industry code and business characteristics of 2 million establishments in Los Angeles from 1990-2014.[12]

In order to map establishments into sectors, I first create 28 separate "local" sectors. I define a "local" sector based on the idea that households physically go to an establishment to pur-

[9]A notable exception is Tsivanidis (2018) who evaluates the distributional effects of infrastructure investment in a model of commuting by skill groups with Stone-Geary preferences.

[10]I will use Los Angeles Metropolitan Area and Los Angeles interchangeably. In the data I treat all urban contiguous Census tracts in Los Angeles County as the Los Angeles Metropolitan Area.

[11]I will refer to the ACS 2012-1016 as 2014 for the remainder of the paper.

[12]Throughout this paper I will use the terms establishment and firm interchangeably.

chase/consume different goods and services. In defining these sectors, I account for quality differences as far as possible. For example, households can eat out at fast food establishments versus full-service restaurants or buy groceries at supermarkets or specialty food stores. In both cases, I allow for two different sectors. However, due to data limitations mostly in household expenditure microdata, I cannot account for finer quality differences, like the difference between a regular supermarket and a specialty or upscale supermarket, such as Whole Foods Market. Next, I create a crosswalk between NAICS six-digit sectors in the NETS data and my 28 local sectors, as well as a crosswalk between my sectors and items in household expenditure microdata. In all crosswalks, I assign firms to local sectors based on where a typical household buys a good or service versus in which sectors certain goods are produced. For example, most food items are produced in agricultural sectors but predominantly purchased by consumers in grocery stores. All establishments or expenditure items that cannot be mapped to a local sector are assigned to a "frictionless" sector. I assume firms in the frictionless sector to be equally accessible to all consumers.

To estimate how demand by high skilled and low skilled households varies for local sectors, I use three datasets on household-level expenditures. First, I capture expenditure on service sectors in the Consumer Expenditure Survey (CEX) Interview data, which provides quarterly expenditures across roughly 700 unique expenditure categories. Second, I can break up some sectors that are aggregated in the quarterly data like "food away from home" using biweekly expenditures in the Diary data of the CEX. This breaks food away from home into restaurants, fast food, and bars. Lastly, I use the Nielsen Consumer Panel data to capture demand patterns across retail sectors. This dataset provides detailed expenditures across retail chains and is organized in retail channels, which correspond to 13 of my 28 local sectors for around 40-60k households between 2004-2017.

Data on the geographic distribution of housing comes from the Los Angeles County Tax Assessor. The tax assessor collects a variety of information on parcel size, building square footage, number of living units, number of stories in a building, year built, and usage for every parcel in the county. I compute the residential housing stock by aggregating the square footage of the main building on every residential parcel in a census tract. Since available tax assessor data only goes back to 2006, I impute the housing stock for 1990 and 2000 by removing all buildings built after the respective census year in a census block. I then assign the average size of the remaining units to all units that are reported in the census block data from NHGIS.[13]

Lastly, I use data from Lee & Lin (2017) to account for natural amenities like average slope, temperatures or distance to shore for each tract.

2.4 Motivating Evidence

To provide context and to motivate my modeling choices, I document stylized evidence on how households with a college-educated head and firms that provide goods and services preferred by

[13]For 2014, the number of living units in the ACS data corresponds almost perfectly to the number of living units in the tax assessor data. To keep the housing stock consistent over time I also assign the average unit size in a tract to all units reported in the 2014 ACS data.

richer households collocate throughout Los Angeles. In doing so, I point to endogenous differences in access to consumption varieties for high and low skilled consumers.

First, residence of high and low skilled households are strongly segregated in LA. Figure 2.1 plots the ratio of high-skilled residents over low skilled residents in a census tract, henceforth skill ratio, in Los Angeles in 2014. For example, South and East Los Angeles are almost exclusively populated by low skilled residents whereas college educated residents can be found along the coast (Santa Monica or Malibu) and in the hilly parts of Los Angeles. Figure 2.2 shows a similar relationship for the number of firms operating in local sectors for which I estimate income elasticities above the median over the number of firms in sectors with below median income elasticity.[14] Although the spatial distribution of firms is noisier, we can observe a similar pattern: locations with more high skilled households tend to be locations with more firms in sectors that are disproportionately preferred by richer households (or they are highly income-elastic). In Table 2.1, I report the results from regressing the log ratio of establishment counts by income elasticity on the skill ratio in a tract. A doubling of the skill ratio is associated with 12.6% higher ratio of firms in income-elastic sectors over firms in inelastic sectors.

To give a more specific example of how firms and households collocate based on demand patterns, I compare the two sectors with the highest income elasticities, recreation and education services, with the two sectors that I find to be the least income-inelastic, liquor/tobacco stores and convenience stores. Figure 2.3 plots the log number of establishments for both pairs of sectors against the log skill ratio in each tract. Recreation and education services are much more prevalent in locations with more high skilled residents as compared to liquor and convenience stores. In columns 2 and 3 of Table 2.1, I document the same relationship in linear regressions. In richer census tracts, the number of establishments in recreation and education is four times larger than the number of liquor and convenience stores. Furthermore, in columns 4 and 5, I find that the likelihood to observe any establishment in the two highly income-elastic sectors relative to the two inelastic sectors is six times higher as a function of the local skill ratio.

Figure 2.4 reports coefficients and 95% confidence intervals in black for all sectors from regressing log establishment counts on the log local skill ratio in a tract. I order the point estimates by my sector-level income elasticity estimates on the vertical axis. The ranking of regression coefficients and the ranks of the income elasticities are quite correlated (Spearman Rank Correlation: .495 with p-value of .007). The positive relationship implies that the number of firms in sectors that offer goods and services preferred by rich consumers, e.g. high income elastic sectors, is associated with locations populated by high skilled households. A major concern in this stylized analysis is that location choices of firms are differently impacted by supply factors instead of demand factos, such as having access to high skilled workers. To alleviate this concern, albeit imperfectly, I report coefficients from the same regression in Figure 2.4, but I control for the log ratio of high skilled to low skilled employees, total employment, and population density in each tract. The positive association between the effect of the skill ratio on the number of firms across

[14] I estimate income elasticities by sector with household expenditure microdata in the estimation section below. However, the ordering of sectors by income elasticity is intuitive: examples of highly income-elastic sectors are recreation, education, amusement or apparel stores. Liquor stores, dollar stores, fast food restaurants, or gas stations, I find to be income-inelastic.

sectors in a tract and the ordering of how much sectors are preferred by richer households remains stable. To sum up the stylized evidence, residential location choices of skill groups and location choices of firms operating in sectors that are preferred differently by skill groups are correlated. This points to endogenous differences in access to consumption varieties for low and high skilled households, a key ingredient to my model.

2.5 Model

Motivated by these correlations in the data, I develop a spatial general equilibrium model. It characterizes the forces leading to the spatial sorting of skill groups in a city, and it also guides my theoretical and empirical analysis. In particular, the model features two key skill-location-specific agglomeration forces: endogenous relative local price indices due to two-sided sorting and reduced-form spillovers.

Setup

The city consists of N neighborhoods, indexed n. It is populated by K types of heterogeneous households, indexed k, with fixed mass L_k. There are J sectors whose products differ in income elasticities of demand in household preferences. Households choose the location of their residence, consume housing h_{kn} and a bundle of goods across J sectors, $C_{kn}(g)$.[15] Housing is consumed in the neighborhood of residence whereas goods can be consumed everywhere in the city at iceberg trade costs, which I refer to as shopping frictions. Within each goods-sector there is a continuum of profit-maximizing firms that produce differentiated varieties and decide in which neighborhood to locate. Households can work anywhere in the city and provide labor inelastically for which there are no commuting costs.

Household Problem

A household ι of type k with preference draw $b_{kn}(\iota)$ choosing to live in n has utility,

$$\mathscr{U}_{kn}(\iota) = U_{kn} b_{kn}(\iota) = \frac{I_{kn}}{P_{kn}} b_{kn}(\iota) \tag{2.1}$$

where U_{kn} is real consumption, I_{kn} is type k's income and P_{kn} a type-neighborhood specific price index. Each household ι draws an idiosyncratic preference $b_{kn}(\iota)$ for every neighborhood n that is distributed Fréchet, $b_{kn}(\iota) \sim e^{-B_{kn} z^{-\kappa}}$. Conditional on their preference draws, households choose a neighborhood of residence n that provides the highest utility.

Importantly, real consumption U_{kn} follows a *non-homothetic* CES aggregator between housing and goods and is implicitly defined as

$$\left(a_h U_{kn}^{\varepsilon_h}\right)^{\frac{1}{\eta}} h_{kn}^{\frac{\eta-1}{\eta}} + \left(a_g U_{kn}^{\varepsilon_g}\right)^{\frac{1}{\eta}} C_{kn}(g)^{\frac{\eta-1}{\eta}} = 1 \tag{2.2}$$

[15]When using the term "goods", I am referring to all goods and services other than housing for brevity.

where η is the elasticity of substitution between housing and goods.
Depending on the relative size of ε_h and ε_g consumers shift expenditure between housing and goods when real consumption changes. For example, if housing is a necessity ($\varepsilon_h < \varepsilon_g$) then consumers with a higher level of real consumption spend a larger fraction of income on goods at given relative prices.[16] Comin *et al.* (2018), Borusyak & Jaravel (2018), and Matsuyama (2019) provide detailed discussions of non-homothetic CES preferences. I assume that the non-homotheticity in the model operates only on real market consumption in a neighborhood n, U_{kn}, as opposed to idiosyncratic utility $\mathscr{U}_{kn}(\iota)$.[17]

Consumption across goods sectors j also follows a *non-homothetic* CES aggregator with elasticity of substitution γ,

$$C_{kn}(g) = \left(\sum_{j=1}^{J} \left(\alpha_j U_{kn}^{\nu_j} \right)^{\frac{1}{\gamma}} c_{kn}(j)^{\frac{\gamma-1}{\gamma}} \right)^{\frac{\gamma}{\gamma-1}}.$$

As for the upper nest, the implied CES weights are functions of total real consumption U_{kn}. Given prices, as a consumer's real consumption increases, she shifts expenditures to goods from sectors with higher income elasticity parameters ν_j.[18]

Within each sector j there is an endogenous set of differentiated varieties $\Omega_n(j)$ being offered in neighborhood n.[19] A variety is denoted by ω. Households aggregate varieties from all neighborhoods in each sector j with a homothetic CES aggregator and elasticity of substitution σ,

$$c_{kn}(j) = \left(\sum_{n'=1}^{N} \left(\int_{\Omega_{n'}(j)} c_{knn'}(j,\omega)^{\frac{\sigma-1}{\sigma}} d\omega \right) \right)^{\frac{\sigma}{\sigma-1}}.$$

Households in n can access sector j varieties in another neighborhood n' at shopping costs $\tau_{knn'}(j)$ that takes iceberg form. The price of variety ω offered by a firm in n' faced by a household of type k in neighborhood n is

$$p_{knn'}(j,\omega) = \tau_{knn'}(j) p_{n'}(j,\omega)$$

where $p_{n'}(j,\omega)$ is the price of this variety at the shopping destination n'. I can write the expenditure share of a household of type k living in n on variety ω offered by a firm in n' in sector j according to

$$\tilde{s}_{knn'}(j,\omega) = \tau_{knn'}(j)^{1-\sigma} \left(\frac{p_{n'}(j,\omega)}{p_{kn}(j)} \right)^{1-\sigma},$$

where

$$p_{kn}(j) = \left(\sum_{n'=1}^{N} \left(\int_{\Omega_{n'}(j)} \tau_{knn'}(j)^{1-\sigma} p_{n'}(j,\omega)^{1-\sigma} d\omega \right) \right)^{\frac{1}{1-\sigma}} \tag{2.3}$$

[16] Note that when $\varepsilon_h = \varepsilon_g = 1-\eta$ the expression reduces to the regular CES consumption aggregator.

[17] It is not ex-ante clear whether systematic variation in tastes for different goods is due to real market consumption or overall well-being of a consumer. I choose the former for tractability. Alternative to making this assumption directly, I could assume that households make consumption decisions before the idiosyncratic preference shocks are realized.

[18] Similarly, the aggregator takes the homothetic CES form if $\nu_j = 0, \forall j$.

[19] The household takes set $\Omega_n(j)$ as given. It is determined in equilibrium as the interaction of household and firm problem.

is the price index of sector j faced by a household k in n.[20] Utility maximization implies that the expenditure share on all varieties in sector j follows

$$\tilde{s}_{kn}(j) = \alpha_j U_{kn}^{\nu_j} \left(\frac{p_{kn}(j)}{P_{kn}(g)} \right)^{1-\gamma}, \tag{2.4}$$

where the price index across goods is

$$P_{kn}(g) = \left(\sum_{j=1}^{J} \alpha_j U_{kn}^{\nu_j} p_{kn}(j)^{1-\gamma} \right)^{\frac{1}{1-\gamma}}. \tag{2.5}$$

A convenient property of non-homothetic CES preferences is that the household's expenditure shares on housing and goods as compensated demand is,

$$s_{kn}(h) = a_h \left(\frac{r_n}{I_{kn}} \right)^{1-\eta} U_{kn}^{\varepsilon_h} \quad \text{and} \quad s_{kn}(g) = a_g \left(\frac{P_{kn}(g)}{I_{kn}} \right)^{1-\eta} U_{kn}^{\varepsilon_g} \tag{2.6}$$

where r_n stands for the residential rent in n. The overall price index for type k in n is

$$P_{kn} = \left(a_h U_{kn}^{\varepsilon_h-(1-\eta)} r_n^{1-\eta} + a_g U_{kn}^{\varepsilon_g-(1-\eta)} P_{kn}(g)^{1-\eta} \right)^{\frac{1}{1-\eta}}. \tag{2.7}$$

Applying the Fréchet distribution ($b_{kn}(\iota) \sim e^{-B_{kn} z^{-\kappa}}$) to equation 2.1 I can write the mass of households of type k that resides in neighborhood n as

$$L_{kn} = \frac{B_{kn} U_{kn}^{\kappa}}{\Phi_k} L_k \tag{2.8}$$

where $\Phi_k \equiv \sum_{n'} B_{kn'} U_{kn'}^{\kappa}$. This term is related to expected utility of a household of type k given by

$$\bar{\mathscr{U}}_k = \Gamma \left(\frac{\kappa - 1}{\kappa} \right) \left(\sum_{n'} B_{kn'} U_{kn'}^{\kappa} \right)^{\frac{1}{\kappa}} = \gamma \Phi_k^{\frac{1}{\kappa}}.$$

Finally, skill types supply different levels of efficient units of labor: type k provides ρ_k units of labor. Since labor is freely mobile throughout the city (no commuting costs), the wage w per efficient unit of labor is equalized across all job locations. As a result, income of households with the same skill-level is constant and follows $I_k = w\rho_k + T_k$ where T_k is a lump-sum transfer that is independent of the location of a household. In this model, I abstract from spatial income differences within skill groups to focus on differences in price indices of consumption, which is the new mechanism in my model.

[20] For the remainder of the paper I denote expenditure shares within a sector or within the goods bundle with tilde. Expenditure shares out of total consumption are denoted without tilde.

Firm Problem

There is an infinite mass of potential entrepreneurs outside the city who have access to a single variety ω in a given sector j. To enter the city entrepreneurs have to incur fixed costs $f^e(j)$ in terms of labor. First, entrepreneurs observe expected profits from entering the city in their respective sector $E(\pi_n(j,\omega))$ and decide to do so if it exceeds the fixed entry cost $f^e(j)$. Conditional on entry, each entrepreneur ω in sector j receives an idiosyncratic productivity $z_n(j,\omega)$ to produce in neighborhood n drawn iid from a Fréchet distribution ($z_n(j,\omega) \sim e^{-A_n(j)z^{-\theta}}$ with $\theta > 1$).

An entrepreneur with variety ω in sector j and location n produces output with labor $l_n(j,\omega)$ according to

$$y_n(j,\omega) = z_n(j,\omega) l_n(j,\omega). \tag{2.9}$$

In order to determine the mass of firms in each sector j and neighborhood n we can write variable profits of ω choosing n as

$$\begin{aligned} \pi_n(j,\omega) &= \frac{1}{\sigma}\sum_{n'=1}^{N}\sum_{k'=1}^{K} s_{k'n'n}(j,\omega) I_{k'} L_{k'n'} \\ &\equiv z_n(j,\omega)^{\sigma-1}\tilde{\pi}_n(j). \end{aligned} \tag{2.10}$$

where $s_{knn'}(j,\omega)$ denotes share of total expenditure of type k living in n on variety ω located in n'. Profits can be decomposed into idiosyncratic productivity $z_n(j,\omega)$ and a neighborhood-sector-specific profit term $\tilde{\pi}_n(j)$.

Applying the Frechét distribution of $z_n(j,\omega)$ to 2.10 we get the mass of entrepreneurs in sector j that choose neighborhood n

$$M_n(j) = \frac{A_n(j)\tilde{\pi}_n(j)^{\frac{\theta}{\sigma-1}}}{\Pi(j)} M(j), \tag{2.11}$$

where $M(j)$ is the total mass of firms in sector j operating across the city and

$$\Pi(j) \equiv \sum_{n'}^{N} A_{n'}(j)\tilde{\pi}_{n'}(j)^{\frac{\theta}{\sigma-1}}.$$

I can rewrite the price index of sector j for household k in n in equation 2.3 as

$$p_{kn}(j) = \frac{\sigma}{\sigma-1} w \bar{\gamma}^{\frac{1}{1-\sigma}} M(j)^{-\frac{1}{\theta}} \left(\sum_{n'=1}^{N} \tau_{knn'}(j)^{1-\sigma} A_{n'}(j)^{\frac{\sigma-1}{\theta}} M_{n'}(j)^{1-\frac{\sigma-1}{\theta}} \right)^{\frac{1}{1-\sigma}} \tag{2.12}$$

where $\bar{\gamma} = \Gamma\left(1 - \frac{\sigma-1}{\theta}\right)$ is the Gamma function. We can directly see the effect of firm sorting on local prices: with $\theta > \sigma - 1$ neighborhoods closer to other locations (low $\tau_{knn'}(j)$) with more firms in a sector j face a lower price index due to a larger number of varieties $M_{n'}(j)$. Furthermore, as we can interpret the Frechét scale parameter $A_n(j)$ as a location-sector-specific average productivity, locations with higher $A_n(j)$ experience lower prices for j. Despite idiosyncratic productivity differences across firms and spatial mobility, all endogenous differences in sector-level prices across locations can be summarized by differences in the number of varieties.

The share of total expenditure of households in n in a given sector j on all varieties in location n' is then,

$$s_{knn'}(j) = s_{kn}(g)\tilde{s}_{kn}(j)\frac{\tau_{knn'}(j)^{1-\sigma}A_{n'}(j)^{\frac{\sigma-1}{\theta}}M_{n'}(j)^{1-\frac{\sigma-1}{\theta}}}{\sum_{n''=1}^{N}\tau_{knn''}(j)^{1-\sigma}A_{n''}(j)^{\frac{\sigma-1}{\theta}}M_{n''}(j)^{1-\frac{\sigma-1}{\theta}}}. \tag{2.13}$$

The shopping behaviour of households follows a gravity structure similar to workhorse trade models, but the allocation of expenditures across shopping locations is determined by the endogenous number of available varieties in a destination, shopping frictions, and average productivities.

Note that due to sorting, average profits per variety in a sector are equalized across locations in equilibrium,

$$\frac{1}{M_n(j)}\int_{\Omega_{n'}(j)}\pi_n(j,\omega)d\omega = \bar{\gamma}\Pi(j)^{\frac{\sigma-1}{\theta}}. \tag{2.14}$$

Lastly, we can pin down the mass of active varieties in each sector j, $M(j)$, by equalizing expected profits $E(\pi_n(j))$ with fixed costs of entry $f^e(j)$,

$$\bar{\gamma}\Pi(j)^{\frac{\sigma-1}{\theta}} = f^e(j). \tag{2.15}$$

Frictionless Sector

To allow for part of goods consumption that is independent of location, I include a frictionless sector (the Jth sector) within the goods bundle that operates in the same way as all other sectors. The exception being that consumers do not face shopping frictions for this sector such that $\tau_{knn'}(J) = 1, \forall n, n'$. In the absence of shopping frictions the firm location problem is exclusively determined by productivities $A_n(J)$,

$$M_n(J) = \frac{A_n(J)}{\sum_{n'}A_{n'}(J)}M(J)$$

and the price index for sector J collapses to

$$p_{kn}(J) = p(J) = \frac{\sigma}{\sigma-1}w\left(\bar{\gamma}M(J)\right)^{\frac{1}{1-\sigma}}.$$

Since households do not face shopping frictions the sector price index in any location is determined by the number of varieties city-wide $M(J)$, which is given by the free entry condition for the Jth sector as in 2.15. Since this result holds for any sector that does not face shopping frictions, I can remove relative price index differences as a force that creates sorting of households across locations by setting $\tau_{knn'}(j) = 1, \forall n, n', \forall j$.

Housing Markets

Atomistic landlords own a fixed amount of residential housing H_n in each neighborhood n. They are the claimant on the returns to housing but are fully taxed by the city government. All households in location n consume housing at rents r_n.

In equilibrium, expenditure on residential housing in n needs to equal $r_n H_n$

$$r_n H_n = \sum_k^K s_{kn}(h) I_k L_{kn} = a_h (r_n)^{1-\eta} \sum_k^K I_k^{\eta} U_{kn}^{\varepsilon_h} L_{kn}.$$

City Government

The city government collects all housing expenditures in the city by fully taxing landlords. It redistributes revenues net of any expenditures D, for example place-based subsidies to firms or renters. I will describe this below in detail when I discuss policy shocks. I assume that the city government returns the leftover budget to households as lump-sum transfer proportional to household labor endowment such that

$$T_k = \rho_k \frac{\left(\sum_{n=1}^N r_n H_n\right) - D}{\sum_{k=1}^K \rho_k L_k}. \tag{2.16}$$

The transfer scheme ensures that relative income differences are invariant to policy shocks.[21]

Reduced-form Spillovers

Since I study the role of relative price indices and reduced-form spillovers in explaining spatial sorting of households, I allow for direct spillovers within and across skill groups and locations. Similar to the previous literature (Diamond (2016), Su (2018b), Tsivanidis (2018), Fajgelbaum & Gaubert (2019)), I model reduced-form amenity spillovers for type k as returns to the number of residents of their type k and other types k'. However, I also allow spillovers to operate across neighborhoods, for example from n' to n, to make spillovers consistent with the notion that relative price indices depend on a neighborhood's geographic location, similar to Ahlfeldt *et al.* (2015) who model reduced-form spillovers operating on population density,

$$B_{kn} = \bar{B}_{kn} \mathscr{L}_{kn} = \bar{B}_{kn} \prod_{n'} \prod_{k'} L_{k'n'}^{\delta_{k'n',kn}}, \tag{2.17}$$

where $\bar{B}_{kn}$ represents exogenous amenities and $\mathscr{L}_{kn}$ stands for spillovers. Elasticities $\delta_{k'n',kn}$ govern how strongly amenities on households of type k in n respond to the number of residents of type k' in neighborhood n'. This formulation nests spillovers from the local skill composition, as in Diamond (2016) or Su (2018b), by setting $\delta_{k'n',kn} = 0$ for all $n \neq n'$.

[21]This formulation is isomorphic to assuming that households own a share in the city-wide housing stock proportional to their labor endowment. To finance policies, the city government taxes households lump-sum proportional to labor endowment.

Competitive Equilibrium

The equilibrium of this economy is defined by a distribution of households by neighborhood and skill group with $\sum_{n' \in \{1,2,...,N\}} L_{kn'} = L_k, \forall k$, a distribution of firms by neighborhood and sector with $\sum_{n' \in \{1,2,...,N\}} M_{n'}(j) = M(j), \forall j$, mass of firms in sectors $M(j), \forall j$, prices in all sectors and neighborhoods $\{p_n(j)\}$, sector price indices $\{p_{kn}(j)\}$, neighborhood-skill goods price indices $\{P_{kn}(g)\}$, neighborhood-skill price indices $\{P_{kn}\}$, wage w, residential rents $\{r_n\}$ and transfers $\{T_k\}$ such that:

1. Each type k in a neighborhood n maximizes utility given w, r_n, $p_n(j)$, $p_{kn}(j)$, $P_{kn}(g)$, P_{kn} and T_k and chooses the neighborhood that provides the highest utility with probabilities given in equation 2.8.

2. Firms in sectors j in neighborhood n maximize profits in 2.10 taking P_{kn}, $P_{kn}(g)$, p_{kn}, w, and the distribution of households as given and choose the neighborhood that maximizes profits with probabilities given in 2.11.

3. In each sector j, the mass of varieties is such that fixed cost of entry equals expected profits from entering the city.

4. Markets for residential housing clears in each n.

$$\sum_{k=1}^{K} \lambda_{kn} L_k h_{kn} = H_n, \forall n. \tag{2.18}$$

5. The labor market clears in the city

$$\sum_k \rho_k L_k = \sum_{j=1}^{J} \sum_{n=1}^{N} \int_{\Omega_n(j)} l_n(j,\omega) d\omega + \sum_{j=1}^{J} M(j) f^e(j). \tag{2.19}$$

6. Transfers are given by

$$T_k = \rho_k \frac{\left(\sum_{n=1}^{N} r_n H_n\right) - D}{\sum_{k=1}^{K} \rho_k L_k}.$$

Discussion of Uniqueness of the Equilibrium

The model supports multiple equilibria if the skill-specific agglomeration externalities (two-sided sorting and reduced-form spillovers) dominate the various dispersion forces present in the model.[22] On the household side, inelastic housing supply and idiosyncratic preferences for neighborhoods ensure that all neighborhoods are populated by high and low skilled households. Similarly, firms locate in all neighborhoods due to local competition forces and idiosyncratic productivity draws.

[22] If externalities are stronger than dispersion forces, some neighborhoods may attract predominantly high skilled households and firms in income-elastic sectors in one equilibrium configuration; however, the same neighborhoods may be populated by low skilled households and firms in income-inelastic sectors in an alternative equilibrium. Which equilibrium is reached depends on the starting values when I compute the model.

In section 2.7, I calibrate the parameters of the model to ensure that the household and firm distributions are unique. Although a formal proof of the necessary conditions for uniqueness is still work in progress, I perform a number of numerical simulations to test whether the my baseline model calibration supports a unique equilibrium. These tests suggest that the pecuniary externality from two-sided sorting and spillovers are weaker than the dispersion forces if preference and productivity draws are sufficiently dispersed (small κ and θ), the elasticity of substitution between housing and goods η is less than one, and real consumption is concave in expenditure.

2.6 Model Properties

Sorting Patterns

In the following section, I build intuition for the model's main contribution, namely, that firms in sectors with high income elasticity collocate with high-skilled households based on demand patterns arising from non-homothetic preferences in the model. This feature creates price index differences that endogenously lead to sorting patterns of households with different incomes. To keep the exposition and notation simple, I assume for now that shopping frictions outside the location of residence are infinite, $\tau_{knn'}(j) = \infty, \forall n' \neq n$, meaning households can only access varieties in their residence.[23] Furthermore, in this section and for the remainder of paper, I assume that $K = 2$, e.g. the city is populated by high and low skilled households.

The key variable summarizing underlying sorting patterns is the local expenditure share by skill group k in location n on goods from a sector j, $s_{kn}(j) = s_{kn}(g)\tilde{s}_{kn}(j)$ and is described in the following proposition:

Proposition 1. *Given prices, the expenditure share of households of skill k in location n on goods of sector j, $s_{kn}(j)$, is log-supermodular in real consumption U_{kn} and sector income elasticity parameter ν_j.*

Proposition 1 states that as households get richer[24] they value goods from sectors with higher income elasticity relatively more and that the difference is increasing with real consumption (see Matsuyama (2019) for a similar argument). As a consequence, high-skilled households' expenditure is tilted towards income-elastic sectors relative to low-skilled households. The top graph of Figure 2.5 shows a stylized graphical representation of this finding: I plot the log of expenditure shares for three sectors with decreasing income elasticity ($\nu_1 > \nu_2 > \nu_3$) on the vertical axis and household income on the horizontal axis. High skilled households with income I_{high} spend a larger fraction of income on the first sector and less on the other sectors in comparison to low skilled households with I_{low}.

[23] If shopping frictions do not vary by skill type $\tau_{knn'}(j) = \tau_{nn'} \forall k, j$ which I assume for the calibration of the model below, results of this section are unaffected since spatial consumption patterns are independent of skill type. However, defining local demand faced by firms in a location and the price index faced by households become complex functions of geography when shopping frictions are finite outside the residence, which makes the exposition less tractable.

[24] I assume that real consumption is increasing in nominal income I_k.

By relating this property to firm profits in n, I can rearrange profits of all varieties in sector j and location n in equation 2.10 as a function of the share of high skilled residents x_n in local population L_n,

$$\frac{\Pi_n(j)}{L_n} = \frac{1}{\sigma}\left(s_{high,n}(j)I_{high}x_n + s_{low,n}(j)I_{low}(1-x_n)\right).$$

Now, I can relate proposition 1 to average profits by resident of sector j according to the following corollary:

Corollary 1. *Given prices, total profits by resident of firms in sector j in location n is log-supermodular in high-skilled share x_n and sector income elasticity parameter ν_j.*

Intuitively, since high-skilled households spend more on income-elastic sectors, locations with a larger share of high skilled residents offer larger profits to firms in income-elastic sectors relative to income-inelastic sectors. Applying equation 2.14 and corollary 1 it follows immediately that the number of varieties $M_n(j)$ in income-elastic relative to income-inelastic sectors in locations with more high skilled residents must be larger than in locations with a lower share thereby keeping prices and total residents equal.[25] We can conclude that $M_n(j)$ is also log-supermodular in the high-skilled share x_n and sector income elasticity parameter ν_j implying that: $\frac{M_n(j)}{M_{n'}(j)}$ is non-decreasing in ν_j if $x_n > x_{n'}$. This result establishes that firms offering varieties in income-elastic sectors collocate with high income residents. Figure 2.5 summarizes how the number of varieties $M_n(j)$ increase faster in sector 1 (the highly elastic sector) than for the two less elastic sectors as the average income per resident in n on the horizontal axis increases.[26]

Next, I can combine proposition 1 and corollary 1 to characterize how residents respond to the distribution of varieties in a location. Since the price index of goods is a function of the high skilled share x_n and real consumption U_{kn}, $P_{kn}(g)$ has the following property.

Corollary 2. *Taking M_n, $M_{n'}$ and x_n as given and $\sigma, \gamma > 1$, households' price index of goods consumption, $P_{kn}(g)^{1-\gamma}$, is log-supermodular in real consumption U_{kn} and high-skilled share x_n or*

$$\frac{P_{high,n}(g)}{P_{low,n}(g)} < \frac{P_{high,n'}(g)}{P_{low,n'}(g)}$$

if $x_n > x'_n$.

Corollary 2 combines the intuition of both earlier findings and is graphically depicted in the bottom picture of Figure 2.5. Since richer households have higher expenditure shares on income-elastic sectors and locations with a larger share of high skilled households attract disproportionately more varieties in such sectors, the relative price index of goods between high skilled and low-skilled households must be lower in such neighborhoods compared to locations with more low skilled households.

[25] Profits are also increasing with total number of residents but at given expenditure shares and prices, in equal proportions for all sectors, such that only the composition of residents is relevant for relative sorting of varieties by sector.

[26] Note that with fixed income by type, the average income per resident is a sufficient statistic for x_n.

In equilibrium, the high-skilled share and real consumption are related through the location choice of households. Locations with lower relative goods prices between high and low skilled consumers attract more high skilled households. This, then, increases the high skilled share in the population and further reduces the relative price due to more local varieties in income-elastic sectors. Due to the interaction of location choice of households and firms, the model endogenously produces relative price differences that generates a pecuniary externality on residents. As we will see in the next section, this externality is separate from reduced-form spillovers as captured by $\mathscr{L}_{kn}$ in the model.

Local Decomposition of Price Index Effects and Reduced-form Spillovers

In this section, I characterize the forces in the model that link the location choice problems of different households across skill types and locations. In particular, relative price index effects and reduced-form spillovers generate externalities that amplify the mobility response of households to shocks. I consider a small shock in location n', for example an exogenous change in fixed amenities or place-based policy shock, that leads to a change in the population of skill group k' in n', which I refer to as the "shocked population." Then, I decompose the mobility response of residents of type k in location n ("target population") to this change. Taking expenditure shares as fixed,[27] I start from the expression for L_{kn} in equation 2.8. I take logs and differentiate with respect to $\log L_{k'n'}$ to get

$$\begin{aligned}\frac{d\log L_{kn}}{d\log L_{k'n'}} &= - \underbrace{\kappa\frac{(1-\eta)}{\bar{\varepsilon}_{kn}}}_{\text{Marginal Utility}}\underbrace{s_{kn}(h)\frac{d\log r_n}{d\log L_{k'n'}}}_{\text{Rent Congestion}} + \underbrace{\delta_{k'n',kn}}_{\text{Reduced-Form Spillover}} \\ &+\underbrace{\kappa\frac{(1-\eta)}{\bar{\varepsilon}_{kn}}}_{\text{Marginal Utility}}\frac{1}{\theta}\frac{I_{k'}L_{k'n'}}{Y_c}\left(\sum_j \underbrace{\frac{s_{kn}(j)s_{k'n'}(j)}{s_c(j)}}_{\text{Non-Homotheticity}}\left(1+\left(\frac{\theta}{\sigma-1}-1\right)\sum_{n''}\underbrace{\frac{\tilde{s}_{nn''}(j)\tilde{s}_{n'n''}(j)}{s_{n''}(j)}}_{\text{Spatial Frictions}}\right)\right) \\ &+ \underbrace{c_k}_{\text{Terms independent of } n}\end{aligned} \tag{2.20}$$

where $\bar{\varepsilon}_{kn} = s_{kn}(h)\varepsilon_h + s_{kn}(g)\left(\varepsilon_g + \frac{1-\eta}{1-\gamma}\bar{\nu}_{kn}\right)$ and $\bar{\nu}_{kn} = \sum_j \tilde{s}_{kn}(j)\nu_j$ are expenditure weighted average income elasticity parameters. Y_c stands for city-wide total income, $s_c(j)$ is the city-wide average expenditure share on sector j and $s_n(j)$ is the share of expenditure on varieties in n out of

[27]This decomposition holds only locally e.g. for a infinitesimal shock. In response to a larger shock expenditure shares adjust, making the expression intractable without necessarily conveying more information.

total expenditure in j. The term $\tilde{s}_{nn'}(j)$ is the expenditure share of a household in n on varieties in n' out of all expenditures on j.[28]

Terms in the first line correspond to forces present in many quantitative urban models. First, if the shocked population lives in n then additional residents bid up housing rents and congest the location muting the mobility response of the target population. In a model with non-homothetic preferences, price changes are evaluated at marginal utility since changes in expenditure do not translate one-to-one into utility.[29]

Second, depending on the sign of $\delta_{k'n',kn}$, reduced-form spillovers from the shocked population make location n more or less attractive to the target population. Stated differently, if the target population likes living close to the shocked population then the attractiveness of n increases. In this sense, reduced-form spillovers are a black box as they create sorting without any specific economic force underlying them.

The second line of 2.20 summarizes the effect of a change in the shocked population on the goods price index of the target population, the key new sorting force in the model.[30] First, if the distribution of firms is exogenously given, then price indices are are also exogenous and the term in the second line of 2.20 is zero. Thus, for prices to adjust an extensive margin in the number of varieties, e.g. sorting of firms, is necessary. When interpreting the first term in parentheses, price index effects are more pronounced for target populations that spend more on goods relative to housing (higher $s_{kn}(g) = \sum_j s_{kn}(j)$), for example for high skilled households if goods consumption is more income-elastic than housing. Price index effects are stronger if expenditure shares across sectors, $s_{kn}(j)$, of target population and shocked population are more correlated as the term in parentheses is similar to the covariance of expenditure shares. Due to non-homothetic preferences, households with similar incomes have more correlated expenditure patterns; hence, price indices respond more within skill groups than across groups.

The second sum in parentheses has a similar interpretation, but instead of variation in expenditure shares due to non-homothetic preferences, price index effects are stronger for populations that have spatially correlated expenditure shares. The "Spatial Frictions" term is large if target and shocked population both spend a lot in n'' (high $\tilde{s}_{nn''}(j)$ and $\tilde{s}_{n'n''}(j)$) relative to the overall importance of this location in j, as measured by $s_{n''}(j)$. For example, an increase in population and the associated entry of firms reduces price indices in neighboring locations more than in far away locations.

It is instructive to think about two special cases. First, if preferences are homothetic ($\varepsilon_h = \varepsilon_g =$

[28] Since I assume for simplicity that shopping frictions are the same for all skill types k, expenditures by destination do not vary with k.

[29] Under the assumption that utility in concave, but increasing in income, the importance of real consumption e.g. rents in determining the attractiveness of a location is diminishing relative to reduced-form spillovers or fixed amenities.

[30] The strength of the effect also depends, similar to rents, on the marginal utility of consumption, a set of elasticities (firm supply elasticity θ and love of variety $\frac{1}{\sigma-1}$) and the relative economic size of the shocked population, $\frac{I_{k'}L_{k'n'}}{Y_c}$.

$1-\eta$ and $\nu_j = 0$) and there is only one sector[31] then the expression in 2.20 simplifies to

$$\frac{d\log L_{kn}}{d\log L_{k'n'}} = -\kappa s_n(h)\frac{d\log r_n}{d\log L_{k'n'}} + \delta_{k'n',kn} + \kappa\frac{s_n(g)}{\theta}\frac{I_{k'}L_{k'n'}}{Y_c}\left(1+\left(\frac{\theta}{\sigma-1}-1\right)\sum_{n''}\frac{\tilde{s}_{nn''}\tilde{s}_{n'n''}}{s_{n''}}\right)+c_k$$

As a result of homothetic preferences, endogenous price effects on the target population are independent of skill and act solely as an agglomeration force on the local population due to love of variety and free entry, as in Krugman (1991).

Second, let us assume shopping is frictionless ($\tau_{nn'} = 1, \forall n, n'$) then

$$\frac{d\log L_{kn}}{d\log L_{k'n'}} = -\kappa s_{kn}(h)\frac{d\log r_n}{d\log L_{k'n'}} + \delta_{k'n',kn} + \kappa\frac{1}{\sigma-1}\frac{I_{k'}L_{k'n'}}{Y_c}\sum_j\frac{s_{kn}(j)s_{k'n'}(j)}{s_c(j)}+c_k$$

Without shopping frictions, price index effects only operate through entry of firms at the city border. More firms in sectors preferred by the shocked population enter the city leading to a stronger fall in prices if the target population's expenditure shares across sectors are more correlated. When I simulate counterfactuals with only reduced-form spillovers below I assume that shopping is frictionless, hence, this special case of 2.20 applies.

To sum up, price indices of consumption endogenously respond to changes in the spatial income or skill composition through sorting of firms. In addition, the strength of the effect depends on relative expenditure patterns of the target and shocked population leading to differential sorting of skill groups across space in response to a shock. In contrast to reduced-form spillovers, relative price index effects feature rich heterogeneity based on geography and income inequality of a city. Furthermore, the effect of relative price indices on mobility is not invariant to the initial equilibrium or context. Suppose we are able to compare the outcomes of the same policy shock for two different cities: one with little and the other with very strong income inequality. In the first city, consumption baskets of households would be similar, therefore, there would be little difference in mobility due to the price index channel. However, in the more unequal city we would observe that the price index channel leads to larger mobility responses due to more variation in expenditure patterns. Lastly, the strength and direction of the price index channel depends on the shock itself.[32] A shock to real consumption of local residents might cause variation in expenditure shares. As a result, this then changes how households' mobility responds to the shock. In the presence of non-homothetic preferences, the pecuniary externality generated by two-sided sorting cannot be captured by spillovers with constant elasticities. In other words, the price index effects, as in the second line of expression 2.20, are not constant such that they could be subsumed in $\delta_{k'n',kn}$.[33]

Empirically, changes in the relative price index of consumption and reduced-form spillovers are ex-ante not separable from observed mobility responses of a target population to exogenous shocks without information on changes in the price index or expenditure shares. In the next section I will separate both forces in the data.

[31] There is qualitatively no difference if there is more than one sector other than exogenous differences in productivity driving the spatial distribution of firms.

[32] Generally, my model suggests that reduced-form spillovers are subject to the Lucas Critique as they are invariant to context and shocks.

[33] However, I have not yet formally proved this claim.

2.7 Bringing the Model to the Data

In this section, I take the model to data from the Los Angeles Metropolitan Area. First, I describe a few parameters I take from the existing literature. Next, I estimate the key elasticities of the model with household-level microdata from the Nielsen Consumer Panel and the Consumer Expenditure Survey (CEX), firm level microdata from NETS, and Census tract level information. In particular, I empirically quantify how much relative price indices and reduced-form spillovers contributed to the mobility response of skill groups to shocks over time. Lastly, I discuss some additional pieces of information I need in order to invert the model to recover fundamentals and simulate counterfactuals.

Calibrated Parameters

Shopping frictions

In the model, shopping frictions capture how demand from households for establishments in distant locations falls relative to close locations. I assume that shopping frictions between locations n and n' are an increasing function of distance. Furthermore, I assume that this function is independent of local sector j and household type k and follows

$$\tau_{knn'}(j)^{1-\sigma} = \tau_{nn'}^{1-\sigma} = d_{nn'}^{\phi(1-\sigma)}, \forall j \in \{1, 2, ...J-1\}, \forall k$$

where $d_{nn'}$ is the straight line distance between the centroids of two tracts n and n'.[34] For the composite distance elasticity $\phi(1-\sigma)$ I choose a value of -1.5 within the range of estimates in the literature. For example, Couture *et al.* (2019) find values between -1.17 and -1.57 using smartphone movement data. Davis *et al.* (2019) use the location of consumers and restaurants from Yelp reviews to estimate a similar elasticity based on travel time and find values between -1 and -2.

Elasticity of Substitution within Sectors σ and across Sectors γ

The existing literature provides several estimates of the elasticity of substitution within service or retail sectors σ. Couture (2016) finds a value of 8.8 for restaurants, Atkin *et al.* (2018) 3.9 for retailers in Mexico, Dolfen *et al.* (2019) find 6.1 for offline stores, and Redding & Weinstein (2019) estimate a median σ to be 6.5 across disaggregated retail categories in Nielsen data. Su (2018a) reports values between 3.69 and 16 for disaggregated sectors, similar to my sector definition. As my sectors are quite aggregated and about half are retail sectors, which tend to have lower levels of substitution compared to services, I calibrate $\sigma = 5$ more towards the lower end of estimates. I will, however, report model results with higher σ as robustness.

To my knowledge, there exist fewer estimates for the elasticity of substitution across service or retail sectors γ.[35] For now, I rely on estimates from the trade literature and calibrate $\gamma = 2$

[34] For internal distances I rely on Helliwell & Verdier (2001) who find that internal distances are well approximated by $distance_{nn} = .52\sqrt{area_n}$ for a square city. I use this approximation since census tracts are close to square.

[35] See Borusyak & Jaravel (2018) for a short discussion.

which is in the middle of estimates from Redding & Weinstein (2017) who estimate the elasticity of substitution across 4-digit NAICS sectors using trade data to be 1.36 and Hottman & Monarch (2018) who find 2.78 for HS4 sectors.

Skill Premium

To create differences in expenditure shares between high skilled and low skilled households in the model, I need to take a stance on the skill premium which, in turn, creates nominal income differences between skill groups. To this end, I regress log after-tax household income in the Los Angeles sample of the ACS 2014 on a dummy for high skilled household head controlling household size, age, sex, and survey year fixed effects. As reported in Table 2.13, the coefficient is highly significant and implies that households with a high skilled head earn on average 70% higher nominal income compared to low skilled households. Hence, I set $\rho_{high} = 1.7$ and $\rho_{low} = 1$ in my model calibration.

Estimation of Income Elasticity Parameters

Income Elasticities across Goods Sectors

The model requires two broad sets of income elasticity parameters. First, sector-specific income elasticity parameters ν_j govern how households reallocate expenditures across goods sectors as a function of real market consumption. Second, parameters ε_h and ε_g capture how households shift expenditures between housing and goods when they have higher market consumption. I begin by estimating Engel curves for each of the 28 local sectors and the frictionless sector using consumer expenditure data from the CEX and Nielsen. I can write equation 2.4 as the expenditure on sector j relative to the expenditure on a reference sector j^* for household i in location n at time t and taking logs as

$$\log\left(\frac{p_{n,t}(j)c_{i,n,t}(j)}{p_{n,t}(j^*)c_{i,n,t}(j^*)}\right) = \log\left(\frac{\alpha_{i,j,t}}{\alpha_{i,j^*,t}}\right) + (1-\gamma)\log\left(\frac{p_{n,t}(j)}{p_{n,t}(j^*)}\right) + (\nu_j - \nu_{j^*})\log U_{i,n,t},$$

where the demand shifters $\alpha_{i,j,t}$ can be household and time dependent. We can note that $\nu_j - \nu_{j^*}$ is the elasticity of relative expenditures with respect to $U_{i,n,t}$. Since I cannot directly observe real consumption $U_{i,n,t}$ in the data but nominal income $I_{i,n,t}$ is commonly reported, I can locally approximate $\log U_{i,n,t}$ with the product of $\log I_{i,n,t}$ and the elasticity or real consumption with respect to nominal income,

$$\log\left(\frac{p_{n,t}(j)c_{i,n,t}(j)}{p_{n,t}(j^*)c_{i,n,t}(j^*)}\right) = \log\left(\frac{\alpha_{i,j,t}}{\alpha_{i,j^*,t}}\right) + (1-\gamma)\log\left(\frac{p_{n,t}(j)}{p_{n,t}(j^*)}\right) + (\nu_j - \nu_{j^*})\frac{\partial \log U_{i,n,t}}{\partial \log I_{i,n,t}}\log I_{i,n,t}.$$

Furthermore, denoting the average elasticity of real consumption with respect to nominal expenditure by $\varepsilon = \frac{\partial \log U}{\partial \log I}$, I can write the regression specification

$$\log\left(\frac{p_{n,t}(j)c_{i,n,t}(j)}{p_{n,t}(j^*)c_{i,n,t}(j^*)}\right) = \iota_{n,j,t} + (\nu_j - \nu_{j^*})\varepsilon \log I_{i,n,t} + u_{i,n,j,t}, \tag{2.21}$$

where $u_{i,n,j,t} = \log\left(\frac{\alpha_{i,j,t}}{\alpha_{i,j^*,t}}\right) + (\nu_j - \nu_{j^*})\left(\frac{\partial \log U_{i,n,t}}{\partial \log I_{i,n,t}} - \varepsilon\right)\log I_{i,n,t}$ and $\iota_{n,j,t}$ is a location-sector-time fixed effect capturing relative prices between j and j^* in a given location n and time t.

Data: I estimate regression 2.21 using household-level expenditure data from three data sources covering 2012-2016.[36] For retail sectors, I take annual expenditures by sector in the consumer panel data in Nielsen. For restaurants, bars and fast food, I rely on biweekly data from the CEX diary survey. For all other sectors I use quarterly expenditures from the CEX Interview survey. Table 2.2 reports the source for each sector in parentheses after the sector description. Since total household expenditure is reported in neither the Nielsen consumer panel nor the CEX diary data, I proxy total expenditure by nominal annual household income reported in each data source. To make the estimates across the three samples comparable I choose grocery stores as the reference sector, since expenditure on groceries is consistently reported across all samples. I restrict the sample to households living in an MSA and with a household head aged between 25 and 64.

Identifcation: Like Aguiar & Bils (2015) and Comin *et al.* (2018), I include dummies for household size (≤ 2, 3-4, ≥ 5), age of household head (25-37, 38-50, 51-64), and number of earners (1, $\geq$2) that are interacted with sector dummies to account for heterogeneity in preferences across cells defined by household characteristics. I also control for sector-MSA-time fixed effects to capture differences in relative prices and aggregate preference shocks across regions, sectors, and time. Lastly, to deal with measurement error in nominal income and endogeneity concerns, I instrument nominal income with a dummy for high skill or whether the household head has at least a four-year college degree. To make progress, I assume that, conditional on controls and the instrument, the elasticity of real consumption with respect to nominal income is orthogonal to the average elasticity, $E\left[\frac{\partial \log U_{i,n,t}}{\partial \log I_{i,n,t}} - \varepsilon | X, Z\right] = 0$. Albouy *et al.* (2016), Aguiar & Bils (2015), and Hubmer (2018) have to make similar assumptions to estimate income elasticities in non-homothetic demand systems.

Results: Figure 2.6 shows the estimated income elasticities by sector relative to grocery expenditure with 95% confidence intervals. In addition, Table 2.2 reports these results in columns two and three. The ordering of the estimates is quite intuitive. Liquor stores, dollar stores, convenience stores, or fast food have the lowest income elasticities; whereas, education, recreation (gyms, sports activities), and clothing exhibit the highest elasticities. Aguiar & Bils (2015) and Hubmer (2018) reassuringly find similar orderings using slightly different sector definitions. Confidence intervals are fairly tight except for bars and legal services for which I rely on few observations in the data. To interpret the magnitude of the estimates, consider the expenditure on fast food relative to full-service restaurants; a doubling of nominal income reduces relative expenditure by around 45%. In Table 2.15 and Figures 2.19 and 2.20 I report baseline results and two robustness checks. First, one implication of the model is that the location of residence matters for relative prices between sectors since local demand is likely correlated with the incomes of residents, potentially leading

[36]One concern could be that income elasticities are not stable over time. Aguiar & Bils (2015) discuss this issue and find that income elasticities are quite stable over time.

to biased estimates. Since I can observe the zip code of households in the Nielsen dataset, I can replace the sector-MSA-time fixed effect by a sector-zip code-time fixed effect to better account for local relative prices. Figure 2.19 shows estimated income elasticities for retail sectors in Nielsen with zip code level fixed effects. The point estimates and the ordering are broadly similar. Another concern may be that the three samples are fundamentally different and, hence, would give different results if they all covered the same set of sectors. Some of the sectors I use to estimate the baseline results can be approximately found in another sample. For example, the CEX reports expenditure on apparel, which I can assign to apparel stores and estimate the elasticity using CEX instead of Nielsen. Columns 5 and 6 of Table 2.15 and Figure 2.20 show results for some sectors where the alternative source is either Nielsen or CEX Interview and are alternative to the letter in parentheses after the sector name. Again, the results are broadly similar with a few outliers. For example, the estimate for appliances/electronics is considerably smaller in the CEX, which can be due to the fact that those goods can be bought in a variety of stores such as discount or hardware stores.

Income Elasticities between Housing and Goods

Two theoretical insights are useful to better understand the calibration of the income elasticity parameters in the upper nest of the preference specification in expression 2.2. First, in Appendix I show that all income elasticities ($\varepsilon_g, \varepsilon_h, \nu_j \forall j$) and the migration elasticity κ are defined up to a constant factor. Economic choices are unaffected if all elasticities are multiplied by a constant. At given prices, consumption and migration choices of households in response to higher nominal income are determined by their respective income elasticity parameter ($\varepsilon_g, \varepsilon_h, \nu_j \forall j$ and κ for location choice) relative to the elasticity of real consumption with respect to nominal income, which itself is a function of all income elasticity parameters. Second, with only expenditure data, goods-sector elasticity ε_g cannot be separately identified from the expenditure weighted sum of all sectoral elasticities ν_j. The reason for the latter is that when consumers get richer, they shift expenditure between housing and goods as a result of non-homothetic preferences in the upper nest. However, they also reallocate expenditures within the goods bundle due non-homothetic preferences that affect the price of goods relative to housing, leading to changing relative expenditures on goods and housing. To be able to pin down the values of $\varepsilon_g, \varepsilon_h, \nu_j \forall j$, I assume that preferences in the upper nest are homothetic, which implies

$$\varepsilon_h = 1 - \eta \quad \text{and} \quad \varepsilon_g = 1 - \eta.$$

An immediate implication of this assumption is that the non-homotheticity in housing and goods demand operates exclusively through the price index for goods relative to the price of housing. With homothetic preferences in the upper nest, the elasticity of real consumption with respect to nominal income collapses to

$$\frac{\partial \log U_{kn}}{\partial \log I_{kn}} = \frac{1-\eta}{\bar{\varepsilon}_{kn}} = \frac{1}{1 + s_{kn}(g)\frac{\bar{\nu}_{kn}}{1-\gamma}}$$

where $\bar{\varepsilon}_{kn} = s_{kn}(h)\varepsilon_h + s_{kn}(g)\left(\varepsilon_g + \frac{1-\eta}{1-\gamma}\bar{\nu}_{kn}\right)$ and $\bar{\nu}_{kn} = \sum_j \tilde{s}_{kn}(j)\nu_j$. My estimated sectoral income elasticities from equation 2.21 are relative to a reference sector (groceries) and relative to the average income elasticity ε. To recover specific values of ν_j for all sectors, I need values for the elasticity of substitution η and the difference between ε_h and the composite income elasticity parameter for goods, $\varepsilon_g + \frac{1-\eta}{1-\gamma}\bar{\nu}$, evaluated at average expenditure shares by sector. For these last two pieces I rely on values from Albouy *et al.* (2016) who estimate non-homothetic CES preferences between housing and goods using variation in housing expenditure shares and returns to skill across MSAs. I take their estimates with renters and owners, the average expenditure share on goods from IPUMS microdata for Los Angeles ($s_c(g) = .6663$), and citywide sales shares by sector from NETS for 2014 reported in Table 2.2. I calibrate $\eta = .493$ and $s_c(g)\frac{\bar{\nu}_c}{1-\gamma} = .839$.[37] The calibration implies that housing and goods are complements ($\eta < 1$) and that housing is a necessity relative to goods ($\varepsilon_h < \varepsilon_g + \frac{1-\eta}{1-\gamma}\bar{\nu}$).

In Table 2.14, I report some reduced-form evidence, namely that the expenditure share on housing indeed falls with income (or skill). I regress the expenditure share on housing in the ACS microdata (colums 1 and 2), measured as housing expenditure out of after-tax HH-income, and CEX microdata (columns 3 and 4), measured as housing expenditure out of total expenditure, on a dummy for skilled household, a set of time-location fixed effects, and dummies for age, size, sex, and home ownership. Consistent with housing being a necessity, I find that high skilled households spend around 5ppt less on housing than low skilled households in the ACS data and 1-2ppt in the CEX data. Table 2.2 reports implied values of ν_j, based on $\gamma = 2$, and sector sales shares used in the calibration. Note that implied ν_j are negative; however, they follow the same ordering as the estimated relative income elasticities. Holding prices constant, the latter implies that expenditures on high-ν_j sectors increase with income relative to low-ν_j sectors. One implication of the former is that the goods price index increases with real consumption. The expenditure share on goods also increases with higher income, since goods and housing are complements. Moreover, an increase in the price index of goods consumption leads to an increase in the expenditure share on goods, everything else being equal.

Estimation of Resident Supply Elasticity κ and Reduced-Form Spillovers Elasticities

After characterizing the endogenous sorting channels of the model in section 2.6, I now empirically decompose the mobility response of households to exogenous shocks in real income into the price index channel and reduced-form spillovers. I estimate the resident supply elasticity κ, which governs how strongly households' location choices responds to spatial differences in real consumption and reduced-form spillover elasticities $\delta_{k'n',kn}$. The motivation behind the estimation is twofold. First, I can assess what portion of observed changes in spatial inequality, usually explained with reduced-form spillovers, can be attributed to endogenous local price index differences. Second, to perform policy counterfactuals in the model, I require values for κ and two sets of reduced-form spillover elasticities. "True" reduced-form elasticities net out the effect of relative price indices and "biased" elasticities that encompass endogenous price index differences.

[37] Albouy *et al.* (2016) estimate $\frac{1-\eta}{1-\gamma}\bar{\nu} = .6358$. Dividing by $1-\eta$ and multiplying by $s_c(g)$ gives this value.

Starting with the location choice of households in equation 2.8, reduced-form spillover definition in equation 2.17, and taking log changes over time t, denoted by the hats, I get

$$\log \hat{L}_{kn,t} = \kappa \log \hat{U}_{kn,t} + \log \frac{\hat{L}_{k,t}}{\hat{\Phi}_{k,t}} + \hat{\mathscr{L}}_{kn,t} + \log \hat{\bar{B}}_{kn}, \tag{2.22}$$

where $\hat{\mathscr{L}}_{kn,t}$ is a measure of the change in reduced-form spillovers. For example, Diamond (2016) assumes that spillovers $\hat{\mathscr{L}}_{kn,t}$ are a function of the local skill ratio $\frac{L_{high,n}}{L_{low,n}}$. Next, I replace the change in real consumption from expression 2.1 by its components and locally linearize the change in the price index in 2.7 around values of the last period,[38]

$$\log \hat{U}_{kn,t} = \log \hat{I}_{kn,t} - s_{kn,t-1}(h) \log \hat{r}_{n,t} - s_{kn,t-1}(g) \log \hat{P}_{kn,t}(g). \tag{2.23}$$

Equation 2.22 summarizes the drivers of a changing neighborhood population through the lens of the model. Locations become attractive if they offer higher real consumption $\hat{U}_{kn,t}$, stronger spillovers $\mathscr{L}_{n,t}$ or improving exogenous amenities $\hat{\bar{B}}_{kn}$. Real consumption itself can change due to nominal income, housing rents, or the local skill-specific price of goods. Previous work in the literature has made two broad assumptions on goods prices in estimating this type of regression. First, goods prices are independent of household type, and, second, either price indices are identical across locations or perfectly correlated with the local price of housing. With constant goods prices across locations 2.23 reduces to

$$\log \hat{U}_{kn,t} = \log \hat{I}_{kn,t} - s_{kn,t-1}(h) \log \hat{r}_{n,t} - s_{kn,t-1}(g) \log \hat{P}_t(g) \tag{2.24}$$

and under perfectly correlated local prices

$$\log \hat{U}_{kn,t} = \log \hat{I}_{kn,t} - \log \hat{r}_{n,t}. \tag{2.25}$$

As a result, in both specifications, the goods price index does not affect the differential location choice of skill groups. One contribution of this paper is to assess whether changes in goods price indices are by themselves driven by a changing neighborhood population or composition through the location choice of firms. Suppose, firms in income elastic sectors tend to collocate with high skilled households. Then, if there is an influx of high skilled households, relative goods prices falls more for high skilled than for low skilled households in the same location due to more varieties in sectors preferred by the high skilled. This leads to a negative correlation between changes in the price index and spillovers for high skilled, as well as a positive correlation for low skilled households.[39] Hence, omitting skill-specific goods price changes can lead to upward bias in the

[38] Note that I used $\varepsilon_h = \varepsilon_g = 1 - \eta$ in section 2.7

[39] In Figure 2.21, I show evidence that changes in the skill-mix surrounding a tract are negatively correlated with changes in relative price indices for high and low skilled households. Since I cannot directly observe price indices, I use relative expenditure shares on goods in the top graph of the upper panel of 2.21 as a sufficient statistic for relative price indices (see below). In the bottom graph, I construct relative CPIs broadly consistent with the model.

estimated spillover elasticities for high skilled and a downward bias for low skilled households, effectively overestimating the role of spillovers in explaining observed sorting patterns.[40]

Since I do not have access to data on expenditures across service sectors at skill-tract level and corresponding sector price indices, I cannot construct skill-location-specific goods price indices in the data.[41] Hence, I rely on the model to find a sufficient statistic for changes in price indices. I can rearrange equation 2.6 to solve for price index changes in the goods sector:

$$\log \hat{P}_{kn}(g) = \log \hat{I}_{kn} - \frac{\varepsilon_g}{1-\eta} \log \hat{U}_{kn} + \frac{1}{1-\eta} \log \hat{s}_{kn}(g) - \frac{1}{1-\eta} \log \hat{a}_{kn}(g). \tag{2.26}$$

Intuitively, conditional on income and rent changes as well as constant relative tastes for goods and housing ($\hat{a}_{kn}(g) = \hat{a}_{kn}(h)$), all variation in the goods price index is captured by the expenditure share on goods. Hence, I can use variation in the expenditure share on goods as a sufficient statistic for variation in the price index of goods. Under the earlier assumption on the income elasticity parameter for housing and goods consumption ($\varepsilon_g = \varepsilon_h = 1-\eta$) we can plug 2.26 into equation 2.23 and combine with equation 2.22 to arrive at the main regression specification,

$$\log \hat{L}_{kn,t} = \kappa \log \hat{I}_{kn,t} - \kappa \log \hat{r}_{n,t} - \kappa \frac{1}{1-\eta} \frac{s_{kn,t-1}(g)}{s_{kn,t-1}(h)} \log \hat{s}_{kn,t}(g) + \hat{\mathscr{L}}_{kn,t} + \iota_{k,t} + u_{kn,t}, \tag{2.27}$$

where I collect skill-specific terms in a skill-time fixed effect and the error terms capture changes in exogenous amenities and, potentially, changing tastes for goods and housing.[42] If housing and goods are complements ($\eta < 1$), an increase in local goods price index leads to an increase in the expenditure share on goods. Hence, exogenous variation in $\frac{s_{kn,t-1}(g)}{s_{kn,t-1}(h)} \log \hat{s}_{kn,t}(g)$ and a value of η identifies the resident supply elasticity κ.

Lastly, I assume a parametric form for reduced-form spillovers $\hat{\mathscr{L}}_{kn,t}$ in equation 2.17 in the model,

$$\hat{\mathscr{L}}_{kn,t} = \delta_k \sum_{n'} \frac{d_{nn'}^{\psi}}{\sum_{n''} d_{nn''}^{\psi}} \log \frac{\hat{L}_{high,n',t}}{\hat{L}_{low,n',t}}. \tag{2.28}$$

I assume that reduced-form spillovers operate on the distance-weighted skill ratio.[43] Despite being highly parametric this formulation has two advantages. First, it seems sensible to think that

[40] Only relative spillovers matter for differences in sorting between low and high skilled households. Consider the case when spillovers are positive, but identical for high and low skilled. Then, taking the difference of equation 2.22 between high and low skilled cancels spillovers. However, the level of spillovers matters for spatial sorting within group.

[41] Some datasets, for example Nielsen Homescanner data, provide barcode-level expenditures and prices for retail. However, such data is not available for most services.

[42] Note that I can write the term $\frac{s_{kn,t-1}(g)}{s_{kn,t-1}(h)} \log \hat{s}_{kn,t}(g) = -\log \hat{s}_{kn,t}(h)$ as the negative change in the housing share.

[43] I can write 2.17 as a Cobb- Douglas function of the number of residents in all locations

$$B_{kn} = \bar{B}_{kn} \mathscr{L}_{kn} = \bar{B}_{kn} \left(\prod_{n'} L_{high'n'}^{\omega_{nn'}} L_{low'n'}^{-\omega_{nn'}} \right)^{\delta_k}$$

and $\omega_{nn'} = d^{\Psi}$. Assuming CRS with respect to the skill composition ($\sum_{n'} \omega_{nn'} = 1$) inside the parentheses and taking δ_k as the overall degree of spillovers returns equation 2.28.

spillovers operate beyond the skill composition of a tract. Second, the location of a neighborhood relative to the changing skill distribution across the city gives variation independent from the local skill composition and avoids a purely mechanical relationship between changes on the left and right-hand side of 2.27. With an estimate of δ_k and a value for ψ I can recover spillover elasticities $\delta_{k'n',kn}$ in the model with

$$\delta_{k'n',kn} = -\delta_{kn',kn} = \delta_k \frac{d_{nn'}^{\psi}}{\sum_{n''} d_{nn''}^{\psi}}. \tag{2.29}$$

Data: To estimate regression 2.27, I pool changes in the number of households by skill group between 1990-2000 and 2000-2014 in census tracts in LA. I exclude tracts with a population of less than 1000 in 1990 to avoid capturing newly developed neighborhoods. The key independent variable is the change in the expenditure share on goods measured as one minus the expenditure share on housing, which I impute from census tract data on the distribution of expenditure shares by income group (see **??**). In order to construct the proxy for endogenous amenities in 2.28, I choose the same distance elasticity as for the calibrated shopping frictions, $\psi = -1.5$ but also consider a higher value of -3 as robustness check. In the main specification, I control for changes in household income and residential rents as well as natural amenities such as log distance to the center of Los Angeles (City of Los Angeles City Hall), log average slope in a tract, and log population density in 1990.

Identification: To identify the resident supply elasticity κ in equation 2.27, I need a source of variation that shifts the price index on goods but is uncorrelated with unobserved exogenous amenity shocks and taste shocks in a neighborhood. For example, a neighborhood that receives a subway station might attract more residents due to better labor market access. However, firms also locate in this neighborhood since they too benefit from better access to workers and consumers, in turn, affecting access to services. Similarly, reductions in local crime rates might attract both households and firms. For the identification of the resident supply elasticity, I use plausibly exogenous variation in a tract's access to service establishments.

In particular, I construct a shift-share instrument based on cross-sectional variation in the share of establishments in service sectors in a location that I interact with sector growth rates in establishments from the other two large urban areas in California, namely San Francisco Bay Area and San Diego. The initial sector shares in the number of establishments contain information about supply-side characteristics of a location, such as access to suppliers, specific worker pools, or natural advantages. The idea is that a sector's growth in urban areas due to changing tastes or technological improvements will lead to a larger increase in the number of establishments in locations that provide such supply-side advantages to firms in this sector. For example, many establishments in the recreation sector concentrate along the beach since many recreation activities are related to water. We would expect that overall growth in the number of recreation establishments bring more businesses to locations by the water than inland tracts. Conceptually, the instrument captures the idea that an exogenous increase in locally available consumption varieties lowers the goods price index in a location. If goods consumption and housing are complements, as discussed in Section

2.7, an increase in locally available varieties should lower the expenditure share on goods and attract more residents.

Formally, I construct the following average price index instrument,

$$P^{IV}_{n,t} = \sum_j \left(\sum_{n'} \frac{M_{n',t_0}(j)\mathbb{1}(distance_{nn'} < b)}{M_{t_0}(j)} \right) \log \hat{M}^O_t(j)$$

where b is a distance buffer, t_0 refers to a base period 1990 and superscript O stands for urban areas other than LA. Motivated by Agarwal *et al.* (2018), who report that consumers travel only short distances to consumption venues, I choose $b = 5km$, which includes an average of 56 tracts (median of 50) in the main specification but I also report additional results for a larger buffer of $10km$.[44]

The identifying assumption is that, conditional on controls, goods sector growth rates in San Francisco and San Diego are orthogonal to tract-level changes in amenities and tastes in Los Angeles, for example, changes in labor market access or crime rates. In other words, I argue that growth rates are as good as randomly assigned to sectors even though exposure of a location to sectors is endogenous. First, since the sector shares do not add to one, locations with initially more service establishments on average, such as downtown tracts, are more exposed to overall growth in services. They might also experience faster population growth due to changing preferences for such locations, specifically, by skilled households (Couture & Handbury (2017)). Thus, I control for the sum of establishment shares in each location interacted with time dummies to isolate the effect of the local composition of sectors. Moreover, I relate growth rates in San Francisco and San Diego with sector characteristics in Figure 2.7 and Table 2.3. To assess whether some sectors have grown faster than others due to shifting demand caused by overall economic growth leading to, for example, higher growth of sectors in initially richer tracts, I relate sector growth rates with estimated income elasticities. I find that sector growth rates are uncorrelated with estimated income elasticities. Similarly, if sectors that demand more high skilled workers grew faster, tracts with initially more high skilled residents could experience an increase in labor demand for high skilled workers. Again, I find that sector growth rates are uncorrelated with initial skill intensities.[45] Lastly, growth rates are negatively correlated with the initial citywide number of establishments across sectors, though the negative relationship is mostly due to few outliers (e.g. Dollar/Discount Stores are a relatively new sector).

At this point, it is instructive to draw parallels to previous work in this area. For example, Diamond (2016) identifies the labor supply elasticity (resident supply elasticity in my context) from local wage changes using shift-share labor demand shocks based on the initial industry composition and plausibly exogenous industry-level wage growth. In my analysis, I identify the corresponding elasticity from changes in local price indices with an instrument that is based on the initial sector

[44] Agarwal *et al.* (2018) find that the median credit card transaction occurs at nine kilometers from home. However, their results do not differentiate between more rural areas and dense urban areas like LA where travel distances are likely to be shorter.

[45] I compute skill intensities by taking the national share of high skilled workers in each sector over the total number of workers in the ACS microdata for 1990 and 2000.

composition of consumption varieties in a location and plausibly exogenous sector-level growth rates in varieties. In a nutshell, in both approaches the supply elasticity captures sensitivity to real incomes across locations. The difference is that I identify the resident supply elasticity from exogenous variation in price indices whereas Diamond (2016) identifies a similar elasticity from exogenous variation in nominal incomes.

In order to identify spillover elasticities δ_k, I require an additional source of exogenous variation that causes movements in the skill-mix surrounding a location. Intuitively, many shocks to exogenous amenities or tastes might lead to spatial correlated movements in the local population and the skill composition surrounding a tract. Suppose improvements in local school quality attract more high skilled residents into a cluster of neighborhoods. This amenity shock leads to correlation between changes in the skill-mix surrounding a tract in the cluster and changes in populations by skill in the tract itself. Hence, I interact the establishment shares in the average price shift-share instrument with the difference in sector expenditure shares by high and low skilled households derived from citywide income differences in 1990, estimated income elasticities of demand and citywide sales shares by sector $s_{city,t_0}(j)$ (see Table 2.2) according to

$$\Delta P^{IV}_{n,t} = \sum_j \left(\tilde{s}_{high,t_0}(j) - \tilde{s}_{low,t_0}(j)\right) \left(\sum_{n'} \frac{M_{n',t_0}(j)\mathbb{1}(distance_{nn'} < b)}{M_{t_0}(j)}\right) \log \hat{M}^O_t(j), \tag{2.30}$$

where

$$\tilde{s}_{k,t_0}(j) = \frac{s_{city,t_0}(j)\left(\frac{I_{k,t_0}}{I_{city,t_0}}\right)^{\nu_j}}{\sum_{j'} s_{city,t_0}(j')\left(\frac{I_{k,t_0}}{I_{city,t_0}}\right)^{\nu_{j'}}}.$$

The relative price instrument exploits differences in expenditure shares by skill group due to non-homothetic demand resulting in differential exposure of skill groups to growth in varieties across sectors. This leads to differential impacts on the goods price index, which affects the migration response of skill groups in and surrounding a tract. Lastly, I interact the relative price instrument with a skill dummy to recover type-specific spillover elasticities.

Results: Panel D of Table **??** reports the main regression results. First, I estimate regression 2.23 with both instruments and all controls under the assumption that goods prices are identical across the city as in equation 2.24 (column 2) or perfectly correlated with local rents in equation 2.25 (column 3). The main focus is on the estimate of $\delta_{high} - \delta_{low}$ since only the difference in spillover elasticities determines relative sorting patterns by skill group. My estimates of the relative spillover elasticities, $\delta_{high} - \delta_{low}$, are large in both specifications with values of around 1.7 and broadly in line with estimates from Diamond (2016) and Su (2018b).[46] The estimated spillover elasticity on low-skilled households is small and insignificant. Columns 1 and 2 in panels B and C of Table **??** show that the relative price instrument pushes up skill ratios in surrounding tracts in

[46]Despite using a different definition of amenity spillovers and working with MSAs instead of Census Tracts Diamond (2016) finds .7 for low skilled and 1.9 for the difference between high and low skilled. Su (2018b) estimates spillovers for low skilled to be .45 and 1.4 for the difference using US Census Tracts.

the first stage for $\hat{\mathscr{L}}_{kn,t}$. Moreover, the coefficient on the average price instrument is also positive and significant since an average improvement in access to consumption is more valuable to high skilled households because they consume more services than the low skilled.

Next, I jointly estimate $\frac{\kappa}{1-\eta}$ and spillover elasticities in equation 2.27 with both instruments treating changes in price indices of goods as endogenous. In column 5, I report my main results controlling for changes in income and rents, as well as natural amenities and population density in 1990. My estimate of the resident supply elasticity $\hat{\kappa} = 2.4$, conditional on $\eta = .493$ from Albouy *et al.* (2016), is comparable albeit smaller than existing estimates of similar elasticities.[47] The first stages for $\frac{s_{kn,t-1}(g)}{s_{kn,t-1}(h)} \log \hat{s}_{kn,t}(g)$ in panel A of Table **??** confirm that the price index instrument lowers the expenditure share on goods consistent with an increase in local varieties under goods and housing being complements. In column 4, I report results for regression in 2.27 without controls and find similar coefficients though the first-stage F-Statistic is smaller.

I can now compare estimated reduced-form spillover elasticities by skill groups in both models to assess the contribution of relative price index differences to explaining spatial sorting of skill groups. When I treat changes in price indices as endogenous my estimates of the difference in spillover elasticitites between high and low skilled households, $\delta_{high} - \delta_{low}$, as reported in columns 4 and 5, fall significantly by 30-50% relative to the model with exogenous price indices of consumption. This finding implies that the importance of reduced-form externalities from the skill composition is overestimated if endogenous spatial variation in price indices is not taken into account. Moreover, it suggests that local price index differences are quantitatively important in explaining observed sorting patterns by skill groups and similar in magnitude compared to reduced-form spillovers.

The estimate of κ and the difference in spillover elasticities between both models are broadly consistent across several robustness checks reported in Appendix Table 2.16. Changing the size of the distance buffer to 10km leads to almost identical results and to an even larger difference in $\delta_{high} - \delta_{low}$ in between models. Doubling the distance elasticity ψ in constructing the spillovers variable gives similar results, though the size of the spillover elasticity needs to be scaled down since spillovers are more concentrated locally. Using the skill distribution based on population over 25 instead of skill of household heads leaves the estimates unchanged. Next, I weight sector firm locations in constructing the average price instrument with sector sale shares and find similar results. Lastly, when I use expenditure shares on goods of renters instead of all households as sufficient statistic for price index changes, the difference in spillover elasticities between skill groups is comparable in both models; however, the average price instrument loses predictive power as shown by a low first stage F-Statistic.

Estimation of Firm Supply Elasticity θ

The shape parameter θ of the sector productivity distribution becomes the elasticity subject to which firms substitute between neighborhoods in response to differences in profits and, hence,

[47] Tsivanidis (2018), Su (2018b), Diamond *et al.* (2018) and Couture *et al.* (2019) estimate migration elasticities of 3-4.

differences in the demand for a sector in a neighborhood. To derive the relationship underlying the estimation of θ I begin with the model expression 2.10 for profits of a firm ω in location n and sector j in time t, replace profits per unit of productivity $\tilde{\pi}_n(j)$ with equation 2.11, and taking logs to get

$$\log \pi_{n,t}(j,\omega) = \frac{\sigma-1}{\theta}\log M_{n,t}(j) + \frac{\sigma-1}{\theta}\log\frac{\Pi_t(j)}{M_t(j)} - \frac{\sigma-1}{\theta}\log A_{n,t}(j) + (\sigma-1)\log z_{n,t}(j,\omega).$$

Since firm profits are not directly observable in the data, I rely on the model relationship between firm profits and firm size or employment, $l_n(j,\omega) = \frac{1}{w}(\sigma-1)\pi_n(j,\omega)$. Replacing constant and sector-specific terms with a sector-time fixed effect I estimate,

$$\log l_{n,t}(j,\omega) = \frac{\sigma-1}{\theta}\log M_{n,t}(j) + \iota_t(j) + u_{n,t}(j,\omega). \tag{2.31}$$

where $u_{n,t}(j) = -\frac{\sigma-1}{\theta}\log A_{n,t}(j) + (\sigma-1)\log z_{n,t}(j,\omega)$ is an error term. Identification of $\frac{\sigma-1}{\theta}$ is based on the location choice of firms. If many firms are observed in a location n, then the location must offer high profits for firms conditional on average productivity $A_n(j)$ and idiosyncratic productivity $z_{n,t}(j,\omega)$ in sector j. Specifically, location n has high demand for varieties in j. Therefore, we would expect a firm in n to be larger than a firm with the same idiosyncratic productivity in a location with less firms and the same average productivity as in n.

Data: To estimate structural equation 2.31 in the data, I use the census tract and employment for private, for-profit establishments in Los Angeles in 28 local sectors over the period 1992-2014 from NETS. Although NETS has been shown to capture the spatial firm distribution well in the cross-section (Barnatchez *et al.* (2017), Neumark *et al.* (2005)) a large share of employment numbers are imputed and the data cannot capture employment dynamics well. To overcome these limitations, I, first, use only establishments with directly reported employment numbers and, second, I restrict my analysis to the cross-sectional relationship of establishment size and counts of establishments.

Identification: There are two main identification concerns. First, equation 2.11 directly states that $M_n(j)$ is positively correlated with $A_n(j)$ in the error term causing downward bias of $\frac{\sigma-1}{\theta}$. Second, the regression suffers from selection bias, namely, if few firms are observed in a location the idiosyncratic productivities $z_{n,t}(j,\omega)$ in the error term must be high conditional on average productivity $A_n(j)$ leading to downward bias in the estimate of $\frac{\sigma-1}{\theta}$. To address the first concern, I include a tract-time fixed effect that captures location supply shocks like availability of retail space, labor market access, or commercial rents. Furthermore, I instrument the log number of establishments $\log M_n(j)$ with the log average slope in a location interacted with my estimates of the income elasticity of demand by sector, $\hat{v}_j$. The slope of a location is a very strong predictor of household income or skill composition in Los Angeles because households seem to prefer living in steeper locations, such as locations with a better view. Regressing log average household income or log skill ratio in a location on the log slope yields an R^2 of over 20% and highly significant positive coefficients. Conditional on tract-time and sector-time fixed effects the instrument picks up differential exposure of sectors to higher household income due to higher slope, a natural amenity highly valued by households. Since $\hat{\nu}_j$ are estimated parameters that are re-scaled as described in the previous section, I construct an alternative instrument replacing the value of $\hat{\nu}_j$ with the rank of ν_j across all sectors.

I address selection bias by comparing the size of similar establishments across locations. In particular, I assume that establishments ω^m of multi-establishment firms or chains m have a common productivity component independent of location,

$$z^m_{n,t}(j,\omega^m) = z^m_t(j,\omega^m).$$

The literature provides evidence that retail chains follow uniform pricing across stores (DellaVigna & Gentzkow (2019)), which is consistent with this assumption. By restricting the sample to chain establishments, defined as having the same headquarter in the NETS data I can include a chain-time fixed effect and identify $\frac{\sigma-1}{\theta}$ only with variation across locations serviced by the same chain.[48] Conditional on the assumption of common productivity, variation in size of establishments within chain across locations is either due to local demand, consequently due to variation in the number of establishments, or average sector productivity differences. By instrumenting the number of establishments in a location with the slope instrument, I isolate how demand differences affect relative sizes of establishments within the same chain.

Results: Panel A of Table 2.5 reports first stage results of regressing the log number of establishments on log average slope interacted with ν_j. Across all specifications I find a strong positive relationship between the instrument and the number of establishments in an area indicating that establishments in sectors with high income elasticity locate in tracts with high slope and, consequently, close to high income residents. Panel B shows the second stage estimates of 2.31 and the main estimate is reported in column 1. Doubling the number of establishments in a location implies that establishments are 25% larger on average. This estimate informs the baseline value of $\frac{\theta}{\sigma-1} = 4$ in my model calibration. With my assumed elasticity of substitution within sectors, $\sigma = 5$, the implied supply elasticity θ takes a value of 16. When I split the sample into retail and service sectors in columns 3 and 4, I find similar point estimates; although, the point estimate for retail sectors becomes insignificant. Results with the rank-based instrument are reported in column 4 and imply a slightly larger value of θ.

Table 2.6 summarizes all parameters of the model. I set up two different versions. In the model with endogenous relative price index differences, my baseline specification denoted with superscript 1, shopping frictions restrict access to consumption establishments. In the second calibration without endogenous price effects, akin to the previous literature, households can consume everywhere in the city at no cost; however, high skilled households are subject to higher reduced-form spillovers. Since the point estimate of the spillover elasticity for the low skilled is indistinguishable

[48] Some multi-establishment firms operate in several sectors in the NETS data. I restrict each chain to the sector with most establishments in each year.

from zero in my estimation, I set $\delta_{low} = 0$ in both calibrations. In the baseline calibration, I use the estimated spillover elasticity of δ_{high}, reported in column 5 of Table **??** (panel D). For the model without relative price index effects, the corresponding estimates of δ_{high} in columns 2 or 3 fall into a parameter range for which the equilibrium is not necessarily unique. Specifically, if δ_{high} exceeds 1.3 (given all other parameters), the initial equilibrium in the data becomes unstable.[49] To ensure the model without price index effects has a unique solution, I choose a value of $\delta_{high} = 1.25$, just below this cutoff.

Equipped with the full set of model elasticities, I can invert the model to recover the fundamentals of the economy such as fixed amenities and sector productivities. Then, I can perform model-based counterfactuals of different policy shocks to assess the implications of allowing for two-sided sorting for household welfare and sorting.

Model Inversion

To be able to perform counterfactuals in the model, I require location-specific exogenous amenities by skill, $\bar{B}_{kn}$ and sector-location productivities $A_n(j)$ as well as fixed cost of entry by sector and sector demand shifters. My model falls into the set of quantitative urban economics models (e.g. Tsivanidis (2018), Monte *et al.* (2018), and Ahlfeldt *et al.* (2015)) that are fully saturated with structural residuals or "fundamentals", which cover all variation in the data unexplained by the inherent model structure. Thus, with sufficient data moments I can invert the model to recover the set of residuals as stated in the next proposition.

Proposition 2. *Given data on residents by skill and location, L_{kn}, number of firms by sector and location, $M_n(j)$, citywide revenue shares by sector, $rs_c(j)$, citywide expenditure share on goods, $s_c(g)$ and a normalized wage per unit of labor $w = 1$ there exist unique vectors of model fundamentals, namely, exogenous amenities $\bar{B}_{kn}$, composite demand and productivity shifters $\bar{A}_n(j) = A_n(j) a_g^{\theta\frac{1-\sigma}{1-\eta}} \alpha_j^{\theta\frac{1-\sigma}{1-\gamma}}$, fixed entry costs by sector $f^e(j)$ and transfers T_k that replicate the observed equilibrium in the data.*

The process follows several steps. First, using the assumption that all housing expenditures are redistributed to households according to the labor endowment ρ_k, normalized citywide wages transfers are characterized by

$$T_k = \rho_k \frac{1 - s_c(g)}{s_c(g)}.$$

With values for transfers and wage equalization due to free labor mobility, I can directly compute nominal income $I_k = \rho_k w + T_k$. Combining the fact that the city is closed so citywide expenditure

[49] If I slightly perturb the observed household and firm distributions as starting values in the computation, the model fails to converge to the observed equilibrium. Instead, the model finds alternative configurations of the city. When I simulate counterfactuals in the model, this multiplicity makes it difficult to separate the effects of a policy from such an alternative equilibrium. I want to emphasize that the multiplicity of equilibria in an urban context is an interesting area of research; however, it is beyond the scope of this paper.

on goods need to equal citywide income and the free entry condition of firms implies

$$f^e(j) = \frac{1}{\sigma} \frac{rs_c(j) \sum_k \rho_k w L_k}{M(j)}.$$

Next, rearranging equation 2.11 and combining with free entry in 2.15 gives a set of conditions that allows me to recover the set of demand-productivity composites $\mathscr{A} = \{\bar{A}_n(j)\}_{n,j}$ for all sector and locations, namely,

$$\begin{aligned} \bar{A}_n(j) &= \xi_n(j) \left(\sum_{n'} \bar{A}_{n'}(j) \breve{\pi}_{n'}(j)^{\theta} \right) \breve{\pi}_n(j)^{-\theta} && \forall j, n \\ \sum_{n'} \bar{A}_{n'}(j) \breve{\pi}_{n'}(j)^{\theta} &= (f^e(j))^{\theta} && \forall j \end{aligned}$$

where

$$\breve{\pi}_n(j) = c_1 \sum_{n'} \sum_{k'} \tau_{n'n}^{-1} \tilde{p}_{n'}(j)^{\frac{\sigma-\gamma}{1-\gamma}} \tilde{P}_{k'n'}(g)^{\frac{\gamma-\eta}{1-\eta}} P_{k'n'}^{-\varepsilon_g - \nu_j} I_{k'}^{\eta+\varepsilon_g+\nu_j} L_{k'n'}$$

and $\tilde{p}_n(j) = a_g^{\frac{1-\gamma}{1-\eta}} \alpha_j p_n(j)^{1-\gamma}$, $\tilde{P}_{kn}(g) = a_g P_{kn}^{1-\eta}$ and c_1 is a constant. Since all prices are themselves functions of $\mathscr{A}$ and moments in the data, the system maps the vector of fundamentals $\mathscr{A}$ onto itself. Intuitively, total profits by sector in the economy are given by the fixed cost of entry, hence, bounding and normalizing the set of possible parameters. A formal proof is still work in progress.

In the previous step of the inversion, I recover price indices P_{kn} and nominal income I_k from earlier, hence, real consumption by skill group and location, U_{kn}. With this information and an appropriate normalization, I can rearrange equation 2.8 to solve for the unique set of exogenous amenities $\bar{B}_{kn}$, similar to the previous step.

Model Fit

After estimating key parameters and recovering the fundamentals of the economy, I can exactly recover the moments used in the inversion and compare other moments produced by the simulated model to non-targeted moments in the data. The left panel of Figure 2.8 plots residential rents as simulated in the baseline calibration with local price effects against residential rent per square foot computed as total housing expenditure divided by housing stock in the ACS 2014. The two series are highly correlated (Correlation=.58); although, the model finds a larger spatial dispersion in rents. The right panel plots model based rent per square foot against median rental prices per square foot from Zillow's zip code level data from 2014-2016. Again the series are quite correlated with a value of .43.

In Figure 2.9, I compare expenditure shares on goods from ACS used in the estimation earlier against model predicted expenditure shares for low skilled (left) and high skilled (right). With a correlation of .46 the model performs quite well in replicating the spatial distribution of non-targeted expenditure shares. In columns 5 and 6 of Table 2.14, I report results from regressing the housing expenditure share in tract-level and model-produced data on a skill dummy and tract-FX. Although not specifically targeted, the model reproduces the non-homotheticity in housing demand, if anything understates it compared to the data. Lastly, I can compare the model's prediction for the ratio of expenditure shares between high and low skilled. The model has problems capturing the large spatial variation in the data; however, the model understates the differences between high and low skilled implying that my calibration is conservative. Reassuringly, correlations with the data are consistently smaller when the model is simulated without relative price index effects but larger external spillovers.

2.8 Policy Counterfactuals

With the estimated model I can now assess how urban policies interact with the two different sorting forces in the model, relative price indices and reduced-form spillovers. My counterfactual analysis tries to answer two questions. First, through the lens of my model, what are the effects of two real world urban policies, Opportunity Zones and social housing, in terms of mobility of skill groups, inequality, and aggregate welfare in Los Angeles? Second, do we miss important details of these policies if we treat the endogenous response of households (and firms) to changing neighborhoods as a reduced-form spillover as opposed to changing costs of living? To answer these questions, I perform two counterfactuals in my model.

First, I shock the spatial distribution of firms by simulating the effect of a new tax incentive to invest in a subset of locations in LA to investigate the endogenous response of skill groups to changing price indices of consumption. I do the reverse in the second counterfactual. I shock the distribution of households by adding the existing stock of Social Housing projects to the city from an initial counterfactual equilibrium without such projects in order to understand the general equilibrium reaction of firms. For each counterfactual, I recover the fundamentals of the economy (exogenous amenities and productivities) using the same moments in the census tract data for Los Angeles in 2014, but assume that sorting of skill groups is driven either by relative price indices and "true" estimated reduced-form spillovers (Calibration 1) or only "biased" reduced-form spillovers and no shopping frictions (Calibration 2).[50] Then, I simulate the same policy shocks in both versions of the model fitted to the same observed initial equilibrium.

Opportunity Zone Program

In 2017, U.S. Congress passed the Tax Cuts and Jobs Act and among several tax-related policies, created a new place-based tax credit, the so-called Opportunity Zones (OZ). The stated goal of the program is to lift living standards in economically disadvantaged urban areas by incentivizing businesses and investors to invest unrealized capital gains in designated opportunity zones.[51] U.S.

[50] To further trace out how the predictions of my policy counterfactuals change, I also present results removing and adding other features of the model (non-homothetic preferences, spatial frictions or reduced-for spillovers).

[51] The tax cut has three parts. First, capital gains taxes from previous investments can be deferred until 2026 when reinvested in OZs. Second, the tax base of previous investments increases up to 15% depending on the duration of the

Treasury Secretary Steven Mnuchin recently stated that he expects the total investment in Opportunity Zones to exceed $100B in 2019. According to Theodos *et al.* (2018) out of a total of 42,176 eligible census tracts 8,762 were designated Opportunity Zone status.[52] My Los Angeles sample includes 257 of these tracts. Figure 2.11 shows all 257 OZ census tracts in Los Angeles. Designated census tracts are concentrated in the center of Los Angeles with some scattered zones in the periphery. Table 2.17 reports means and differences for OZ and non-OZ tracts. As expected, OZs are populated by lower-income, less educated households and host around 20% less firms in sectors with high income elasticity relative to non-OZ tracts.

The Economic Innovation Group (2019) estimates that due to the tax benefits OZ investments offer, there should be excess returns of around 30% over 10 years. Based on this estimate, I implement a 30% subsidy on variable profits $\tilde{\pi}_{OZ}(j)$ of firms in all sectors that operate in OZ tracts in the model and assume that the subsidy is financed by the city government that due to the policy redistributes a smaller transfer to all households.[53] Then I simulate the model to predict how mobile firms respond to exogenous profit subsidies in OZs and evaluate the endogenous effect on location choices of skill groups in response to the change in access to consumption varieties in general equilibrium.

Before moving on to the results, I want to discuss a few limitations of this policy counterfactual. First, since labor is freely mobile in the model, my counterfactual cannot speak to the local labor market effects of place-based tax incentives, an interesting and well-studied question. Instead, I focus on isolating the demand-side effects on household composition of attracting more consumption varieties into disadvantaged neighborhoods. Consistent with this notion, Reynolds & Rohlin (2015) argue in recent work on Federal Empowerment Zones, a broadly similar place-based policy, that firm-level incentives made targeted neighborhoods more attractive to high-income, well-educated households. They fail to explain this finding with employment effects on the initial resident population. Secondly, there have been concerns in the media that Opportunity Zones predominantly lead to more investment in high-end real estate, a channel currently absent from the model since I treat the residential housing stock as fixed. However, a recent investor survey by KPMG reports that 39% of potential investors plan on operating a business in OZs and quotes by investment fund managers indicate that developing amenities like retail venues is crucial for creating value in OZs and returns to business investments might be much more profitable than real estate.[54] Lastly, since the Opportunity Zone program has been implemented very recently and the tax bill requires little to no reporting on take-up or costs, I cannot validate any of my counterfactual outcomes using moments in the data.[55] With this caveat in mind I view the counterfactual as

OZ investment. Lastly and most importantly, after 10 years all gains from OZ investment are excluded from taxation.

[52] Eligibility is based on high poverty rate and low family income, (see Theodos *et al.* (2018) for details). State Governors propose the final selection of OZs among the eligible tracts to the U.S. Treasury for approval.

[53] Although the tax benefits are ultimately financed by all U.S. tax payers (current and future) I implicitly assume that the share of LA tax payments in the total tax bill of the reform is equal to LA's share in the population.

[54] *"..the firm plans to cluster investments in individual neighborhoods to create a critical mass of amenities, such as housing and grocery stores, that can increase property values in the area."*, Garett Bjorkman, CIM Group CEO; *"The returns on investing in a high potential company that sets up as a qualified opportunity zone business [...] could be 10 times more profitable than flipping commercial real estate."*, Brian Phillips, Founder Pearl Fund

[55] In future work I plan to validate effects based on similar past policies, like Empowerment Zones.

a prediction of the policy's effects.

Table 2.7 summarizes the main results for the calibration with endogenous price effects. In reaction to the profit subsidy, firms move into OZs such that the number of varieties increases by 80% percent, around one standard deviation of the overall variation in the number of firms across tracts.[56] However, the reallocation of economic activity is at the expense of nearby tracts as shown in Figure 2.13, reducing the local effect on access to varieties. Price indices of consumption fall for low and high skilled households but more so for the latter, leading to almost twice as many high skilled households relocating to OZs due to firm subsidies. The Opportunity Zones program induces gentrification in these disadvantaged areas due to demand-driven effects. There are three reasons for the stronger response of high-skilled households to the policy. First, since OZs are initially populated by relatively few high skilled households, the policy makes OZs more attractive to this group relative to other areas and this effect is larger than for the low skilled.[57] Second, high skilled households value consumption of goods more than low skilled because goods demand is more elastic than housing; hence, a fall in prices benefits high skilled more. Lastly, more firms in sectors with high income elasticity enter endogenously, as shown in Table 2.20, further lowering the price index for high skilled more than for low skilled households. Finally, high skilled households are subject to positive reduced-form spillovers. Both channels amplify gentrification of OZs.

As indicated by the modest R^2 in Table 2.7, effects of the policy are not limited to OZ tracts. Non-OZ tracts are affected by the reallocation of firms into OZs because of two features of the model. Figure 2.17 plots the movements into Non-OZ tracts as a function of the share of shopping that occurs in OZs initially. At around 25% of consumption spending in OZs, the effect of the policy on non-OZ locations is the same as for OZ locations, showing a similar difference in mobility of low and high skilled households. Moreover, I set up reduced-form spillovers to be a function of the location of a tract. Hence, high skilled households in non-OZs close to targeted areas receive spillovers from the gentrification of OZs. Figures 2.22 and 2.23 show that changes in the price index of goods consumption and reduced-form spillovers for high skilled households are highly correlated, but they are not identical and not limited to OZs.

In comparison, Table 2.8 reports the mobility effects of the same policy on households and firms for the model with no shopping frictions but larger reduced-form spillovers. This version of the model predicts almost no movement of households, since location of consumption venues has no bearing on the price index of consumption if shopping is frictionless.[58] Despite being a stark example, the counterfactual emphasizes that endogenous price index effects as opposed to reduced-form spillovers, are not equivalent drivers of spatial inequality in a counterfactual sense and can lead to first-order differences in the outcomes of a common urban policy in terms of inequality or welfare.

[56] There is also entry of firms into the city due to the policy of around 4.5% on average in all local sectors.

[57] This effect is present even without non-homothetic preferences and is due to idiosyncratic preferences for locations. As the policy targets low skilled neighborhoods expected utility falls more (or increases less) for high skilled, increasing the relative attractiveness between the targeted areas and the average larger, hence stronger movements into these areas.

[58] The small movements into OZs in calibration 2 can be explained as follows: Overall welfare falls due to the policy making initially poor OZs more suited to the marginal resident in both groups.

Table 2.9 summarizes spatial inequality and the welfare effect of Opportunity Zones on low and high skilled residents of LA for both calibrations and four other versions, two calibrations without spillovers and two calibrations with homothetic preferences. Spatial inequality effects are strongest in the baseline model and mostly driven by price index effects when compared to the full model without spillovers. In the models with homothetic preferences, the policy does not affect the skill ratio in OZ and Non-OZ differently such that both groups locate in OZ but in equal proportions. Lastly, I turn to welfare effects, measured as the compensating variation on expected utility by each group. The policy leads to modest welfare losses of both groups in all calibrations. The benefits of increased variety in OZs cannot compensate for the lump-sum taxes imposed by the city to finance the subsidy. In the baseline calibration, losses are slightly smaller for the high skilled because the policy leads to stronger within-group spillovers from centrally located OZs. Welfare losses are lowest in the baseline model because Opportunity Zones are targeting within-group populations with high marginal utility as the areas lack access to consumption varieties; hence, they face high consumption prices in the initial equilibrium.[59] In a model with only reduced-form spillovers, we would miss two crucial effects of Opportunity Zones. First, the reallocation of firms in the city affects relative price indices of consumption resulting in reallocation of skill groups; specifically, the influx of high skilled households leads to gentrification of OZs. Second, welfare effects of the policy are different since Opportunity Zones are targeting specific locations that are disproportionately benefiting from more consumption varieties.

Social Housing in Los Angeles

Around 100 thousand out of 3.3 million housing units in Los Angeles receive state or federal housing assistance either in the form of direct housing projects, loans or tax credits for low-income households.[60] For this counterfactual, I rely on data from the California Housing Partnership, an affordable housing think-tank, that provides address-level rent savings by social housing project. Since I cannot distinguish who receives the assistance in the data, but it is common in these policies that affordable housing eligibility is based on income, I aggregate rent savings to the census tract level and assign all savings to the housing expenditure of low skilled households that I observe in the ACS data. Figure 2.12 shows the distribution of social housing units as the share of housing costs of low skilled households covered by rent savings. Social housing in LA is fairly spread out over the city with a higher concentration in the dense center of LA. Since social housing projects tend to be large with several hundred units in small areas, intensity in terms of covered share of expenditure by low skilled households is concentrated in a few tracts.

[59] To a first order, the difference between homothetic and non-homothetic preferences can be explained by average marginal utility in the first four models which I calibrate around .55 as compared to 1 with homothetic preferences.

[60] Source: California Housing Partnership Preservation Database, June 2019. For more information, please visit chpc.net/policy-research/preservation/ ; The California Housing Partnership has provided address-level savings from social housing. The data covers housing assistance from HUD (Project-based Section 8, Project Rental Assistance Contract, Section 202 Direct Loans, Insurance Programs), Low Income Housing Tax Credit and USDA. However, the database does not include other state or local programs. For more detail, see https://chpc.net/affordable-housing-benefits-map/

I implement social housing in the model as subsidy on rents for low skilled households equivalent to the share of rent savings in housing expenditure (shown in Figure 2.12) and financed by the city government budget. First, I invert the model from the observed equilibrium assuming that social housing is present in the current equilibrium and then remove all social housing in a counterfactual. Hence, I report all outcomes going from the counterfactual to the observed equilibrium to assess the effects of the current state of social housing in LA.

In Tables 2.10 and 2.11, I report changes in population by skill and number of firms as a function of the share of rent savings for the baseline calibration and the calibration with only reduce-form spillovers, respectively.[61] In all calibrations social housing leads to higher market rents in targeted tracts. This causes high skilled households to move out because they face the market rent as opposed to the low skilled who locate in targeted areas due to subsidized rent. The skill ratio in targeted areas at mean subsidy level (3.25%) falls by around 3.5%. Figures 2.15 and 2.16 show the reallocation of firms and high skilled households away from the center of LA towards the periphery of the city. In the calibration with price effects, firms, on average, leave areas with social housing leading to an increase in price indices for both groups. However, firms in sectors with low income elasticity leave targeted areas much less than firms in income-elastic sectors as shown in Figure 2.18, endogenously leading to stronger price index changes for the high skilled. Responses by firms to the subsidy are quantitatively limited because of two opposing forces on local firm profits and, consequently, firm mobility. First, low-skilled households become richer in targeted areas due to the subsidy, shifting demand to goods consumption in general and towards income-elastic sectors. Second, the remaining high skilled households behave more similar to poorer households due to high market rents. The first increases profits in areas with social housing and more so in income-elastic sectors, the reverse holds for the second.

Table 2.12 shows that the effects on the skill ratio are broadly similar in both calibrations. It is instructive to consider the effects of the policy on neighborhood composition in the models without spillovers, shown in columns 3 and 4 of Table 2.12. The model without spillovers and price effects describes these effects, absent any endogenous amplification; whereas the model with price index effects traces out the amplification of the policy due to endogenous changes in consumption access in response to the outflow of high skilled households.[62] The firm response reinforces the shift towards less skilled neighborhoods and slightly increases the welfare loss for the high skilled. On the one hand, this welfare loss is small since high skilled households in these areas are now poorer and behave more similar to a low skilled household. On the other hand, spillovers amplify welfare losses for the high skilled since the areas with rent subsidies, which are relatively central in LA, provide less amenity spillovers to the rest of the city. Hence, the model with large spillovers overstates the welfare loss of the high skilled and understates the welfare gain of the low skilled compared to the baseline model.

---

[61] The numbers are reflecting percent changes in number of firms and HHs if a tract has 100% subsidy on rent due to social housing for the low skilled. In fact, the mean subsidy conditional on hosting any social housing is 3.25% with 19% at the 99th percentile.

[62] Price index channel and reduced-form spillovers reinforce each other. Hence, a comparison of column 3 and 4 in Table 2.12 does not fully capture the total amplification from the price index channel.

2.9 Conclusion

Spatial inequality in cities has sparked interest by the public and policy makers. To inform urban policies, it is important to understand its sources, in particular, how the composition of local residents endogenously shapes the attractiveness of a neighborhood. This paper studies how two-sided sorting of heterogeneous households and firms generates pecuniary externalities that amplify clustering of household groups in cities. As a benchmark, I compare these forces to reduced-form spillovers that have been studied in previous literature. First, I develop a quantitative general equilibrium model of the city that features two-sided sorting of skill groups and firms in various local consumption sectors but nests previous work. Second, I combine the model with rich administrative microdata from Los Angeles to quantify the contributions of these pecuniary externalities and reduced-form spillovers to observed sorting patterns in the data. Third, I assess the implications of urban policies when allowing for two-sided sorting.

I find that two-sided sorting is an important driver of spatial inequality in Los Angeles. Spatial variation in local price indices reduces the estimates of reduced-form spillover elasticities by 30-50%. In the first policy counterfactual, I show that subsidizing firm entry in specific locations leads to heterogeneous location choices of skill groups and welfare due to differential changes in relative price indices of consumption. In response to rent subsidies that target specific locations and groups, firms also relocate thereby amplifying the sorting of households and welfare effects of the policy. In a model that relies on reduced-form amenity spillovers to generate strong sorting patterns, these effects are absent since household and firm location choices are independent. In addition, such a model overstates the welfare losses (or understates the gains) of the urban policies. My results suggest that demand linkages between different household groups as well as across neighborhoods are important determinants of spatial inequality in cities and have profound implications for our understanding of urban policies.

2.10 Appendix

Properties of Household Preferences

The following analysis is under the assumptions of given prices and from the view of an individual household of any type. To save on notation I omit location and type subscripts
First, let's look at the price index of goods responds to changes in real consumption

$$\frac{U\partial P(g)}{P(g)\partial U} = \frac{1}{1-\gamma}\sum_{j=1}^{J}\tilde{s}(j)\nu_j = \frac{\bar{\nu}}{1-\gamma}$$

where $\bar{\nu}$ is the expenditure share weighted income elasticity of demand parameter across sectors inside the goods sector. Second, I can compute the expenditure elasticity with respect to real consumption

$$\begin{aligned}\frac{U\partial E}{E\partial U} &= \frac{1}{1-\eta}I^{\eta-1}\left(a_h r^{1-\eta}U^{\varepsilon_h}\varepsilon_h + a_g P(g)^{1-\eta}U^{\varepsilon_g}\varepsilon_g + a_g P(g)^{1-\eta}U^{\varepsilon_g}(1-\eta)\frac{U\partial P(g)}{P(g)\partial U}\right)\\ &= \frac{1}{1-\eta}\left(s(h)\varepsilon_h + s(g)\left(\varepsilon_g + \frac{1-\eta}{1-\gamma}\bar{\nu}\right)\right) = \frac{\bar{\varepsilon}}{1-\eta}\end{aligned}$$

where $\bar{\varepsilon}$ is the expenditure share weighted average income elasticity of demand parameter across housing and goods.
With the above result, I can compute the expenditure elasticity of demand for housing

$$\frac{\partial \log C(h)}{\partial \log E} = \eta + \varepsilon_h \frac{\partial \log U}{\partial \log E} = \eta + (1-\eta)\frac{\varepsilon_h}{\bar{\varepsilon}}$$

and goods,

$$\frac{\partial \log C(g)}{\partial \log E} = \eta + (1-\eta)\left(\frac{\varepsilon_h - \frac{\eta}{1-\gamma}\bar{\nu}}{\bar{\varepsilon}}\right).$$

For the expenditure elasticity of demand for a particular sector in the service industry, it holds that

$$\frac{\partial \log C(j)}{\partial \log E} = (\gamma-\eta)\frac{\log P(g)}{\log E} + \eta + (\nu_j + \varepsilon_g)\frac{\partial \log U}{\partial \log E} = \eta + (1-\eta)\frac{\varepsilon_g + \nu_j}{\bar{\varepsilon}} + (1-\eta)\frac{(\gamma-\eta)}{(1-\gamma)}\frac{\bar{\nu}}{\bar{\varepsilon}}$$

Next, we can compute the mobility elasticity with respect to income. Recall

$$\lambda_n = \frac{B_n U_n^{\kappa}}{\sum_{n'} B_{n'} U_{n'}^{\kappa}}$$

So,

$$\frac{E\partial\lambda_n}{\lambda\partial E} = \kappa \left.\frac{E\partial U}{U\partial E}\right|_n - \frac{E\partial\Phi}{\Phi\partial E} = \kappa(1-\eta)\left(\frac{1}{\bar{\varepsilon}_n} - \sum_{n'}\lambda_{n'}\frac{1}{\bar{\varepsilon}_{n'}}\right)$$

These elasticities imply the following:

- Engel aggregation: $s(h)\frac{\partial \log C(h)}{\partial \log E} + s(g)\sum_j \tilde{s}(j)\frac{\partial \log C(j)}{\partial \log E} = 1$

- Conditional on prices income elasticities of demand parameters $\varepsilon_g, \upsilon(q)$ are defined up to scale. Consumption choices are not affected by scaling the parameters by a constant factor. Furthermore, if κ is scaled by the same factor agents mobility choices are unaffected.

- As a result of the above I can normalize one income elasticity parameter and one taste shifter without affecting the economic choices of agents.

- Sufficient: If $0 < \eta < 1$ and $\gamma > 1$ then $\varepsilon_i > 0, \forall i \in \{h, g\}$ and $\bar{\nu} < 0$ such that utility is increasing in expenditure and the inner price index is increasing in expenditure.

- 1. $\varepsilon_i = 1 - \eta, \forall i$ and $\nu_j = 0, \forall j$: preferences are homothetic nested CES, many trade models
  2. $\varepsilon_i = 1 - \eta, \forall i$ and $\exists \nu_j \neq 0$: upper nest is homothetic and within sectors non-homothetic, Borusyak & Jaravel (2018)
  3. $\varepsilon_i \neq 1 - \eta, \forall i$ and $\nu_j = 0, \forall j$, upper nest is non-homothetic and lower nest homothetic, Comin et al. (2018), Matsuyama (2018)

- In the case of homothetic upper nest ($\varepsilon_g = 1 - \eta$): $\frac{U\partial E}{E\partial U} = 1 + s(g)\frac{\bar{\nu}}{1-\gamma}$

Proofs

Proof of Proposition 1

Proof. The proof is straightforward and can be found in a similar form in Matsuyama (2019). Recall the expression for the expenditure share on goods from sector j in location n' by household k in n and taking logs

$$\log s_{knn'}(j) = \log a_g + \log \alpha_j + (\gamma - \eta)\log P_{kn}(j) + (\eta - 1)\log I_k + (\varepsilon_g + \nu_j)\log U_{kn} + (1-\gamma)\log P_n(j)$$

Taking prices and nominal income as given, I take the derivative with respect to U_{kn}

$$\frac{\partial \log s_{knn'}(j)}{\partial U_{kn}} = \frac{1}{U_{kn}}\left(\varepsilon_g + \nu_j + \frac{\gamma - \eta}{1-\gamma}\bar{\nu}_{kn}\right)$$

Note that $s_{knn'}(j)$ is increasing in real consumption if $\varepsilon_g + \nu_j > \frac{\gamma-\eta}{1-\gamma}\bar{\nu}_{kn}$ which captures the property that as household get richer they allocate more spending to sector with higher income elasticity. For any $\nu_1 > \nu_2$,

$$\frac{\partial \log s_{knn'}(1)}{\partial U_{kn}} - \frac{\partial \log s_{knn'}(2)}{\partial U_{kn}} = \frac{1}{U_{kn}}(\nu_1 - \nu_2) > 0.$$

This establishes log-supermodularity of $s_{knn'}(j)$ in U_{kn} and ν_j. The result holds by the same logic for $\tilde{s}_{kn}(j)$. □

Proof of Corollary 1

Proof. Given $U_{high,n} > U_{low,n}$ Proposition 1 implies for any $\nu_1 > \nu_2$

$$\frac{s_{high,n}(1)}{s_{high,n}(2)} > \frac{s_{low,n}(1)}{s_{low,n}(2)}.$$

With $I_k > 0, \forall k$

$$\frac{\pi_{high,n}(1)}{\pi_{high,n}(2)} > \frac{\pi_{low,n}(1)}{\pi_{low,n}(2)}$$

where $\pi_{kn}(j) = s_{kn}(j)I_k$. We want to show for any $x_n > x_n'$ and $\nu_1 > \nu_2$

$$\frac{\pi_{high,n}(1)x_n + \pi_{low,n}(1)(1-x_n)}{\pi_{high,n}(1)x_n' + \pi_{low,n}(1)(1-x_n')} > \frac{\pi_{high,n}(2)x_n + \pi_{low,n}(2)(1-x_n)}{\pi_{high,n}(2)x_n' + \pi_{low,n}(2)(1-x_n')}$$

Note that the left hand side is increasing in $\pi_{high,n}(1)$ since $x_n > x_n'$. Aplying the log-spm of $\pi_{kn}(j)$ we can write

$$\frac{\pi_{high,n}(1)x_n + \pi_{low,n}(1)(1-x_n)}{\pi_{high,n}(1)x_n' + \pi_{low,n}(1)(1-x_n')} > \frac{\frac{\pi_{low,n}(1)\pi_{high_n}(2)}{\pi_{low,n}(2)}x_n + \pi_{low,n}(1)(1-x_n)}{\frac{\pi_{low,n}(1)\pi_{high_n}(2)}{\pi_{low,n}(2)}x_n' + \pi_{low,n}(1)(1-x_n')} = \frac{\pi_{high,n}(2)x_n + \pi_{low,n}(2)(1-x_n)}{\pi_{high,n}(2)x_n' + \pi_{low,n}(2)(1-x_n')}$$

This completes the proof. □

Proof of Corollary 2

Proof. The proof uses results from Athey (2002) on monotone comparative statistics of sums of log-spm functions. I can write the goods price index as

$$P(U_{kn}, x_n)^{1-\gamma} = \sum_j \underbrace{\alpha_j U_{kn}^{\nu_j}}_{=f(U_{kn},\nu_j)} \underbrace{P_n(j,x_n)^{1-\gamma}}_{=u(\nu_j,x_n)}$$

Theorem 1 in Athey (2002) states that iff $f(U_{kn}, \nu_j)$ is log-spm in U_{kn} and ν_j a.e. and $u(x_n, \nu_j)$ is log-spm in x_n and ν_j a.e. then $P(U_{kn}, x_n)^{1-\gamma}$ is log-spm in U_{kn} and x_n a.e. To show log-spm of $u(x_n, \nu_j)$ I start with equation 2.14 implies

$$\frac{M_n(j)}{L_n} = \frac{\pi_n(j)}{L_n f^e(j)}.$$

By corollary 1 $\frac{\pi_n(j)}{L_n f^e(j)}$ is log-spm in x_n and ν_j, hence $\frac{M_n(j)}{L_n}$ is log-spm in x_n and ν_j.[63] Specifically for $x_n > x_n'$ and $\nu_1 > \nu_2$,

$$\frac{M_n(1;x_n)}{M_n(1;x_n')} > \frac{M_n(2;x_n)}{M_n(2;x_n')}$$

[63]Dividing by a positive constant $f^e(j)$ does not affect log-supermodularity.

Applying equation 2.12 under the assumption that shopping frictions are infinite outside n we can directly see that

$$\frac{P_n(1;x_n)^{1-\gamma}}{P_n(1;x_n')^{1-\gamma}} > \frac{P_n(2;x_n)^{1-\gamma}}{P_n(2;x_n')^{1-\gamma}}$$

and conclude that $P_n(j;x_n)^{1-\gamma}$ is log-spm in x_n and ν_j.
Lastly, log-supermodularity of $f(U_{kn},\nu_j)$ is given by proposition 1. Hence, I can apply therorem 1 in Athey (2002) and conclude that $P(U_{kn},x_n)^{1-\gamma}$ is log-spm in x_n and U_{kn}. $\square$

2.11 Figures and Tables

Figure 2.1: Spatial inequality as measured by skill ratio in Los Angeles, 2014

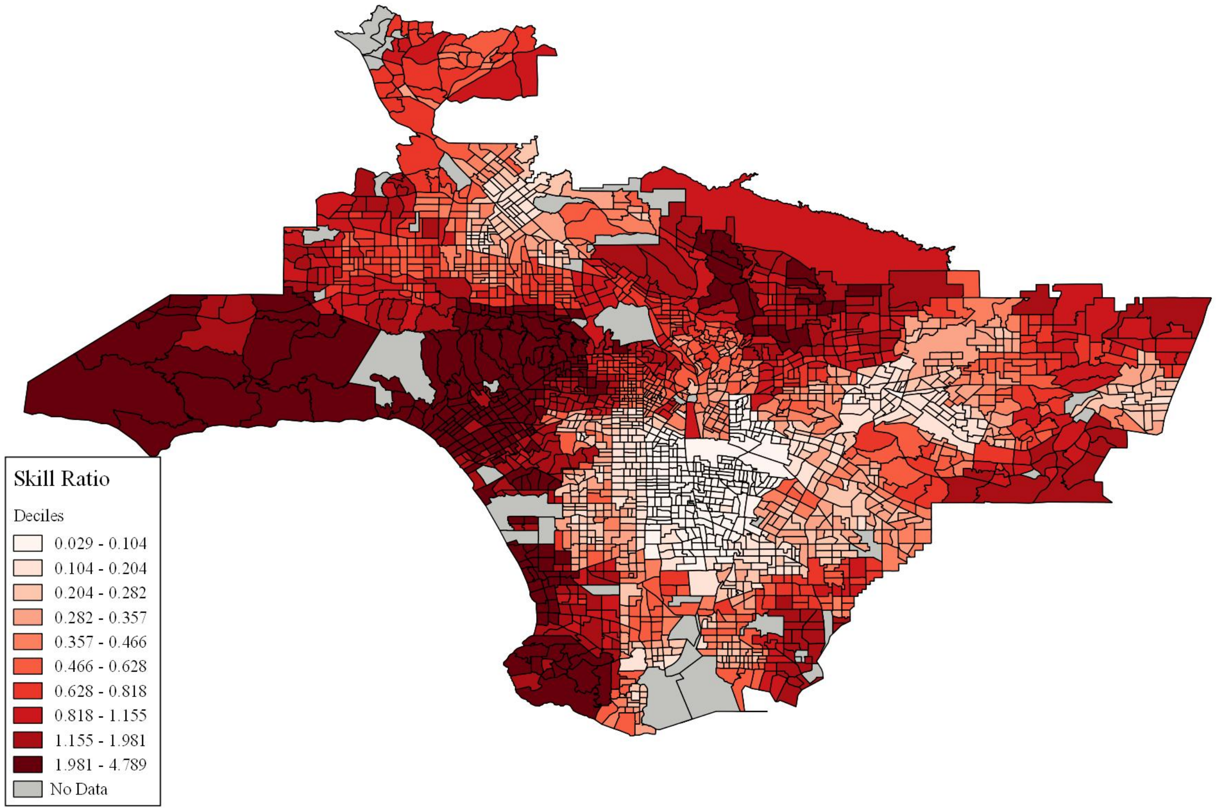

The figure plots the skill ratio in each census tract in Los Angeles, ACS 2014.

Figure 2.2: Distribution of establishments by income elasticity in Los Angeles, 2014

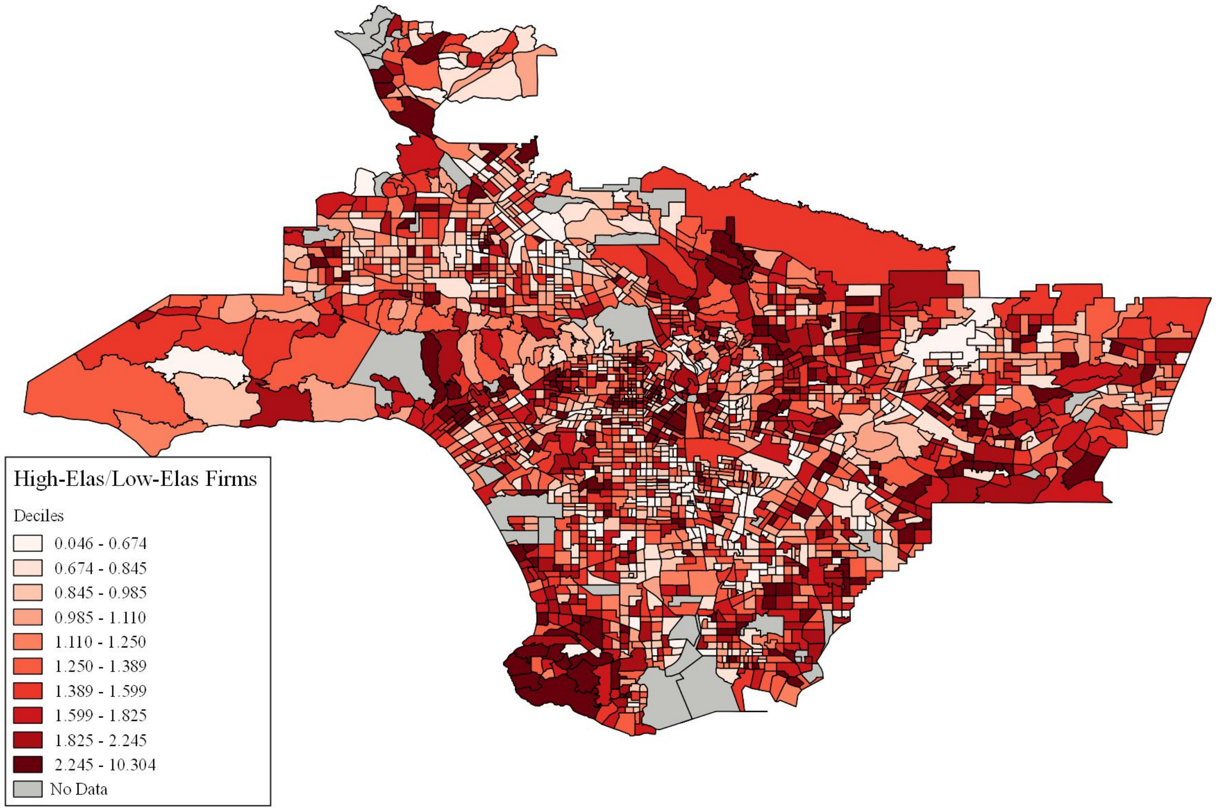

The figure plots the ratio of establishment counts in a census tract with income elasticity above median over establishment counts below median, NETS.

Figure 2.3: Number of establishments in Recreation & Education vs Liquor & Convenience Stores, 2014

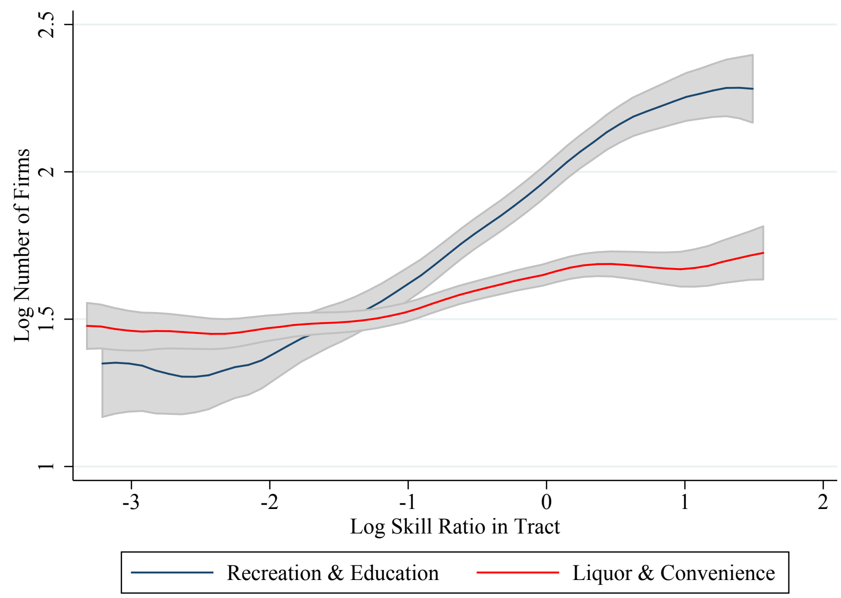

The figure plots the log number of Recreation & Education and Liquor & Convenience Stores against the local skill ratio, NETS and ACS 2014.

Figure 2.4: Number of establishments in all sectors, 2014

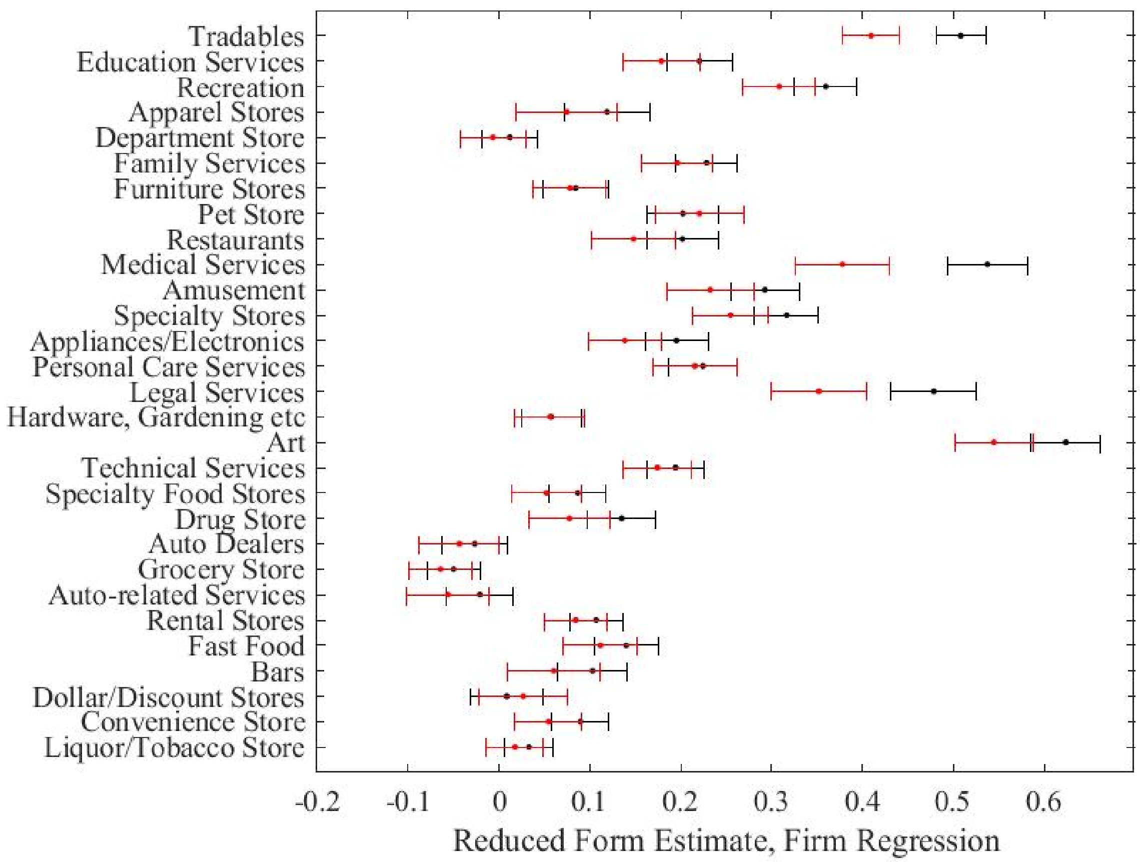

The figure plots coefficients and 95% CI from sector-level regressions of log number of firms in tract against log local skill ratio. Regressions without controls in BLACK and with controls for log population density, ratio of skilled over unskilled employment and total employment in tract in RED. Without controls: Spearman Rank Correlation (p-value): .495 (.007), with controls: Spearman Rank Correlation (p-value): .505 (.006). Data from NETS, ACS 2014 and LODES.

Figure 2.5: Graphical example of sorting patterns in model

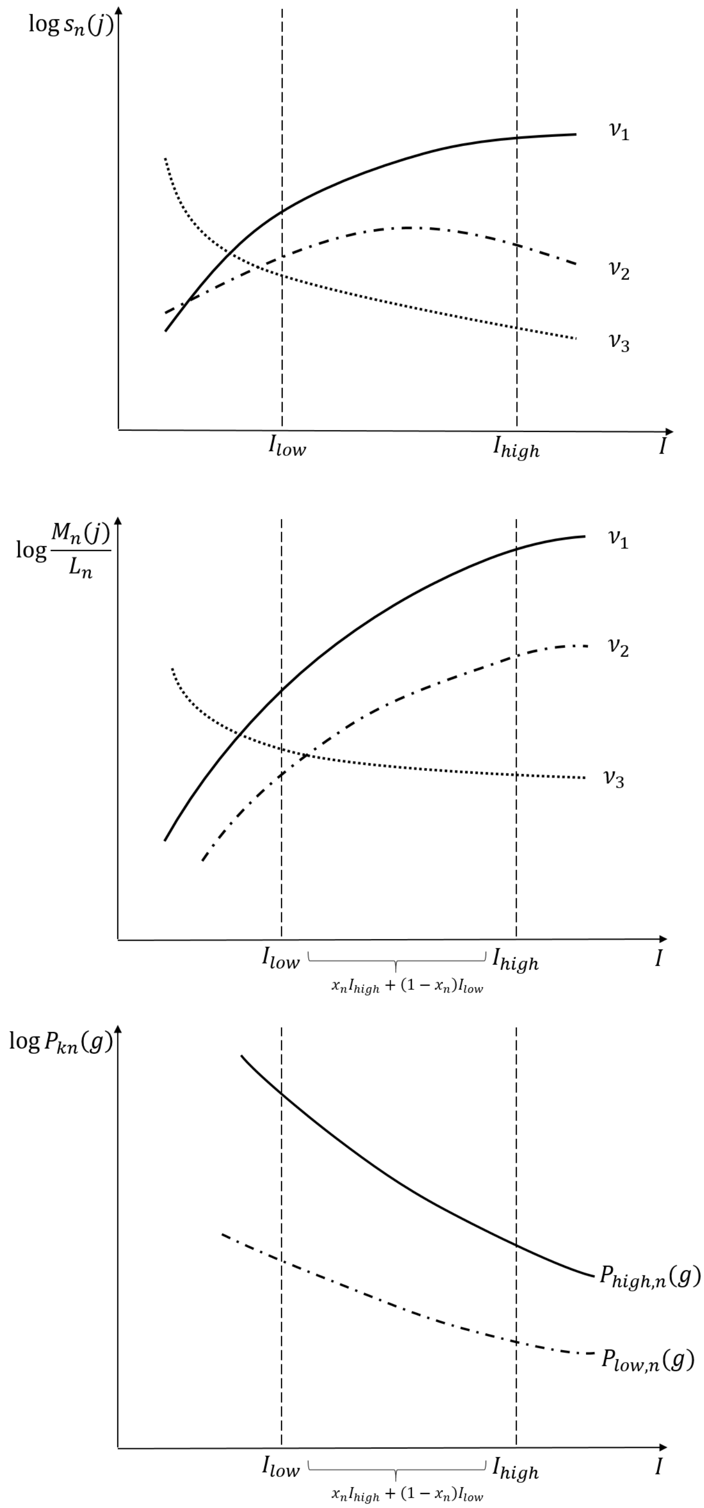

Figure 2.6: Income Elasticities by Sector

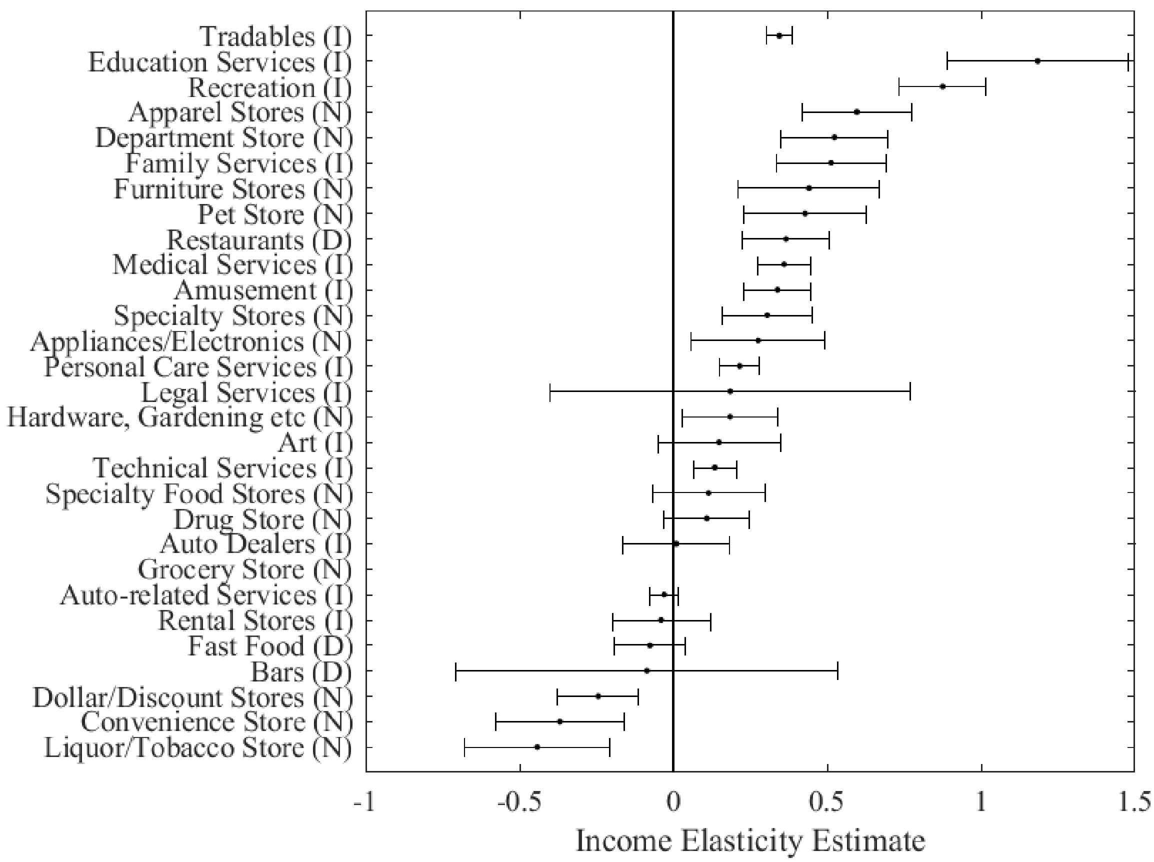

The figure plots estimated point estimates of income elasticities by goods sector using consumer expenditure data and 95% confidence intervals. Data source in parentheses (N=Nielsen, I=CEX Interview, D=CEX Diary).

Figure 2.7: Correlations of Local Sector Growth Rates in Price-Bartik

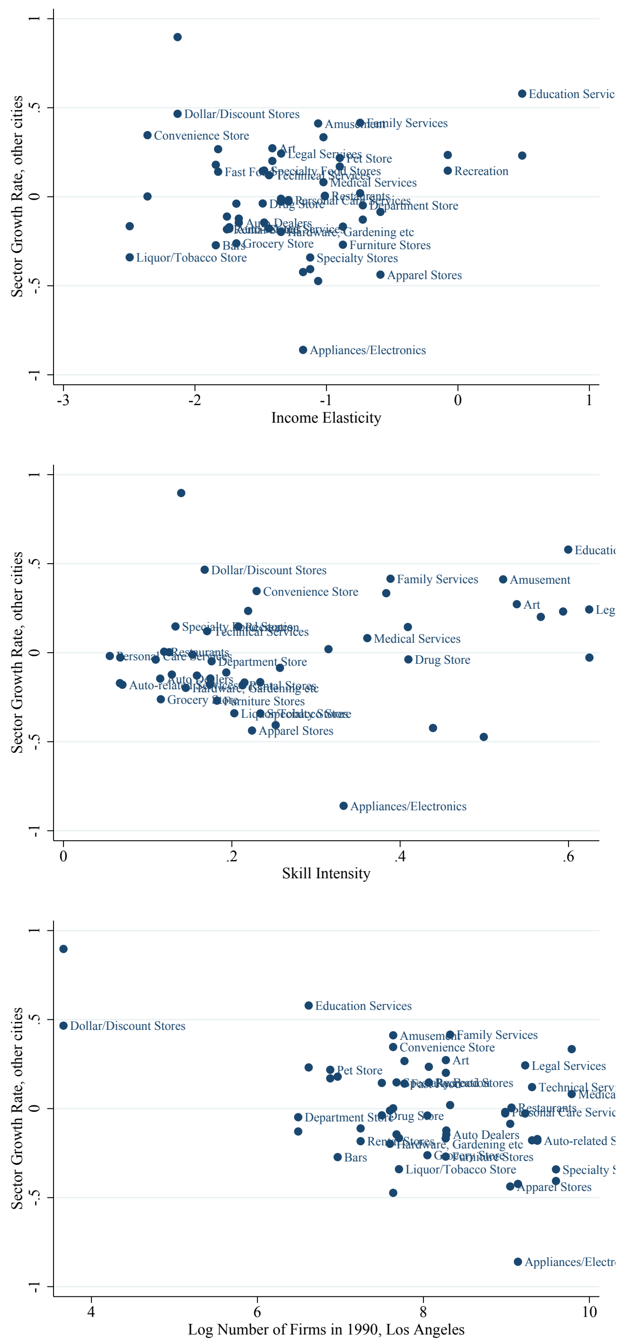

The figure plots sector growth rates in number of establishments from San Francisco Bay Area and San Diego for 1990-2000 and 2000-2014 and income elasticities (top), skill intensities by sector as skilled employment over total employment in sector nation-wide from Census 1990 and 2000 (middle) and initial log number of establishments in LA in 1990 (bottom).

Figure 2.8: Model Fit: Rents

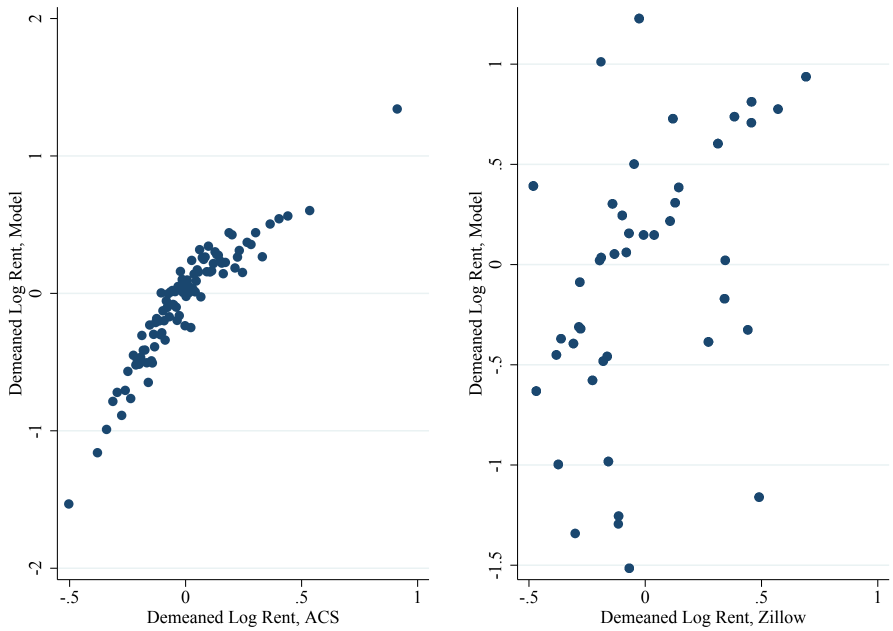

Left Panel: Log Rent by census tract in baseline model with price effects to data from the ACS 2014, binscatter with 100 bins, $Corr = .58$; Right Panel: Log Median Residential Rents from Zillow, $Corr = .43$

Figure 2.9: Model Fit: Expenditure Share on Goods

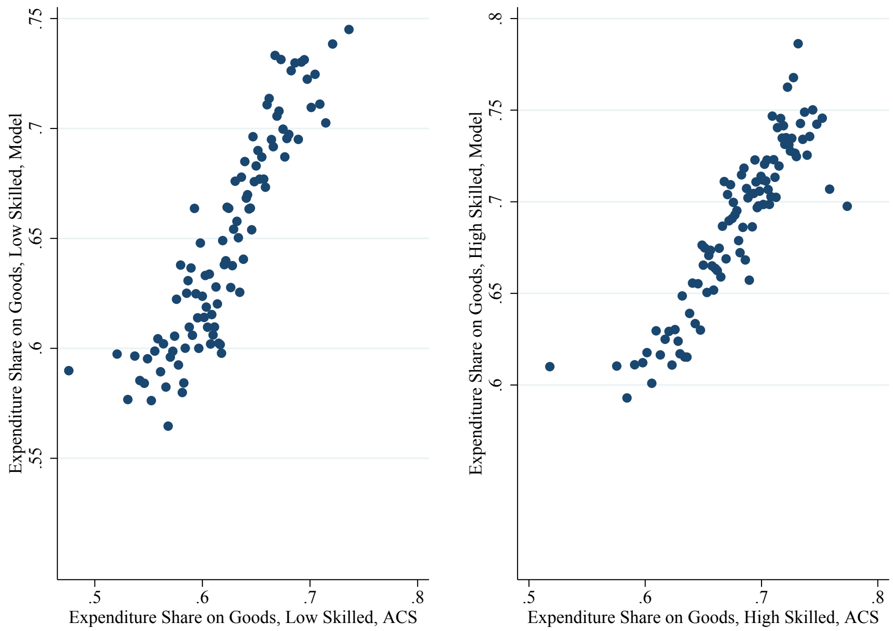

Corr = .46 in both, binscatter with 100 bins, ACS 2014 Data and Model-based.

Figure 2.10: Model Fit: Relative Expenditure Share on Goods

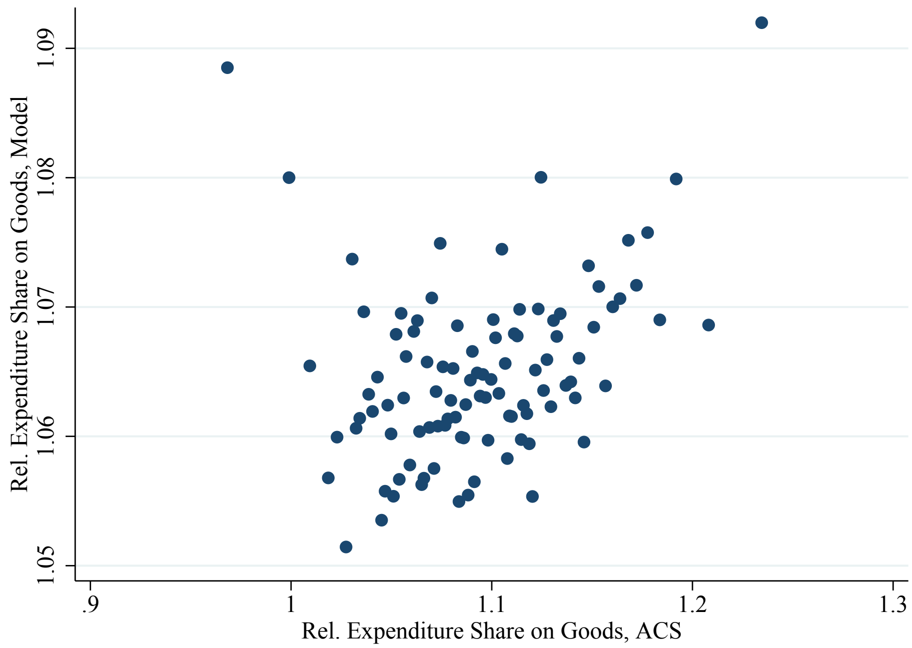

$Corr = .08$, binscatter with 100 bins, ACS 2014 Data and Model-based.

Figure 2.11: Opportunity Zones (OZ) in Los Angeles

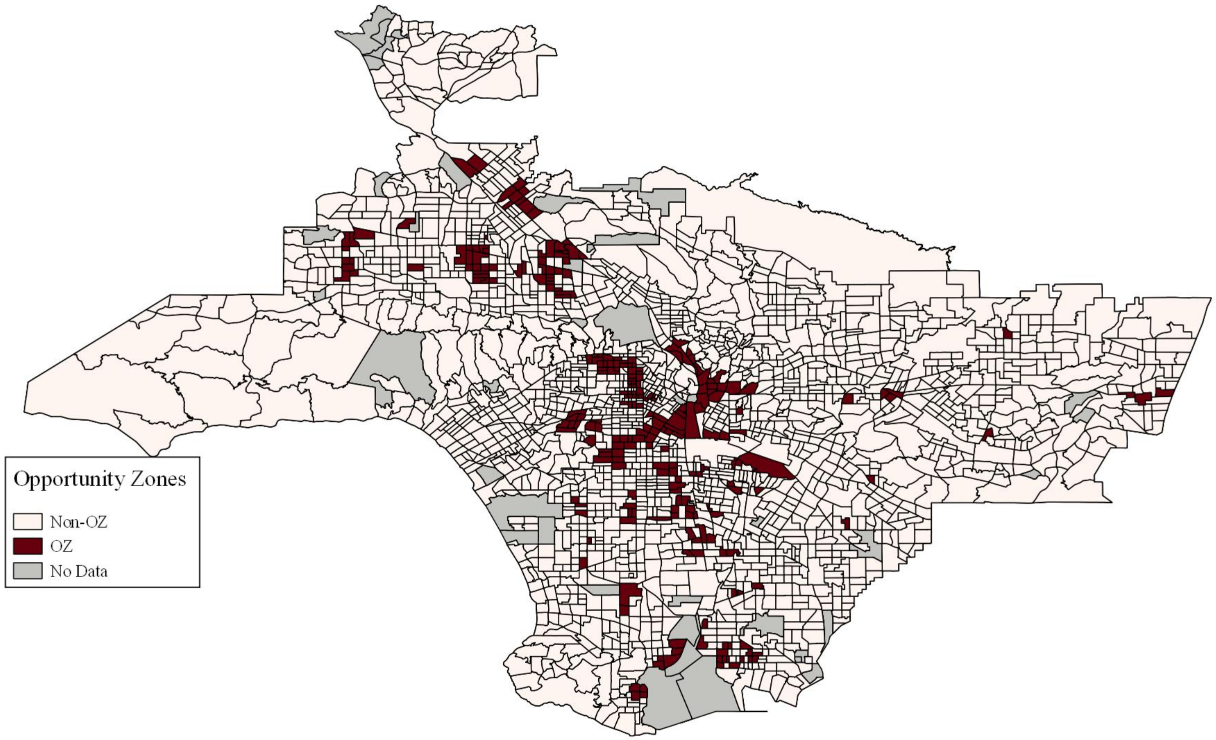

The figure plots designated Opportunity Zones (257 Census Tracts).

Figure 2.12: Social Housing in Los Angeles, 2019

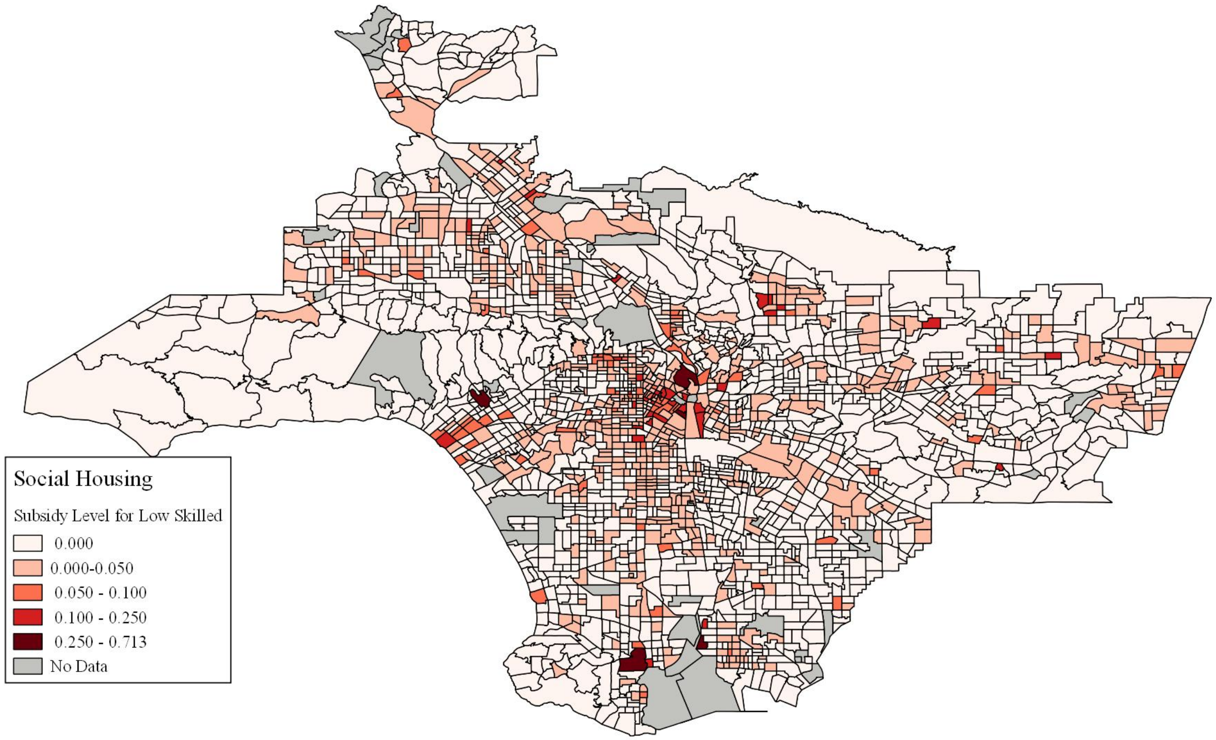

The figure plots the share of total housing costs of low skilled households that is covered by federal and state housing assistance in each Census Tract, Data from California Housing Partnership Preservation Database.

Figure 2.13: Opportunity Zones, % Change in Number of Firms, Baseline Calibration

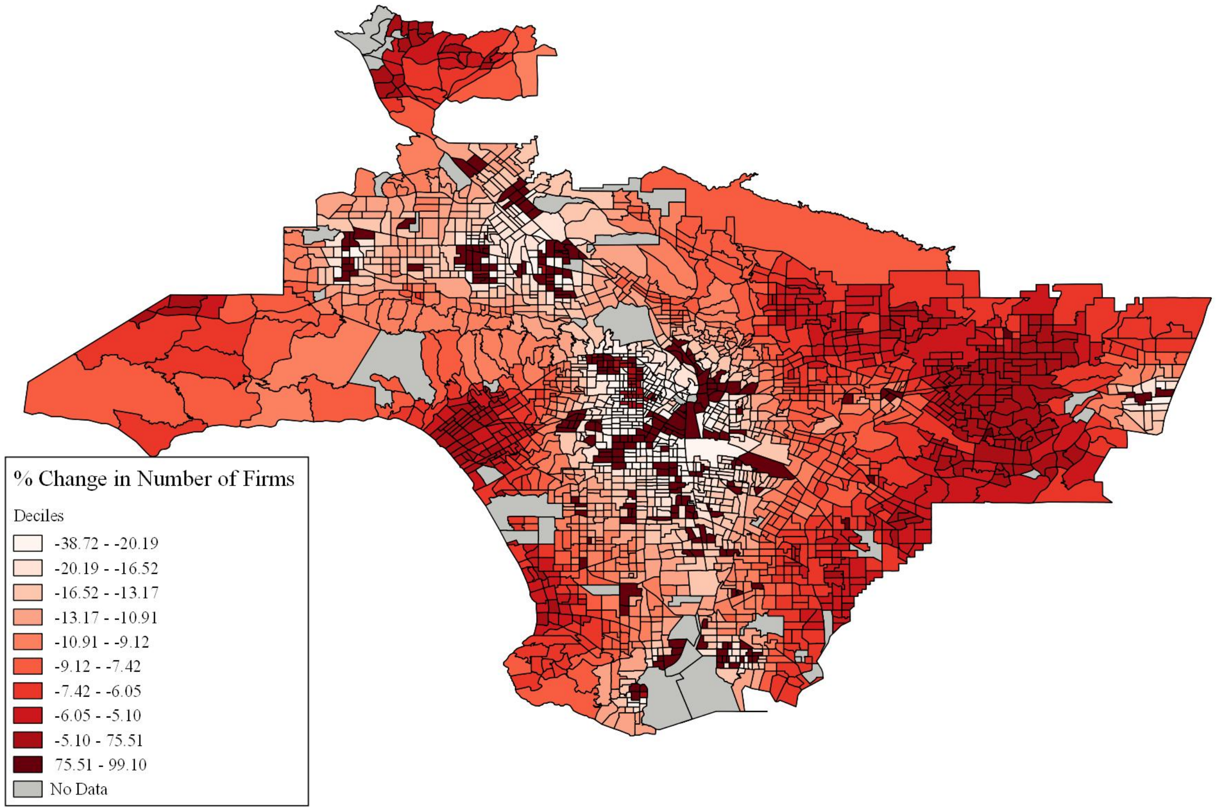

Figure 2.14: Opportunity Zones, % Change in Skill Ratio, Baseline Calibration

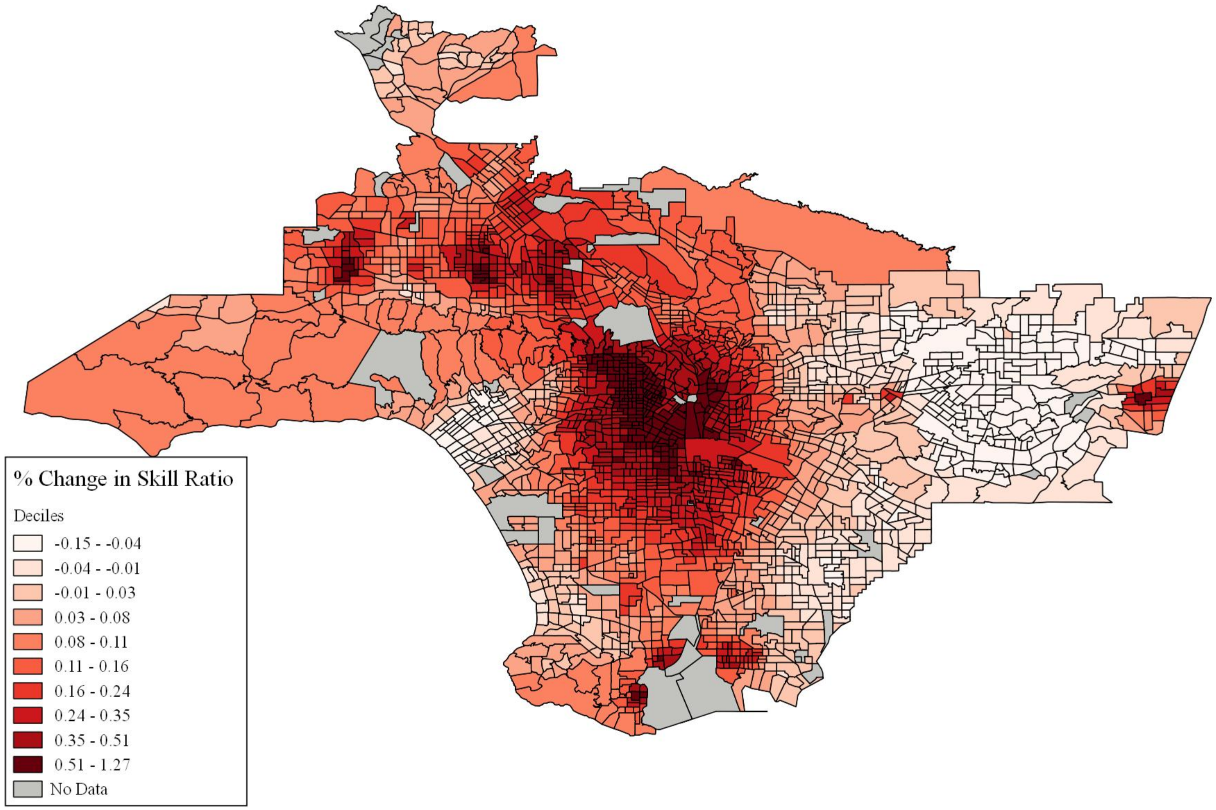

Figure 2.15: Social Housing, % Change in Number of Firms, Baseline Calibration

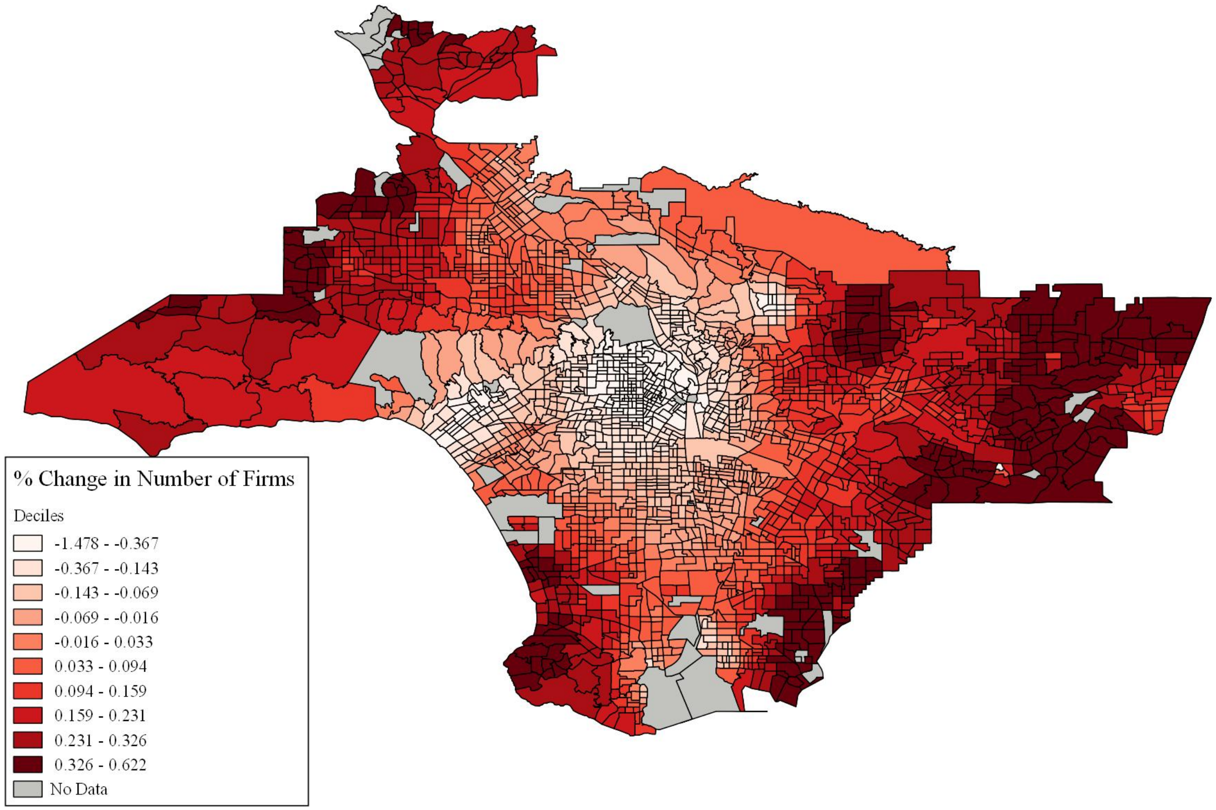

Figure 2.16: Social Housing, % Change in Skill Ratio, Baseline Calibration

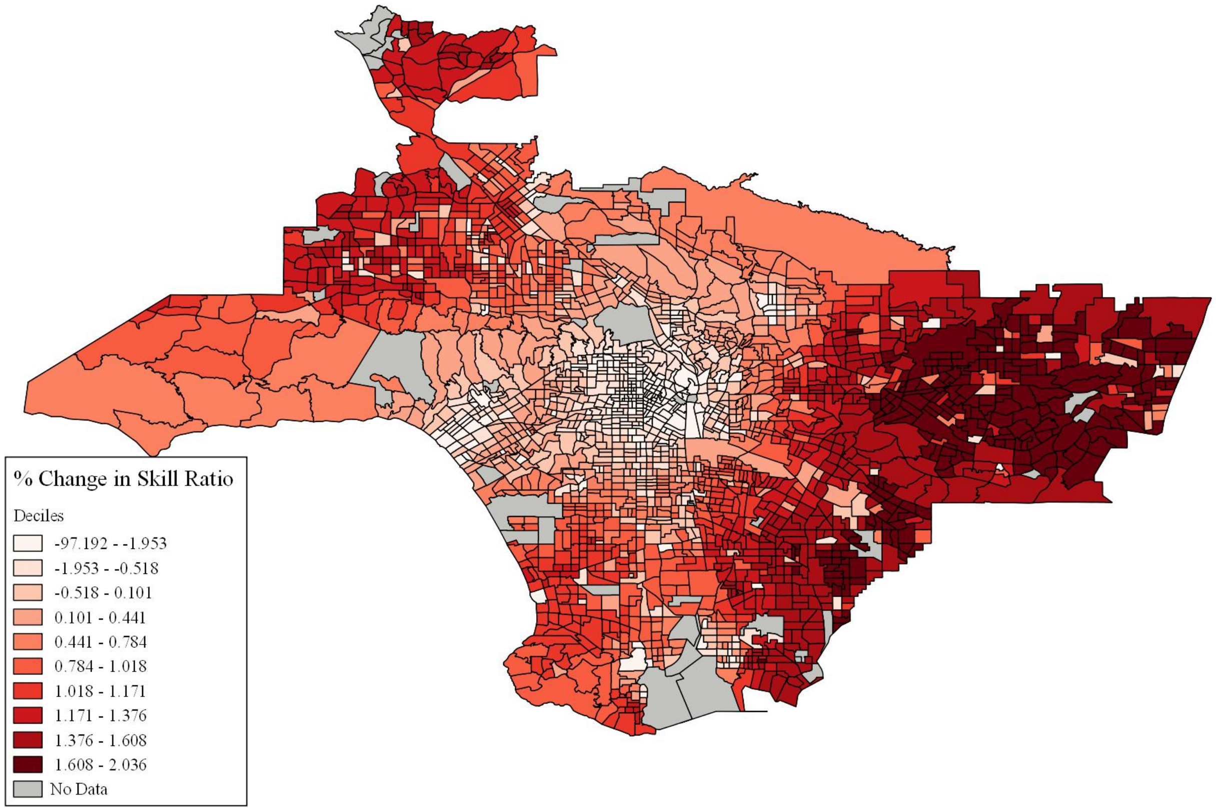

Figure 2.17: Opportunity Zones, Effect on Non-OZ Tracts, Baseline Calibration

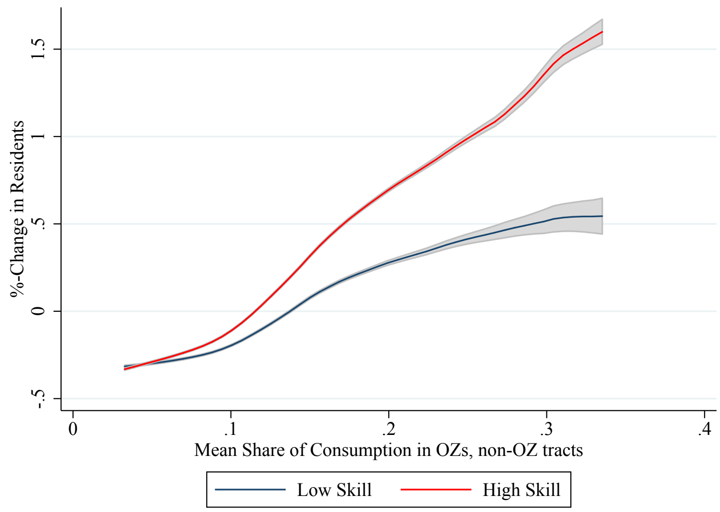

The figure plots the % change in residents as a function of the mean share of consumption across sectors from non-OZ tracts in OZ tracts.

Figure 2.18: Social Housing, Mobility of Firms by Inc-Elasticity, Baseline Calibration

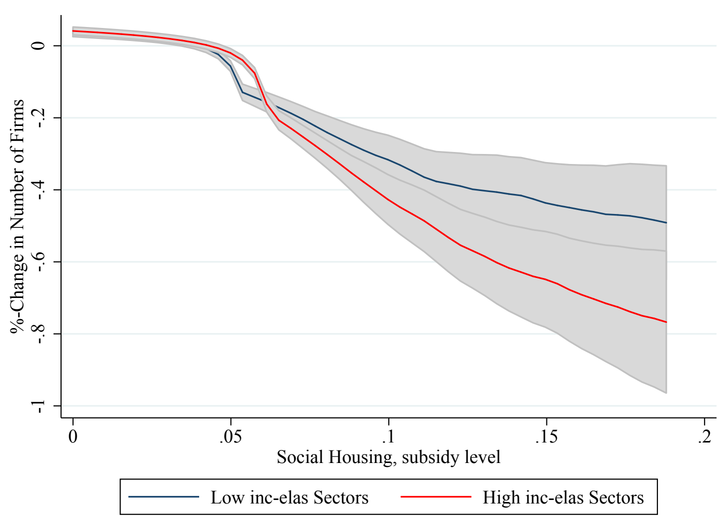

The figure plots the % change in the Number of Firms by Income Elasticity as a function fo Social Housing subsidy level.

Figure 2.19: Income Elasticities by Sector with zip code fixed effects

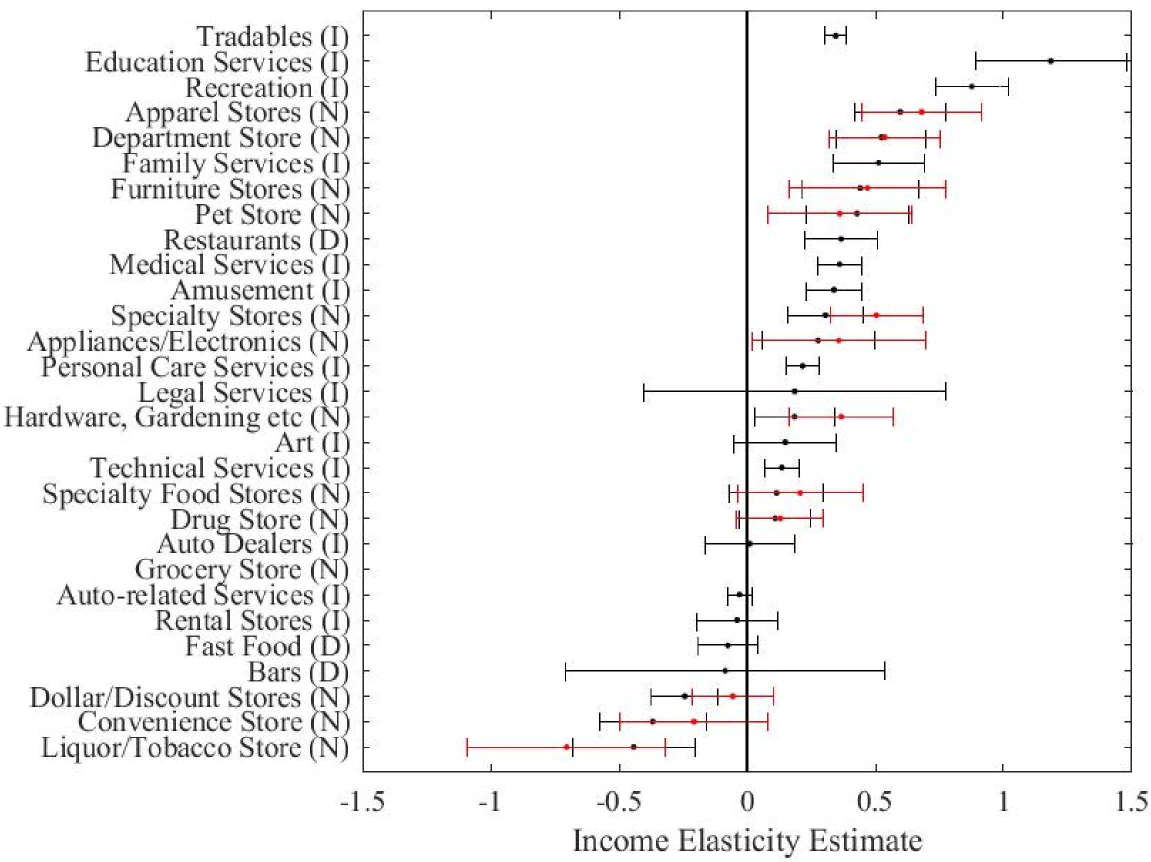

The figure plots estimated point estimates of income elasticities by goods sector using consumer expenditure data and 95% confidence intervals. Estimates from Nielsen with Zip-Code fixed effects in red. Data source in parentheses (N=Nielsen, I=CEX Interview, D=CEX Diary).

Figure 2.20: Income Elasticities by Sector with sector definition from alternative source

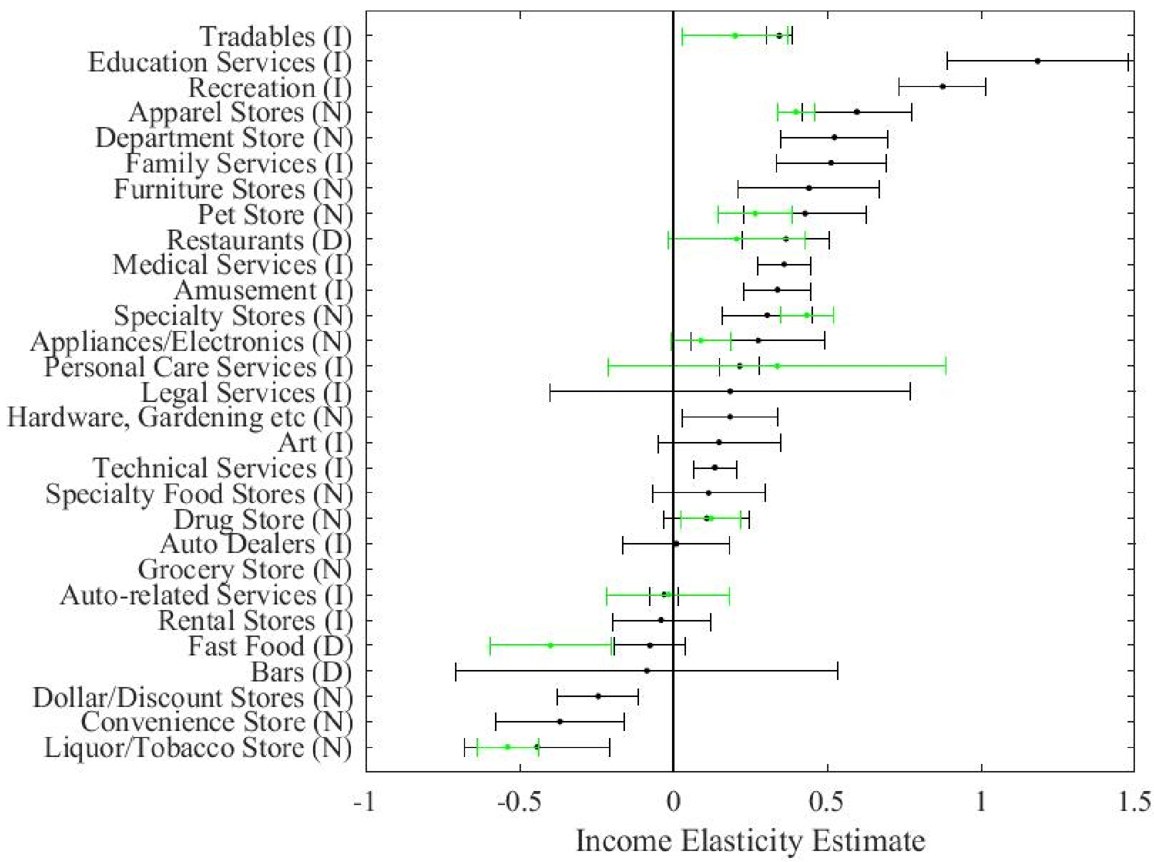

The figure plots estimated point estimates of income elasticities by goods sector using consumer expenditure data and 95% confidence intervals. Estimates from data alternative source covering approximately the same sector in green. Data source in parentheses (N=Nielsen, I=CEX Interview, D=CEX Diary).

Figure 2.21: Tract-level Changes in Relative Expenditure Share and Local Price Index Proxy

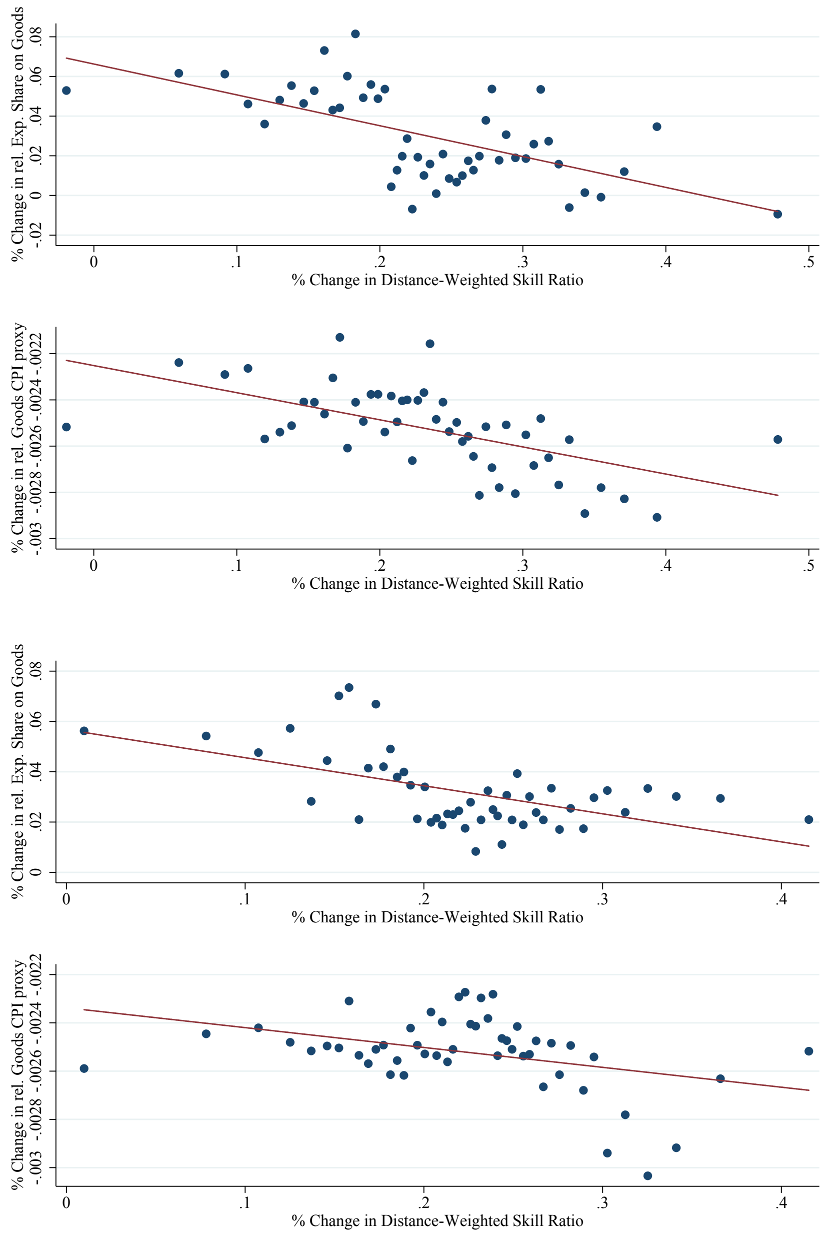

The figure plots % change in ratio of expenditures shares on goods of high over low skilled HHs (upper panel of top graph), and % change in ratio of price index proxy on goods of high over low skilled HHs (lower panel of top graph) as a function of % changes in the distance-weighted skill ratio (measure of spillovers) without controls. Bottom graph shows the same relationships with regression controls. I construct the price index proxy as in 2.30, but I use observed changes in the number of varieties in each tract instead of sector growth rates from other MSAs. I multiply the tract-sector-specific growth rates with $\frac{\theta}{(1-\sigma)}-1$, based on a first-order approximation of the model expression for the price index in 2.12. Binscatter with 50 bins, data from ACS 2014, NETS.

Figure 2.22: Opportunity Zones, % Change in Price Index for High Skilled

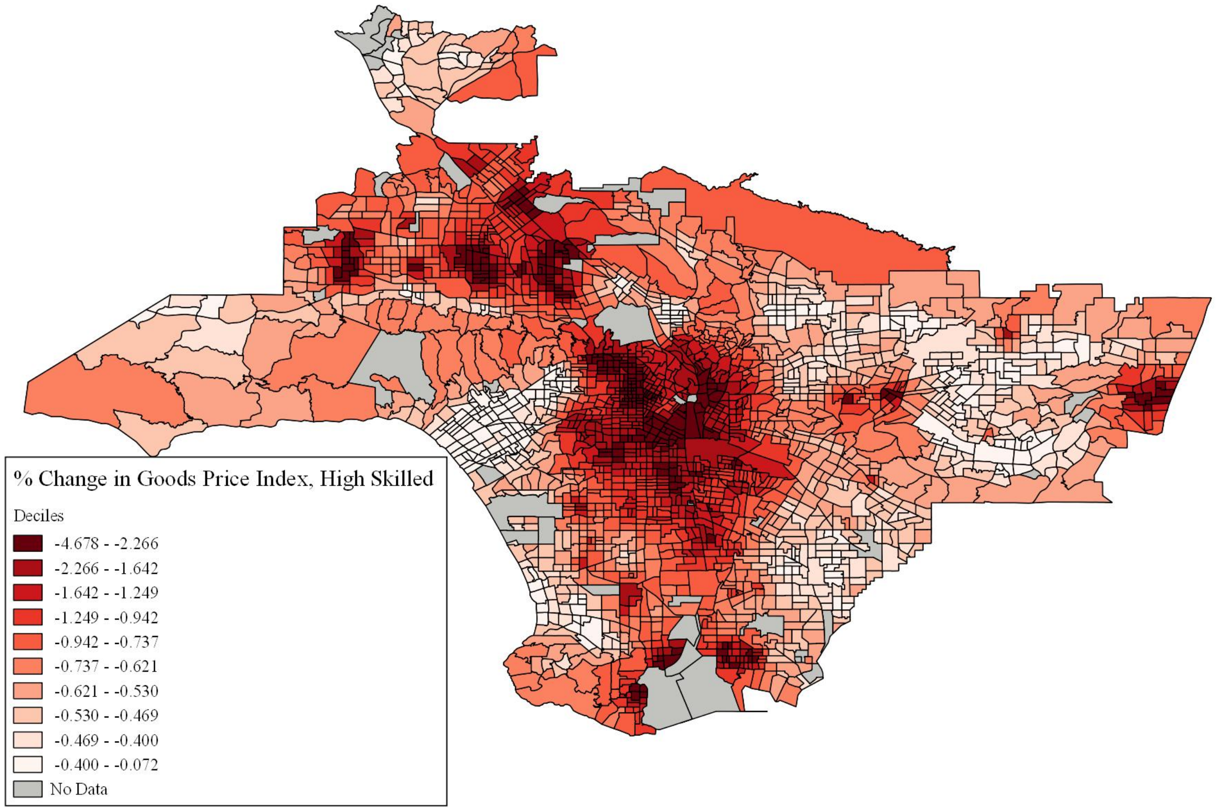

Figure 2.23: Opportunity Zones, % Change in Spillovers

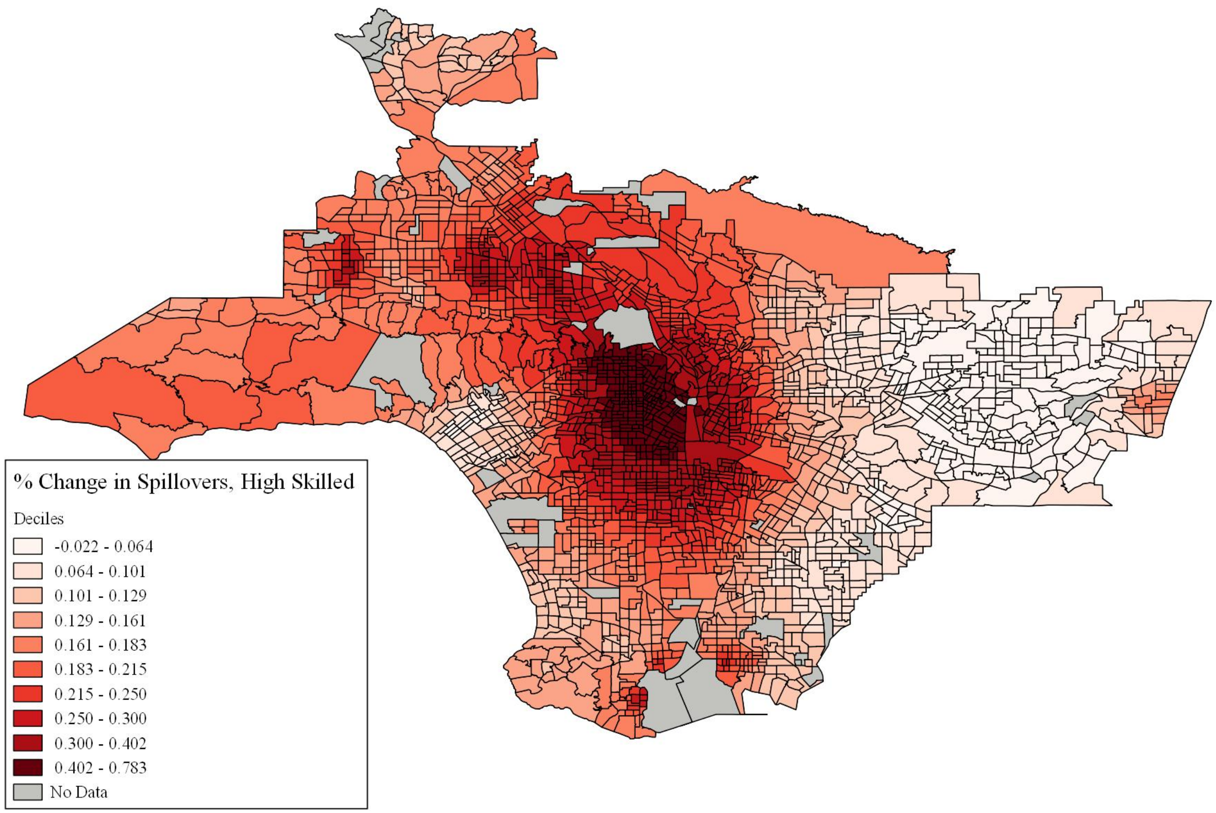

Figure 2.24: Social Housing, Relative % Change in Number of Firms by Income Elasticity

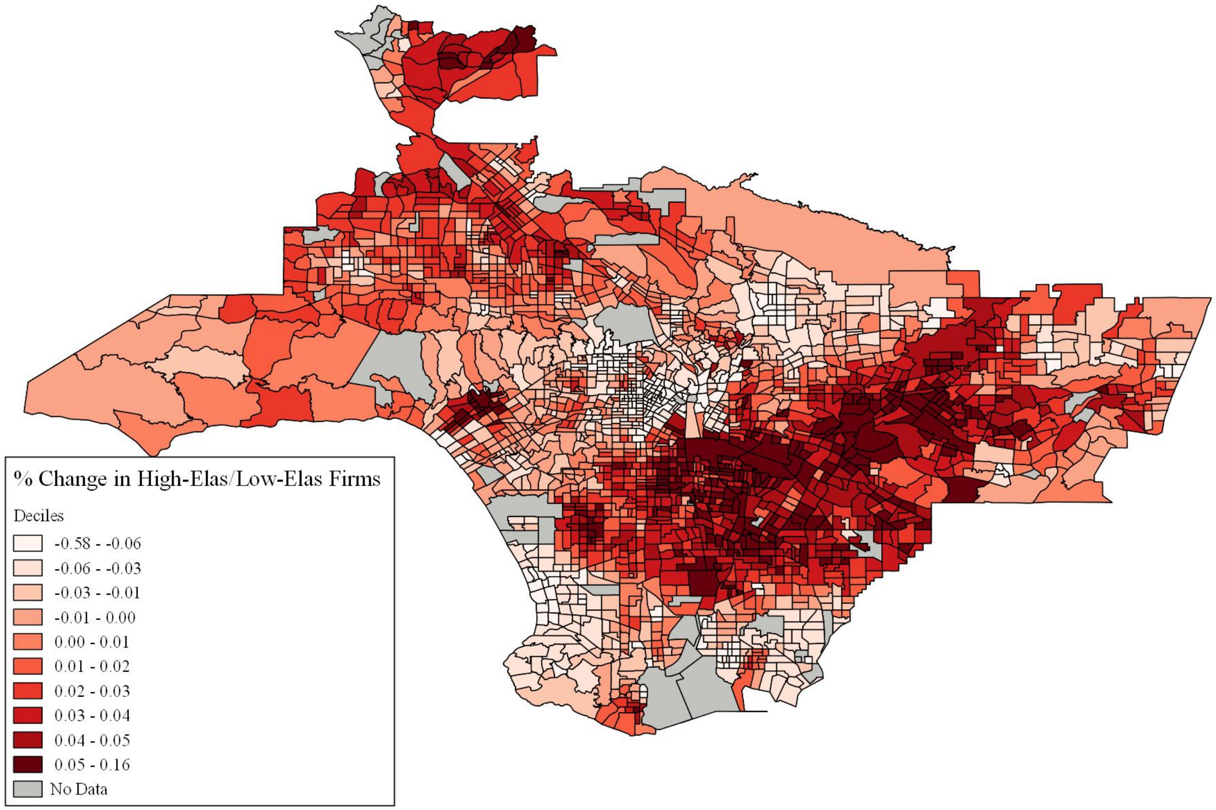

Figure 2.25: Social Housing, % Change in Skill Ratio, Calibration w/o Price Effects

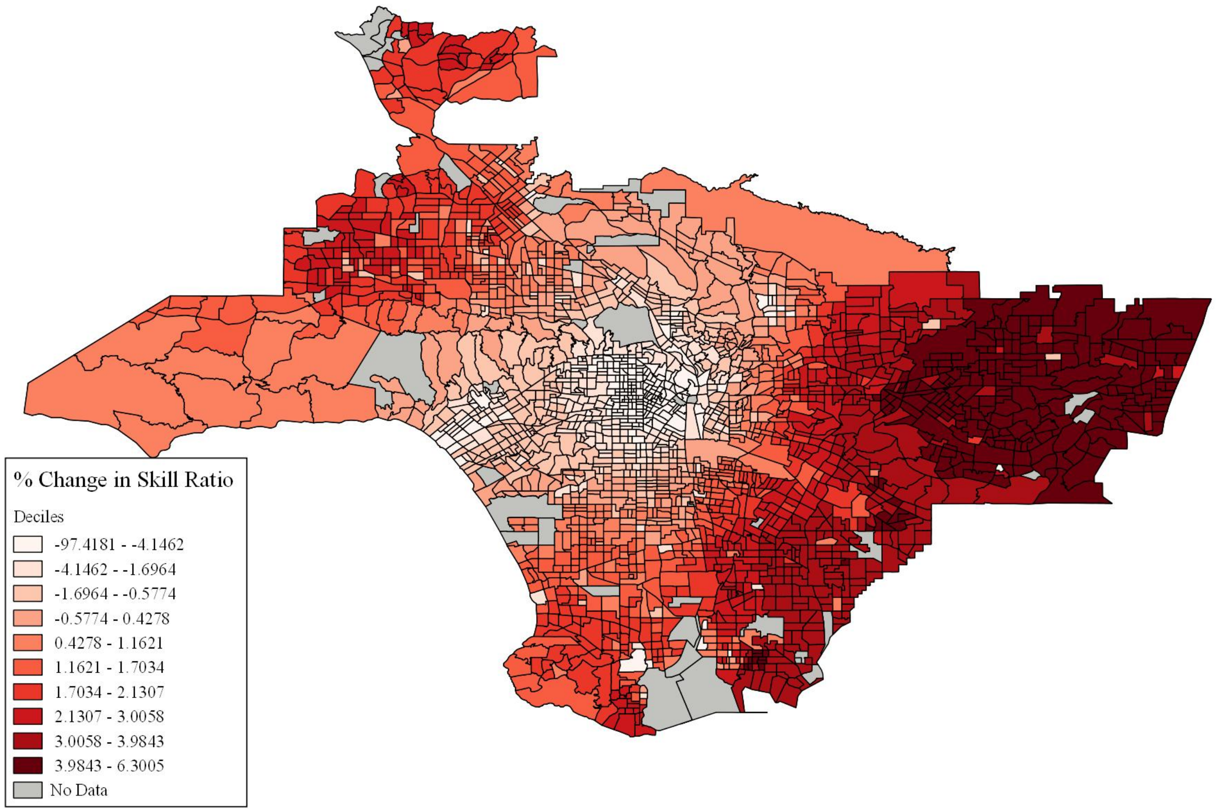

Table 2.1: Number of Establishments in Recreation & Education vs Liquor & Convenience Stores, 2014

VARIABLES	(1) Income-elastic/ Inelastic Log Ratio	(2) Recreation & Education Log Number	(3) Liquor & Convenience Log Number	(4) Recreation & Education Number> 0	(5) Liquor & Convenience Number> 0
Log Skill Ratio	0.126***	0.320***	0.081***	0.181***	0.029***
	(0.010)	(0.021)	(0.017)	(0.008)	(0.010)
Observations	2,194	1,074	1,077	2,182	2,182
R-squared	0.070	0.180	0.022	0.149	0.004

Robust standard errors. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2.2: Sector Income Elasticities and implied ν_j for Goods Sectors, 2014

	Sales Share	$\widehat{(\nu_j - \nu_{j^*})}\varepsilon$	**SE**	**Implied** ν_j
Liquor/Tobacco Store (N)	0.0062	−0.4419	0.1210	−2.4963
Convenience Store (N)	0.0153	−0.3682	0.1064	−2.3610
Dollar/Discount Stores (N)	0.0348	−0.2435	0.0672	−2.1322
Bars (D)	0.0027	−0.0851	0.3169	−1.8414
Fast Food (D)	0.0205	−0.0755	0.0587	−1.8237
Rental Stores (I)	0.0065	−0.0390	0.0813	−1.7568
Auto-related Services (I)	0.0352	−0.0293	0.0235	−1.7389
Grocery Store (N)	0.0612			−1.6852
Auto Dealers (I)	0.0866	0.0102	0.0890	−1.6666
Drug Store (N)	0.0338	0.1092	0.0712	−1.4847
Specialty Food Stores (N)	0.0110	0.1152	0.0928	−1.4737
Technical Services (I)	0.0310	0.1354	0.0349	−1.4367
Art (I)	0.0164	0.1491	0.1020	−1.4115
Hardware, Gardening etc (N)	0.0283	0.1849	0.0793	−1.3459
Legal Services (I)	0.0511	0.1856	0.2997	−1.3446
Personal Care Services (I)	0.0064	0.2162	0.0330	−1.2884
Appliances/Electronics (N)	0.0217	0.2765	0.1116	−1.1777
Specialty Stores (N)	0.0357	0.3059	0.0749	−1.1238
Amusement (I)	0.0129	0.3389	0.0561	−1.0631
Medical Services (I)	0.1452	0.3608	0.0444	−1.0229
Restaurants (D)	0.0442	0.3668	0.0716	−1.0119
Pet Store (N)	0.0046	0.4286	0.1019	−0.8985
Furniture Stores (N)	0.0142	0.4416	0.1167	−0.8747
Family Services (I)	0.0243	0.5128	0.0915	−0.7439
Department Store (N)	0.0170	0.5240	0.0891	−0.7234
Apparel Stores (N)	0.0363	0.5970	0.0907	−0.5893
Recreation (I)	0.0252	0.8760	0.0723	−0.0772
Education Services (I)	0.0044	1.1844	0.1503	0.4889
Tradables (I)	0.1675	0.3451	0.0218	−1.0517

Nominal income instrumented with dummy for high skill. All regressions include dummies for household size, age of householder and number of earners interacted with sector fixed effects, as well as Sector-MSA-Time fixed effects. All regressions are weighted by household weights in respective expenditure survey (N=Nielsen, I=CEX Interview, D=CEX Diary). Standard error are clustered at Household level.

Table 2.3: Correlations of Local Sector Growth Rates in Price-Bartik

VARIABLES	(1) Sector Growth Rate other cities	(2) Sector Growth Rate other cities	(3) Sector Growth Rate other cities
Income Elasticity	0.088		
	(0.075)		
Log Number of Firms, 1990		-0.113***	
		(0.029)	
Skill Intensity			0.369
			(0.248)
Observations	56	56	49
R-squared	0.050	0.232	0.068

Robust standard errors. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2.4: Estimation of κ and δ_k, $\psi = -1.5$, 5km buffer

Panel A: First Stage Regression, $\frac{s_{kn,t-1}(g)}{s_{kn,t-1}(h)} \log \hat{s}_{kn,t}(g)$

	(1)	(2)
	1st	1st
VARIABLES	Price Effects	Price Effects
Avg Price IV	-0.037**	-0.075***
	(0.017)	(0.015)
Rel Price IV	0.502***	0.460***
	(0.063)	(0.070)
Rel Price IV X High	-0.136***	-0.122***
	(0.032)	(0.032)
Observations	8,362	8,362
R-squared	0.221	0.372
Controls	no	yes

Panel B: First Stage Regression, $\hat{\mathscr{L}}_{kn,t}$

	(1)	(2)	(3)	(4)
	1st	1st	1st	1st
VARIABLES	No Price Effects	Local	Price Effects	Price Effects
Avg Price IV	0.119***	0.116***	0.094***	0.116***
	(0.011)	(0.011)	(0.016)	(0.011)
Rel Price IV	0.306***	0.310***	0.284***	0.310***
	(0.061)	(0.061)	(0.063)	(0.061)
Rel Price IV X High	0.008***	0.002	0.000	0.002
	(0.002)	(0.001)	(.)	(0.001)
Observations	8,362	8,362	8,362	8,362
R-squared	0.867	0.867	0.827	0.867
Controls	yes	yes	no	yes

Panel C: First Stage Regression, $\hat{\mathscr{L}}_{kn,t} \times High$

VARIABLES	(1) 1st No Price Effects	(2) 1st Local	(3) 1st Price Effects	(4) 1st Price Effects
Avg Price IV	0.060***	0.058***	0.047***	0.058***
	(0.006)	(0.006)	(0.008)	(0.006)
Rel Price IV	-0.244***	-0.244***	-0.262***	-0.244***
	(0.029)	(0.029)	(0.027)	(0.029)
Rel Price IV X High	0.811***	0.807***	0.807***	0.807***
	(0.031)	(0.031)	(0.031)	(0.031)
Observations	8,362	8,362	8,362	8,362
R-squared	0.906	0.906	0.894	0.906
Controls	yes	yes	no	yes

Panel D: Second Stage Regressions

VARIABLES	(1) OLS	(2) IV No Price Effects	(3) IV Local	(4) IV Price Effects	(5) IV Price Effects
$-\frac{\hat{\kappa}}{1-\eta}$	0.020			-5.717***	-4.695***
	(0.048)			(1.007)	(0.768)
$\hat{\delta}_{low}$	-0.514***	0.385	0.260	-0.920	-0.883
	(0.065)	(0.511)	(0.517)	(1.014)	(0.736)
$\hat{\delta}_{high} - \hat{\delta}_{low}$	2.675***	1.728***	1.714***	0.769**	1.019***
	(0.052)	(0.090)	(0.088)	(0.306)	(0.247)
Observations	8,362	8,362	8,362	8,362	8,362
Controls	yes	yes	yes	no	yes
1st Stage F-Stat		46	44.03	8.80	19.19

All specifications include the sum of shares interacted with year-dummies and skill-year FX. Controls include changes in household income, changes in residential rents, log distance to city center (City of Los Angeles City Hall), log population density in 1990 and log average slope. Standard errors clustered at level of Census Tract. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2.5: Estimation of θ

Panel A: First Stage Regression

VARIABLES	(1) 1st	(2) 1st	(3) 1st	(4) 1st alt. IV
Log Avg Slope X v_j	0.116***	0.061***	0.185***	
	(0.010)	(0.013)	(0.015)	
Log Avg Slope X rank(v_j)				0.008***
				(0.001)
Observations	152,323	79,550	58,783	152,323
Sample	Chains	Retail	Services	Chains
Tract-Year FX	yes	yes	yes	yes
Sector-Year FX	yes	yes	yes	yes
Chain-Year FX	yes	yes	yes	yes
Number of clusters	6113	5285	5601	6113

Panel B: Second Stage Regressions

VARIABLES	(1) IV	(2) IV	(3) IV	(4) alt. IV
$\frac{\sigma-1}{\hat{\theta}}$	0.251***	0.225	0.317***	0.164***
	(0.066)	(0.146)	(0.098)	(0.062)
Observations	152,323	79,550	58,783	152,323
Sample	Chains	Retail	Services	Chains
Tract-Year FX	yes	yes	yes	yes
Sector-Year FX	yes	yes	yes	yes
Chain-Year FX	yes	yes	yes	yes
Number of clusters	6113	5285	5601	6113
1st Stage F-Stat	144.2	20.96	162.3	154.7

Standard errors clustered at level of Year-Zipcode. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2.6: Main Calibration of Model

Parameter	Description	Value	Source
Preferences			
κ	Resident Supply Elasticity	2.4	Estimated
ε_h	Income Elasticity of Housing	.507	Assumed
ε_g	Income Elasticity of Goods	.507	Assumed
ν_j	Sector Income Elasticities	see table 2.2	Estimated
η	EoS between Housing and Goods	.493	Assumed
γ	EoS across Sectors	2	Assumed
σ	Eos within Sector	5	Assumed
Firm Supply			
θ	Firm Supply Elasticity	16	Estimated
Shopping Frictions			
ϕ^1	Distance Elasticity, w/ Price Effects	-1.5	Assumed
ϕ^2	Distance Elasticity, w/o Price Effects	0	Assumed
Spillover Elasticities			
δ^1_{low}	Low Skilled, w/ Price Effects	0	Estimated
δ^2_{low}	Low Skilled, w/o Price Effects	0	Estimated
δ^1_{high}	High Skilled, w/ Price Effects	1	Estimated
δ^2_{high}	High Skilled, w/o Price Effects	1.25	Estimated
Skill Premium			
ρ	Rel. Labor Endowment of High Skilled	1.7	Estimated

Table 2.7: Opportunity Zones, Baseline Calibration, % Changes

	Firms All local	HHs All	HHs Low Skill	HHs High Skill	HHs Skill Ratio
OZ	79.98	0.73	0.59	1.18	0.59
Non-OZ	-10.39	-0.09	-0.09	-0.09	0.09
R-squared	0.95	0.44	0.46	0.43	0.47

Table 2.8: Opportunity Zones, Calibration without Price Effects, % Changes

	Firms All local	HHs All	HHs Low Skill	HHs High Skill	HHs Skill Ratio
OZ	90.86	0.02	0.02	0.03	0.00
Non-OZ	-12.93	-0.00	-0.00	-0.00	0.01
R-squared	0.99	0.07	0.08	0.05	0.71

Table 2.9: Opportunity Zones, % Changes in local Skill Ratio and Welfare by Skill

	w/ Price Effects	w/o Price Effects	w/ Price Effects No Spillover	w/o Price Effects No Spillover	Homothetic Preferences	Homothetic Preferences No Spillover
Skill Ratio						
OZ	0.59	0.00	0.39	0.01	0.21	0.21
Non-OZ	0.09	0.01	0.12	0.01	0.21	0.21
Welfare						
Low Skill	-0.09	-0.14	-0.09	-0.14	-0.19	-0.19
High Skill	-0.08	-0.16	-0.16	-0.16	-0.19	-0.28

Table 2.10: Social Housing, Baseline Calibration, % Changes

	Firms All local	HHs All	HHs Low Skill	HHs High Skill	HHs Skill Ratio
Subsidy	-4.54	2.36	29.73	-50.77	-107.66
Constant	0.05	-0.02	-0.29	0.45	0.80
R-squared	0.23	0.09	0.78	0.61	0.85

Table 2.11: Social Housing, Calibration without Price Effects, % Changes

	Firms All local	HHs All	HHs Low Skill	HHs High Skill	HHs Skill Ratio
Subsidy	0.01	-0.10	32.64	-63.86	-115.84
Constant	0.00	0.00	-0.32	0.59	0.71
R-squared	0.00	0.00	0.63	0.34	0.64

Table 2.12: Social Housing, % Changes in local Skill Ratio and Welfare by Skill

	w/ Price Effects	w/o Price Effects	w/ Price Effects No Spillover	w/o Price Effects No Spillover	Homothetic Preferences	Homothetic Preferences No Spillover
Skill Ratio						
Subsidy	-107.66	-115.84	-95.04	-92.92	-162.47	-143.81
Constant	0.80	0.71	0.76	0.75	1.25	1.19
Welfare						
Low Skill	0.10	0.08	0.10	0.10	0.16	0.17
High Skill	-0.23	-0.35	-0.11	-0.10	-0.38	-0.18

Table 2.13: Skill Premium, ACS 2014

	(1) All US	(2) LA Sample
Skilled HH	0.639***	0.705***
	(0.001)	(0.006)
Obs	4,242,708	137,063
R-squared	0.284	0.254

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 2.14: Expenditure Share on Housing by Skill

	(1) All US ACS 2014	(2) LA Sample ACS 2014	(3) All US CEX 12-16	(4) LA Sample CEX 12-16	(5) LA Tracts ACS 2014	(6) LA Tracts Model
Skilled HH	-0.0449***	-0.0535***	-0.0168***	-0.0072*	-0.0578***	-0.0398***
	(0.0002)	(0.0015)	(0.0015)	(0.0043)	(0.0006)	(0.0002)
Obs	4,078,372	127,523	40,868	5,578	4,388	4,388
R-squared	0.1257	0.0700	0.1438	0.1444	0.9573	0.9984

(1) and (2) include Puma-year FX and dummies for sex and age of HH head, HH size and home ownership. (3) and (4) include MSA-year FX and same dummies as above. (5) and (6) use tract level data and model outcomes and include tract FX. Robust Standard errors. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2.15: Sector Income Elasticities, Robustness

	Baseline		Zipcode Fx		Alternative Source	
	$\widehat{(v_j - v_{j^*})\varepsilon}$	SE	$\widehat{(v_j - v_{j^*})\varepsilon}$	SE	$\widehat{(v_j - v_{j^*})\varepsilon}$	SE
Liquor/Tobacco Store (N)	−0.4419	0.1210	−0.7041	0.1971	−0.5380	0.0508
Convenience Store (N)	−0.3682	0.1064	−0.2072	0.1473		
Dollar/Discount Stores (N)	−0.2435	0.0672	−0.0558	0.0814		
Bars (D)	−0.0851	0.3169				
Fast Food (D)	−0.0755	0.0587			−0.3986	0.1009
Rental Stores (I)	−0.0390	0.0813				
Auto-related Services (I)	−0.0293	0.0235			−0.0149	0.1013
Grocery Store (N)						
Auto Dealers (I)	0.0102	0.0890				
Drug Store (N)	0.1092	0.0712	0.1292	0.0865	0.1239	0.0493
Specialty Food Stores (N)	0.1152	0.0928	0.2069	0.1252		
Technical Services (I)	0.1354	0.0349				
Art (I)	0.1491	0.1020				
Hardware, Gardening etc (N)	0.1849	0.0793	0.3675	0.1027		
Legal Services (I)	0.1856	0.2997				
Personal Care Services (I)	0.2162	0.0330			0.3382	0.2796
Appliances/Electronics (N)	0.2765	0.1116	0.3574	0.1733	0.0899	0.0493
Specialty Stores (N)	0.3059	0.0749	0.5045	0.0926	0.4340	0.0434
Amusement (I)	0.3389	0.0561				
Medical Services (I)	0.3608	0.0444				
Restaurants (D)	0.3668	0.0716			0.2061	0.1135
Pet Store (N)	0.4286	0.1019	0.3610	0.1429	0.2665	0.0616
Furniture Stores (N)	0.4416	0.1167	0.4692	0.1549		
Family Services (I)	0.5128	0.0915				
Department Store (N)	0.5240	0.0891	0.5369	0.1102		
Apparel Stores (N)	0.5970	0.0907	0.6805	0.1190	0.3991	0.0306
Recreation (I)	0.8760	0.0723				
Education Services (I)	1.1844	0.1503				
Tradables (I)	0.3451	0.0218			0.2013	0.0877

As in Table 2.2. Zipcode Fx refers to specification in Nielsen where the Sector-MSA-Time Fx is replaced by Sector-Zipcode-Time Fx. Alternative Source refers to estimates from other samples covering the approximately same sector.

Table 2.16: Estimation of κ and δ_k, Robustness

VARIABLES	(1) 10km Buffer Local	(2) 10km Buffer Price Effects	(3) $\psi = 3$ Local	(4) $\psi = 3$ Price Effects	(5) Pop-based Local	(6) Pop-based Price Effects	(7) Weighted ^{IV} Local	(8) Weighted ^{IV} Price Effects	(9) Rent Share Local	(10) Rent Share Price Effects
$-\frac{\hat{\kappa}}{1-\eta}$		-4.596***		-4.853***		-4.508***		-4.628***		-7.713***
		(0.920)		(0.848)		(0.824)		(0.725)		(2.948)
$\hat{\delta}_{low}$	-1.042***	-1.273***	0.581***	-0.492	1.101***	-0.538	-0.146	-0.927*	0.260	-1.752
	(0.189)	(0.291)	(0.220)	(0.387)	(0.308)	(0.564)	(0.272)	(0.485)	(0.517)	(1.491)
$\hat{\delta}_{high} - \hat{\delta}_{low}$	1.605***	0.693***	1.001***	0.589***	1.763***	1.007***	1.715***	1.029***	1.714***	1.848***
	(0.074)	(0.266)	(0.029)	(0.147)	(0.171)	(0.321)	(0.088)	(0.250)	(0.088)	(0.390)
Observations	8,362	8,362	8,362	8,362	8,358	8,358	8,362	8,362	8,362	8,362
Controls	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
1st Stage F-Stat	210.6	14.28	17.39	11.31	64.69	13.86	67.35	18.48	44.03	2.442

All specifications include the sum of shares interacted with year-dummies (Borusyak *et al.* (2018)) and skill-year FX. Controls include log distance to city center (City of Los Angeles City Hall), log population density in 1990 and log average slope. Standard errors clustered at level of Census Tract.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2.17: Opportunity Zones, Summary Statistics

	OZ	Non-OZ	Difference
Log Number of HHs	6.827	6.907	-0.0795*
Log HH Income	10.74	11.22	-0.484***
Log Skill Ratio	-1.423	-0.692	-0.731***
Log Number of Firms	4.264	4.380	-0.116*
Log Number of Firms, low Inc-Elas	3.491	3.539	-0.0471
Log Number of Firms, high Inc-Elas	3.571	3.765	-0.195**
Log Rent per sqft	0.341	0.330	0.0111
Observations	257	1937	

Table 2.18: Opportunity Zones, Decomposing the HH Response, Baseline Model

	U Low Skill	U High Skill	Rent All	P(g) Low Skill	P(g) High Skill	CPI(g) Low Skill	CPI(g) High Skill	B High Skill
OZ	0.16	0.25	0.61	-2.75	-2.85	-2.94	-3.15	0.38
Non-OZ	-0.13	-0.19	-1.47	-0.90	-0.72	-0.73	-0.48	0.18
R2	0.61	0.74	0.79	0.81	0.80	0.74	0.69	0.74

Table 2.19: Opportunity Zones, Decomposing the HH Response, Model w/o Price Effects

	U Low Skill	U High Skill	Rent All	P(g) Low Skill	P(g) High Skill	CPI(g) Low Skill	CPI(g) High Skill	B High Skill
OZ	-0.13	-0.15	-1.43	-1.02	-1.02	-0.85	-0.83	0.01
Non-OZ	-0.14	-0.16	-1.49	-1.03	-1.03	-0.84	-0.82	0.01
R2	0.99	0.99	1.00	1.00	1.00	1.00	1.00	0.84

Table 2.20: Opportunity Zones, Decomposing the Firm Response

	$M_n(j)$ w/ Price Effects	$M_n(j)$ w/ Price Effects	$M_n(j)$ w/o Price Effects	$M_n(j)$ w/o Price Effects
OZ	91.30		104.95	
Non-OZ X ν_j		1.06		0.00
R2	0.96	1.00	1.00	1.00
Sector FX	yes	yes	yes	yes
Tract FX	no	yes	no	yes

Table 2.21: Social Housing, Decomposing the HH Response, Baseline Model

	U Low Skill	U High Skill	Rent All	P(g) Low Skill	P(g) High Skill	CPI(g) Low Skill	CPI(g) High Skill	B High Skill
Subsidy	12.39	-12.81	62.43	16.62	-15.88	0.30	0.35	-20.03
Constant	-0.02	0.02	-0.04	-0.04	0.02	-0.00	-0.00	-0.14
R2	0.78	0.59	0.78	0.77	0.55	0.11	0.10	0.20

Table 2.22: Social Housing, Decomposing the HH Response, Model w/o Price Effects

	U Low Skill	U High Skill	Rent All	P(g) Low Skill	P(g) High Skill	CPI(g) Low Skill	CPI(g) High Skill	B High Skill
Subsidy	13.60	-10.78	56.01	17.99	-13.76	0.00	0.00	-37.99
Constant	-0.05	0.02	0.01	-0.07	0.02	0.00	-0.00	-0.29
R2	0.63	0.37	0.49	0.63	0.37	0.04	0.10	0.11

Table 2.23: Social Housing, Decomposing the Firm Response

	$M_n(j)$ w/ Price Effects	$M_n(j)$ w/ Price Effects	$M_n(j)$ w/o Price Effects	$M_n(j)$ w/o Price Effects
Subsidy	-4.56		0.00	
Subsidy X ν_j		-1.94		-0.00
R2	0.23	0.97	1.00	1.00
Sector FX	yes	yes	yes	yes
Tract FX	no	yes	no	yes

2.12 Transitional Section

In the previous chapter, I study how the location choices of heterogeneous households and heterogeneous firms can create and amplify spatial segregation in cities. The purpose of this chapter is to show that the pecuniary externality arising from the two-sided sorting of households and firms interacts with place-based urban policy and yields different implications compared to a model with a citywide market for services and strong amenity spillovers. These implications were derived from a quantitative general equilibrium model of the city which has been fit to detailed microdata from Los Angeles. In the next chapter local markets play a crucial role in determining the difference between a local productivity-enhancing intervention in agriculture in developing countries and an intervention scaled to the the national level. As in the previous chapter the key for the difference between local and scaled intervention lies in the general equilibrium forces that operate through prices on the local markets for agricultural goods and labor. Methodologically, the two chapters are related throught the application and estimation of a spatial quantitative general equilibrium model and its estimation and calibration with rich microdata. In both chapters, the models are then applied to conduct counterfactuals to answer questions regarding the general equilibrium effects of policies.

Chapter 3

Scaling Agricultural Policy Interventions: Theory and Evidence from Uganda

3.1 Introduction

Roughly two thirds of the world's population living below the poverty line work in agriculture. In this context, interventions aimed at improving agricultural productivity, such as agricultural extension campaigns providing access, information, training and/or subsidies for modern production techniques, have played a prominent role in the global fight against poverty.[1] In order to inform these policies using rigorous evidence, much of the recent literature in this space has used variation in household outcomes from randomized control trials (RCTs) or natural experiments.

While rightly credited for revolutionizing the field of development economics, experiments and quasi-experiments frequently face the well-known limitation that local interventions or shocks often do not speak to the broader general equilibrium (GE) effects if the policy were to be scaled up to cover all farmers at the regional or national levels. For example, both the average and distributional effects of a fertilizer subsidy on household real incomes could substantially differ between a local intervention that leaves market prices unchanged and a policy at scale that affects output and factor prices across markets. The magnitude and distributional implications of these GE forces depend on a complex interplay of the policy's direct effect on yields across different crops, the pre-existing geography of household consumption and production choices across market places in the country, the size and nature of trade costs between farmers and local markets and across markets, the use of different factors of production across crops and sectors, expenditure shares on crops and sectors across the income distribution, as well as the responsiveness of production and consumption choices to the policy's direct effect on yields and prices.

While much has been written about the challenges of using local variation for informing policy-making at scale (e.g. Heckman & Smith (1995), Moffitt (2009), Ravallion (2009, 2018), Rodrik (2008), Deaton (2010)), these forces have not been studied and quantified using a combination of theory and administrative microdata covering all households and market places for the entire

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

country.[2] In this paper, we take a step in this direction through the lens of a rich but tractable quantitative GE model of farmer-level production and trade. To capture a number of salient features that we document in the agricultural empirical context, the model departs in several dimensions from the workhorse "gravity" structure of models in international trade and economic geography. In this environment, we propose a new solution approach that allows us to quantify GE counterfactuals without relying on structural gravity and without imposing stark new data requirements to be able to do so. We then bring to bear administrative microdata on pre-existing household locations, production, consumption and the transportation network within and across local markets to calibrate the model to the roughly 6 million households populating Uganda, and proceed to explore economy-wide GE counterfactuals at this granular level.

We use these building blocks to answer four central questions: i) To what extent do the average and distributional effects of an agricultural policy on household welfare differ between a local intervention and implementation at scale?; ii) How do these differences behave as a function of increasing rates of saturation going from a small number of farmers to 100 percent coverage?; iii) What is the role of modeling realistic trading frictions between agents in the economy, and the nature of these frictions, for the impact of scaling up?; and iv) What is the role of modeling individual households vs aggregating to regions or markets when quantifying the impact of scaling up? To fix ideas throughout the analysis, we focus on the impact of a subsidy for modern inputs (chemical fertilizers and hybrid seed varieties in our setting), but the framework is set up to study other types of interventions that are targeted at increasing agricultural productivity or farmers' market access more generally.

Our analysis proceeds in five steps. In the first step, we use the Ugandan microdata we have assembled described in Section 3.2 to document a number of stylized facts on farm production, consumption and trade. These stylized facts inform the structure of the model we develop in the second step. In terms of basic context, we document that farmers trade what they produce and consume in local markets, rather than purely living in subsistence, and that farmers adjust the share of land allocated across different crops across space and time. Using trader survey microdata, we find that local markets do not trade with one another on the vast majority of possible bilateral connections, suggesting that agricultural crops are mostly not differentiated across producers in this setting (Sotelo (2017)). We further document evidence for downward-sloping Engel curves in the share of agricultural consumption across rich and poor households facing the same market prices, suggesting non-homothetic preferences and a potential for distributional implications of GE effects on output prices. Next, we document that trade costs from farmers to local markets and between local markets are best captured by additive unit trade costs (charged per unit of weight) rather than ad valorem (iceberg), implying incomplete and decreasing price pass-through as a function of distance between markets. Finally, we show that the adoption of modern inputs, such as chemical fertilizer or hybrid seeds, changes the relative cost shares of traditional inputs (land and labor), suggesting that adoption of modern production techniques would not be well-captured by Hicks-neutral productivity shifts.

After laying out the model in step 2, step 3 proposes a new approach to quantifying GE coun-

[2]See discussion of related literature at the end of this section.

terfactuals in this rich environment. To explain our approach, it is useful to compare it to what is now standard practice using "exact-hat algebra" in the international trade literature (see for instance Adao *et al.* (2017)). This involves using the full matrix of bilateral trade flows and knowledge of the key elasticities or aggregate demand functions to solve for the (hat) changes of the endogenous variables given (hat) changes in the exogenous parameters. However, we do not observe the universe of crop-level bilateral trade flows at the individual farmer-to-farmer level, and even if one did have such data, it would be mostly made up of zeros, so the standard procedure would not be applicable. Instead, we show that solving for counterfactual changes in our environment requires knowledge of the full vector of pre-existing prices faced by agents across all markets. To address this challenge, we show that we can use estimates of trade costs in combination with data on household-level expenditure shares and production quantities to set up a "price discovery problem". This entails solving for equilibrium farm-gate prices and trade flows that rationalize the observed consumption and production decisions given a graph of trade costs connecting households and markets. In turn, with knowledge of farm-gate prices and trade costs, we can express farmer-level excess demand functions in terms of counterfactual prices and hat changes in farmer productivities (along with expenditure shares and production in the original equilibrium). We then use these excess demand functions and the no-arbitrage conditions to form a system of equations that we can use to solve for the counterfactual equilibrium.

This approach has several advantages. First, we are able to solve the model without relying on structural gravity and without imposing stark new data requirements (such as observing the full set of pre-existing market prices). Second, our solution method ensures that the economy is in equilibrium before solving for counterfactuals: the household prices we obtain from the price discovery are by construction consistent with the calibrated trade costs and the consumption and production decisions we observe in the data.[3] Finally, from a computational perspective our solution method is capable of handling high-dimensional GE counterfactuals at the level of individual households in the macroeconomy.

In step 4, we use the Ugandan microdata to calibrate the model to the roughly 6 million households who populate the country. To calibrate cross-market trade costs, we make use of estimates from Bergquist & McIntosh (2019), using newly collected market and trader survey microdata to provide information on bilateral market trade flows and local market prices at origin and destination across crops. To calibrate within-market trade costs between farmers and their local markets, we use observed gaps in the Ugandan National Panel Survey (UNPS) between farm-gate prices and local markets in combination with knowledge of farmer-level trade flows to and from the markets. To estimate the key supply elasticity of the model, we use the UNPS microdata and exploit plausibly exogenous changes in world market prices across different crops that propagate differently to local markets as a function of distance from the nearest border crossing.

In the final step, we use the calibrated model to conduct the counterfactual analysis and present a number of additional robustness and model validation tests using natural experiments over our sample period (road additions and weather shocks). To investigate the difference in changes in

[3]For example, Sotelo (2017) uses province-level crop unit values from agricultural surveys to calibrate and solve the model, but these price data are not model-consistent with the calibrated trade costs.

economic outcomes across local interventions and scaling up, we randomly select a representative sample of ten thousand rural households in Uganda (roughly 0.2 percent of Ugandan households). We then solve for counterfactual GE changes in household-level economic outcomes due to an intervention that only targets the subsidy for modern inputs at these ten thousand households, and compare the effects on those same households to an intervention that scales the subsidy policy to all rural households in the country. In particular, we investigate differences in both the average and distributional welfare effects across households, and document the underlying channels.

In our preliminary set of results (work in progress at the time of this writing), we find that the average effect of a subsidy for modern inputs on rural household welfare can differ substantially when comparing effects in the local intervention compared to effects on the same households when scaling up to the national level. In addition, the policy's distributional implications differ significantly: while the local intervention is strongly regressive (benefiting initially richer rural households the most), the welfare gains are significantly more evenly distributed under the intervention at scale. Underlying these findings, we document GE price effects on crops and factors of production that propagate along the trading network and affect nominal incomes as well as household price indices differently compared to the local intervention.

Recognizing that GE effects at scale can play an important role, recent empirical studies in development economics have used two-stage cluster randomization designs (e.g. Baird *et al.* (2011), Burke *et al.* (2019)), Egger *et al.* (2019)) to measure treatment effects at different levels of saturation across clusters of markets. Due to constraints on statistical power, such designs typically limit the comparison to just two discrete levels of saturation (often chosen *ad hoc*). If the aim is to extrapolate from these two points of saturation to make inference on what treatment effects would look like at 100% saturation, however, one must make the assumption that GE forces are both monotonic and linear with respect to changes in saturation rates.[4] To investigate this question, we quantify GE effects on an identical sample of Ugandan households across various levels of national saturation going from 0-100 percent. We find that the GE forces appear to be both non-linear and non-monotonic as a function of the national saturation rate. This finding suggests caution regarding unobserved non-linearities when extrapolating from the results of randomized saturation designs to policy implications at full scale.

In the third question, we study the importance of allowing for realistic trade costs between agents when aggregating the average and distributional implications of shocks in the macroeconomy. As we discuss below, the recent macroeconomics literature on aggregation mostly abstracts from such frictions. To this end, we estimate two additional sets of counterfactuals (both for the local intervention and the scaled intervention). We first re-estimate our baseline counterfactuals after setting the trading frictions between households and local markets and across local markets to zero. Second, we re-estimate the counterfactuals after specifying and calibrating the model featuring ad valorem (iceberg) trading frictions, instead of our baseline counterfactuals with additive unit costs. We find that the welfare gains from scaling up the policy are affected in both the average

[4]Another necessary assumption for extrapolating results from randomized saturation designs to at-scale policies is that there are no GE spillover effects across the clusters. We plan to present additional results on this question in future versions of this draft.

and distributional effects subject to these alternative assumptions.

In the fourth question, we explore the implications of modeling households at a granular level compared to aggregating them into different clusters of regional representative agents across Uganda. In our preliminary results, we find that preserving trading frictions at a granular level matters for the counterfactual results.

Related Literature In addition to the references discussed above, this paper relates and contributes to a number of different literatures. It relates to a large and growing number of studies using experiments or quasi-experiments to evaluate the impact of agricultural policy interventions in (e.g. Caldwell et al. (2019), Carter *et al.* (2014), Duflo *et al.* (2011), Emerick *et al.* (2016), Magruder (2018)). Relative to the existing literature in this space, our objective is to study and quantify how the impacts found in relatively small-scale interventions are subject to change if the same intervention were to be scaled up at different levels of saturation, and to disentangle the underlying mechanisms.[5] Our hope is to provide a useful methodological toolkit that can be used to complement the results from observed local interventions to evaluate interventions at scale.

From a methodological point of view, our framework contributes to a growing literature in macroeconomics that has sought to quantify the aggregation of observed local shocks if they were to occur to all agents in the economy (e.g. Buera *et al.* (2017), Baqaee & Farhi (2018), Sraer & Thesmar (2018), Fujimoto *et al.* (2019)). A common feature of this literature is its treatment of the macroeconomy as one single integrated market in which all agents interact without trading frictions and face identical prices. In our framework, each household faces trading frictions and imperfect pass-through for buying and selling output to local markets, for both goods and factors of production, and in turn trade flows across local markets are subject to trade costs and imperfect pass-through along the transportation network. We calibrate these trading frictions using survey data on household farm-gate prices, local market prices and information on bilateral trade flows across local markets. By studying counterfactuals before and after setting trade costs close to zero, this allows us to investigate the importance of modeling realistic trading frictions between agents when solving for both the average and distributional implications of a shock in the aggregate.

Our methodology also relates to an earlier literature on computable general equilibrium (CGE) models in development economics (see e.g. de Janvry & Sadoulet (1995) for a review).[6] Our framework and analysis depart from this literature in at least three important respects. First, due to computational constraints as well as much less availability of rich survey microdata at the time, these models are usually based on one (or a small number of) representative agents that make up the macroeconomy. In contrast, our framework embraces the full degree of heterogeneity across individual households that we observe in the initial equilibrium. Second, as in the macroeconomics literature on aggregation discussed above, these models largely abstract from trading frictions across households and markets, and model the economy as one integrated market instead of local market

[5] See Svensson & Yanagizawa-Drott (2012) for a cautionary tale on how estimated impacts in partial and general equilibrium can diverge for agricultural interventions in Uganda.

[6] This literature has also been referred to as "multi-market" analysis, as the impact of shocks is traced across multiple output and factor markets in the economy.

places that are connected along a transportation graph. Third, the CGE literature relied on what sometimes has been referred to as a "black box" of numerous parameters whose values determine the responsiveness to different shocks across different sectors, output markets and factor markets. Given a large number of such parameter values, it becomes hard to judge ex-post which of these parameter combinations affect the simulation results and to what extent. In contrast, our model follows recent work on quantitative GE models in international trade and economic geography (e.g. Eaton & Kortum (2002)), which has the benefit of greater tractability and transparency. Instead of dozens or hundreds of parameters governing counterfactual results, our framework highlights a small set of key elasticities in both production and demand, whose impact on the main findings can be readily assessed across alternative parameter ranges.

Finally, our theoretical framework builds on recent work using quantitative models in international trade and economic geography (e.g. Allen & Arkolakis (2014), Redding (2016)). Given the empirical context, we depart from the workhorse gravity structure most commonly used in this literature. As discussed above, we therein build on recent work by e.g. Costinot & Donaldson (2016), Fajgelbaum & Khandelwal (2016), Adao *et al.* (2017), Sotelo (2017) and Adao *et al.* (2018). Given our focus on agriculture, the paper is in particular related to Sotelo (2017) and Costinot & Donaldson (2016). Relative to Sotelo (2017) the main differences are that we set out to quantify counterfactuals at the level of individual households rather than representative agents at the level of provinces, and that we propose a new solution method allowing us to quantify counterfactual changes in absence of structural gravity without imposing stark new data requirements (such as observing pre-existing prices for all agents).[7] Relative to Costinot & Donaldson (2016) the main differences are, again, households versus regions, that we model trade flows between all markets rather than to one national hub and that we aim at a welfare analysis (requiring a demand side) rather than focusing on the production side. Finally, in terms of calibration and solution method, the key source of information used by Costinot & Donaldson (2016) to construct production possibility frontiers across crops –the FAO GAEZ database– would not be suitable for an analysis at the household level.[8]

The remainder of the chapter is structured as follows. Section 2 presents the data and stylized facts. Section 3 develops the model and the solution method. Section 4 presents the calibration. Section 5 presents the counterfactual analysis, robustness and model validation. Section 6 concludes.

3.2 Data and Stylized Facts

In this section we briefly describe the database we have assembled. We then use these data to document the empirical context and a number of stylized facts.

[7]Our framework also differs by allowing for non-homothetic preferences, additive trade costs and different technology regimes in production.

[8]In the case of Uganda, FAO GAEZ covers roughly 2100 5-minute arc grid cells ($> 100km^2$ on average). In contrast, the survey data suggest rich variation in crop suitability across small plots of land (including within farmers), a feature that our model and calibration aim to embrace.

Data

Our analysis makes use of four main datasets.

Uganda National Panel Survey (UNPS)

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

Uganda Population and Housing Census 2002

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

Survey Data on Cross-Market Trade Flows and Trade Costs

The survey data collected by Bergquist & McIntosh (2019)) can be used to shed light on cross-market trade flows and 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 the markets in which they 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 collected for maize, beans, and other crops. A greater description of the data collection can be found in Bergquist & McIntosh (2019).

GIS Database and World 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 Office of National Statistics. We complement this database with geo-referenced information on crop suitability from the Food and Agricultural Organization (FAO) Global Ago-Ecological Zones (GAEZ) database. This dataset uses an agronomic model of crop production to convert data on terrain and soil conditions, rainfall, temperature and other agro-climatic conditions to calculate the potential production and yields of a variety of crops. We use this information as part of the projection from the UNPS sample to the Ugandan population at large. Finally, we use information on world crop prices (faced by other African nations over time) from the FAO statistics database.

Context and Stylized Facts

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

Figures 3.1, 3.2 and Tables 3.1-3.5 present a number of basic stylized facts about the empirical context. Unless otherwise stated, these are drawn from the UNPS panel data of farmers. First, Table 3.1 documents that the 9 most commonly grown crops (matooke (banana), beans, cassava, coffee, groundnuts, maize, millet, sorghum and sweet potatoes) account for 99 percent of the land allocation for the median farmer in Uganda (and for 86 percent of the aggregate land allocation).

Second, Figure 3.1 and Table 3.2 document a significant degree of regional specialization in Ugandan agricultural production across regions. Table 3.2 provides information that these regional differences translate into meaningful variation in regional market prices across crops: the across-district variation in average crop prices accounts for 20-60 percent of the total variation in observed farm-gate prices.

Third, Table 3.3 documents that the majority of all farmers are either net sellers or net buyers, rather than in subsistence, and this holds across each of the 9 major crops. The table also presents evidence that there are significant movements in and out of subsistence, conditional on having observed subsistence at the farmer level in a given season. Fourth, Table 3.4 documents that farmers buy and sell their crops mostly in local markets, which in turn are connected to other markets through wholesale traders. Finally, Table 3.5 documents that farmers frequently reallocate their land allocations across crops over time.

Product Differentiation Across Farmers

Table 3.6 looks at evidence on product differentiation across farmers. The canonical approach in models of international trade sets focus on trade in manufacturing goods across countries, where CES demand coupled with product differentiation across manufacturing varieties imply that all bilateral trading pairs have non-zero trade flows. In an agricultural setting, however, and focusing on households instead of entire economies, this assumption would likely be stark. Consistent with this, the survey data collected by Bergquist & McIntosh (2019) 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 3.6 provides corroborating evidence that agricultural crops in the Ugandan empirical setting are unlikely well-captured by the assumption of product differentiation across farmers who produce the crops.

Household Preferences

Figure 3.2 reports a non-parametric estimate of the household Engel curve for food consumption. We estimate flexible functional forms of the following specification:

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

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.

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 & McIntosh (2019). Specifically, we estimate:

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

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

Following Bergquist & McIntosh (2019), we estimate these specifications conditioning on market pairs for which we observe positive trade flows in a given month. If trade costs include an ad valorem component, we would expect the coefficient β to be positive and statistically significant. On the other hand, if trade costs are charged per unit of the shipment (e.g. per sack), we would expect the point estimate of β to be close to zero.

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

As reported in Table 3.7, we find that β is slightly negative and statistically significant in the OLS regressions, but very close to zero and statistically insignificant after addressing the concern of correlated measurement errors in the IV specification. Taken together with existing evidence from field work (e.g. Bergquist (2017)), these results suggest that trade costs in this empirical setting are best-captured by per-unit additive transportation costs.

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:

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

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

As reported in Table 3.8, we find that the share of labor costs relative to land costs increases significantly as a function of whether or not the farmer uses modern production techniques. This holds both before and after the inclusion of farmer fixed effects (using variation only within-farmer across crops or over time). These results suggest that modern technology adoption is unlikely to be well-captured by a simple Hicks-neutral productivity shift in the production function.

3.3 Model and Solution Method

In this section we develop a rich but tractable GE model of farm production and trading that is able to capture the stylized facts we document in Section 3.2 above. We present a model that features heterogeneous producers and consumers who interact across a complex geography.

The economy is populated by farmers that are endowed with land of heterogeneous suitability for different homogeneous crops. These producers choose an optimal land allocation across crops taking output and input prices as given. Price differences are driven by trade costs across farmers and their local markets as well as trade costs across local markets that weaken specialization and induce some farmers to stay in subsistence farming for certain crops. Trade costs also reduce the amount of modern intermediate inputs, like fertilizer, used in production. Trade costs are driven by the farmer's location relative to the local market and the position of the local market relative to the rest of the economy.

Farmers are assigned to a local market and can trade goods and labor on that market. These local markets are connected with all other markets and the rest of the world by a graph based on existing infrastructure. Given increasing policy attention to promoting modern technology adoption, like the usage of chemical fertilizer or hybrid seed varieties, the model allows farmers to change their production technology in response to an intervention. We further augment the model by introducing a homogeneous and tradable manufacturing good produced by urban households. The economy is small in the sense that it does not affect international prices of crops, the manufacturing good, or the agricultural intermediate good in the rest of the world.

In contrast to the standard approach in the literature, and consistent with the stylized fact presented in the previous section, we assume that trade costs have both an additive and an iceberg component, and that preferences are non-homothetic. Additive trade costs are in terms of some good (e.g., gasoline) that is imported from the rest of the world, and hence its price is not affected by our small economy. This will simplify our analysis since we do not need trace potential effects of the scaling up of the intervention on local trade costs across the geography of Uganda. As in most of the development literature, we allow for non-homotheticity in preferences to capture the large disparity in the share of income spent on food, and allowing for potential distributional implications through the price index.

Environment

There are two kinds of agents, farmers (indexed by i) and urban households (indexed by h), and two kinds of markets, villages (indexed by v) and urban centers (indexed by u). There will also be an agent that we call Foreign (denoted by F) and stands for the rest of the world. In general, each of these nodes (farmers or markets) in the economy is indexed by o (origin) or d (destination) when dealing with the trade network, and with j (households i or h or Foreign F) or m (market) when dealing with agent behavior or market clearing conditions, respectively.

These nodes trade in outputs ($k \in \mathscr{K}$) and inputs ($n \in \mathscr{N}$). The are two kinds of outputs, agricultural goods ($k \in \mathscr{K}_A$) and a manufacturing good ($k = M$), and two kinds of inputs, intermediate goods ($n \in \mathscr{N}_I$) and labor ($n = L$). We use g as a generic index that encompasses both outputs and inputs, hence $g \in \mathscr{G} \equiv \mathscr{K} \cup \mathscr{N}$.

Farmers own land and labor in quantities Z_i and L_i, and they produce agricultural goods using their own land (i.e., land is not tradable) as well as labor and intermediate goods. Urban households own labor in quantity L_h and produce the manufacturing good using labor. Intermediate goods are imported from Foreign.

Let $\mathscr{O}_g(d)$ denote the set of origin nodes from where node d can obtain good g and $\mathscr{D}_g(o)$ denote the set of destination nodes which can obtain good g from node o. In other words, trade in good g from o to d only happens if $o \in \mathscr{O}_g(d)$ or, equivalently, $d \in \mathscr{D}_g(o)$. We assume that a farmer i can trade only with the village v in which she is located, that is, $\mathscr{O}_g(i) = \mathscr{D}_g(i) = \{v\}$ for all g. Similarly, an urban household h can trade only with the urban center u in which it is located, that is, $\mathscr{O}_g(h) = \mathscr{D}_g(h) = \{u\}$ for all g. Further, while each village can consist of multiple farmers, each urban center consists of one representative household. Labor is not tradable across markets, i.e., $m' \notin \mathscr{O}_L(m)$ for all markets $m \neq m'$.

Let $p_{j,g}$ denote the price at which agent j buys or sells good g, and let $p_{m,g}$ is the price at which good g is bought or sold at market m. Trade in good g from o to $d \in \mathscr{D}_g(o)$ is subject to iceberg and additive trade costs. Iceberg trade costs are $\tau_{od,g}$ and additive trade costs are $t_{od,g}$ in units units of a "transportation good." We use index T for this transportation good and assume that it is produced by Foreign at price $p^*_{F,T}$, and further assume that there are no trade costs for this good, so that all agents can access this good at price $p^*_{F,T}$. Thus, for example, if a farmer buys good g from her village v, her farm-gate price is $p_{i,g} = \tau_{vi,g}\left(p_{v,g} + p^*_{F,T} t_{vi,g}\right)$. We take this "transportation good" as the numeraire and so we set $p^*_{F,T0} = 1$. Finally, we assume that our economy is "small' in the sense that Foreign is willing to buy from or supply to it any amount of any good g at exogenous prices $p^*_{F,g}$.

Preferences

Agent j has an indirect utility function $V_j\left(\left\{a_{j,k}p_{j,k}\right\}, I_j\right)$, where I_j denotes income and $a_{j,k}$ and $p_{j,k}$ denote taste shifters and prices of goods $k \in \mathscr{K}$ for agent j. Let $\xi_{j,k}$ denote the expenditure share of agent j on good k and let $\varphi_{j,k}$ denote the corresponding expenditure share function. Roy's identity implies that

$$\xi_{j,k} = \varphi_{j,k}\left(\left\{a_{j,k'}p_{j,k'}\right\}_{k'}, I_j\right) \equiv -\frac{\frac{\partial \ln V_j\left(\left\{a_{j,k'}p_{j,k'}\right\}_{k'}, I_j\right)}{\partial \ln p_{j,k}}}{\frac{\partial \ln V_j\left(\left\{a_{j,k'}p_{j,k'}\right\}_{k'}, I_j\right)}{\partial \ln I_j}}.$$

Further, letting $\varphi_j\left(\left\{a_{j,k}p_{j,k}\right\}_k, I_j\right) \equiv \left\{\varphi_{j,k}\left(\left\{a_{j,k'}p_{j,k'}\right\}_{k'}, I_j\right)\right\}_k$, we assume that $\varphi_j(\bullet)$ is invertible so that one can obtain $\left\{a_{j,k}p_{j,k}\right\}_k$ (up to a normalization for prices) as

$$\left\{a_{j,k}p_{j,k}\right\}_k = \varphi_j^{-1}\left(\left\{\xi_{j,k}\right\}_k, I_j\right). \tag{3.1}$$

Technology

We start with farmers and then describe urban households. A farmer can produce good $k \in \mathscr{K}_A$ with $\omega \in \Omega$ techniques. For farmer i, technique ω uses inputs $n \in \mathscr{N}$ in a Cobb-Douglas production function with shares $\alpha_{i,n,k,\omega}$ where we assume that $\sum_n \alpha_{i,n,k,\omega} < 1$. It can be easily established that the return to a unit of effective land allocated to good k with technique ω is

$$\tilde{v}_{i,k,\omega} \equiv \tilde{b}_{i,k,\omega}\eta_{i,k,\omega}\left(\frac{p_{i,k}}{\prod_n p_{i,n}^{\alpha_{i,n,k,\omega}}}\right)^{\frac{1}{1-\sum_n \alpha_{i,n,k,\omega}}},$$

where $\tilde{b}_{i,k,\omega}$ is a technology shifter and $\eta_{i,k,\omega}$ is a constant.[9] The function defining land returns for farmer i is given by

$$Y_i\left(\left\{v_{i,k,\omega}\right\}_{k,\omega}\right) \equiv \max_{\left\{Z_{i,k,\omega}\right\}} \sum_{k,\omega} v_{i,k,\omega} Z_{i,k,\omega} \quad \text{s.t.} \quad f_i(\left\{Z_{i,k,\omega}\right\}_{k,\omega}) \leq Z_i,$$

where $v_{i,k,\omega} \equiv \left(\frac{b_{i,k,\omega}p_{i,k}}{\prod_n p_{i,n}^{\alpha_{i,n,k,\omega}}}\right)^{\frac{1}{1-\sum_n \alpha_{i,n,k,\omega}}}$ and $b_{i,k.\omega} \equiv \left(\tilde{b}_{i,k,\omega}\eta_{i,k,\omega}\right)^{1-\sum_n \alpha_{i,n,k,\omega}}$. Here $Z_{i,k,\omega}$ can be understood as the effective units of land allocated to producing agricultural good k with technique ω. We assume that $f_i(\bullet)$ is strictly quasiconvex so that the maximization problem has unique solution.

Consistent with the stylized facts, the input shares are allowed to vary across techniques. When we get to the model calibration in Section 3.4, we will allow input shares to differ across Ugandan regions, and we will allow for only two techniques: traditional, $\omega = 0$, and modern, $\omega = 0$. We will map these two techniques to data in terms of observed use of modern intermediates (such as chemical fertilizer or hybrid seeds) in production: the traditional technique makes use of land and labor (with $\alpha_{k,0} = 0$), whereas the modern technique adopts the use intermediates (with $\alpha_{k,1} > 0$). Thus, the choice of a modern technique will increase the importance of intermediates and decrease the importance of land or labor.

Let $\pi_{i,k,\omega} \equiv \frac{v_{i,k,\omega}Z_{i,k,\omega}}{\sum_{k',\omega'} v_{i,k',\omega'}Z_{i,k',\omega'}}$ denote the share of land returns coming from production of crop k with technique ω and let $\psi_{i,k,\omega}(\bullet)$ denote the corresponding share function. An envelope result implies that

$$\pi_{i,k,\omega} = \psi_{i,k,\omega}\left(\left\{v_{i,k',\omega'}\right\}_{k',\omega'}\right) = \frac{\partial \ln Y_i\left(\left\{v_{i,k',\omega'}\right\}_{k',\omega'}\right)}{\partial \ln v_{i,k,\omega}}.$$

In turn, demand for input n (in value) as a ratio of land returns is

$$\phi_{i,n}\left(\left\{v_{i,k',\omega'}\right\}_{k',\omega'}\right) = \sum_{k,\omega}\left(\frac{\alpha_{i,n,k,\omega}}{1-\sum_{n'}\alpha_{i,n',k,\omega}}\right)\psi_{i,k,\omega}\left(\left\{v_{i,k',\omega'}\right\}_{k',\omega'}\right).$$

---

[9]In particular, $\eta_{i,k,\omega} = \left[\left(1-\sum_n \alpha_{i,n,k,\omega}\right)\prod_n \alpha_{i,n,k,\omega}^{\frac{\alpha_{i,n,k,\omega}}{1-\sum_n \alpha_{i,n,k,\omega}}}\right]^{-1}$

Finally, letting $\psi_i\left(\left\{v_{i,k,\omega}\right\}_{k.\omega}\right) \equiv \left\{\psi_{i,k,\omega}\left(\left\{v_{i,k',\omega'}\right\}_{k',\omega'}\right)\right\}_{k,\omega}$, we assume that $\psi_i(\bullet)$ is invertible so that one can obtain $\left\{v_{i,k,\omega}\right\}_{k,\omega}$ (up to a normalization for prices) as

$$\left\{v_{i,k,\omega}\right\}_{k,\omega} = \psi_i^{-1}\left(\left\{\pi_{i,k,\omega}\right\}_{k,\omega}\right). \tag{3.2}$$

Now we turn to urban households. These households produce the manufacturing good. We keep their technology simple by assuming that production is linear in labor, so that the quantity of the manufacturing good produced is given by $b_{h,M}L_h$. Given that labor supply is perfectly inelastic, we can then simply treat $y_{h,M} \equiv b_{h,M}L_h$ as the urban households' endowment of the manufacturing good.

Equilibrium

We assume that all markets are perfectly competitive. In equilibrium, rural and urban households maximize utility taking prices as given, prices respect no-arbitrage conditions given trade costs, and all markets clear. To formalize this definition, let $\chi_{j,g}\left(\left\{a_{j,k}p_{j,k}\right\}_k, \left\{v_{j,k,\omega}\right\}_{k,\omega}, I_j\right)$ be the excess demand (in value) of agent j for good g given prices of outputs and inputs. The equilibrium is a set of prices, $\left\{p_{j,g}\right\}$ and $\left\{p_{m,g}\right\}$, and trade flows (in quantities), $\left\{x_{od,g}\right\}$, such that excess demand is equal to the difference between purchases and sales for each agent j and good g,

$$\chi_{j,g}\left(\left\{a_{j,k}p_{j,k}\right\}_k, \left\{v_{j,k,\omega}\right\}_{k,\omega}, I_j\right) = p_{j,g}\left(\sum_{o\in\mathscr{O}_g(j)} x_{oj,g} - \sum_{d\in\mathscr{D}_g(j)} x_{jd,g}\right) \quad \forall j \in \mathscr{J}\setminus\{F\}, g, \tag{3.3}$$

$$\chi_{j,g}\left(\left\{p_{j,g}\right\}_g\right) = p_{j,g}\left(\sum_{o\in\mathscr{O}_g(j)} x_{oj,g} - \sum_{d\in\mathscr{D}_g(j)} x_{jd,g}\right) \quad \forall j \in \{F\}, g, \tag{3.4}$$

markets clear,

$$\sum_{d\in\mathscr{D}_g(m)} x_{md,g} = \sum_{o\in\mathscr{O}_g(m)} x_{om,g} \quad \forall m, g, \tag{3.5}$$

and no-arbitrage conditions hold,

$$\tau_{od,g}\left(p_{o,g} + t_{od,g}\right) \geq p_{d,g} \perp x_{od,g} \quad \forall d \in \mathscr{D}_g(o), g. \tag{3.6}$$

Here the symbol $\perp$ between a weak inequality and a variable indicates that the weak inequality holds as equality if the variable is strictly positive. For example, if farmer i sells good k to market v then $x_{iv,k} > 0$ and we must have $p_{v,k} = \tau_{iv,k}(p_{i,k} + t_{iv,k})$, while the converse implies that if $p_{v,k} > \tau_{iv,k}(p_{i,k} + t_{iv,k})$, then $x_{iv,k} = 0$. The excess demand functions $\chi_{j,g}(\bullet)$ for farmers, urban households and Foreign are determined by the results in the previous subsections, and can be found in Appendix 2.

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

Counterfactual Analysis

We are interested in computing the effect of a shock to preferences, technology, and trade costs. Using hat notation (i.e., $\hat{x} = x'/x$), these shocks are given by $\left\{\hat{a}_{j,k}\right\}$, $\left\{\hat{b}_{j,k,\tau}\right\}$, and $\left\{\hat{\tau}_{od,k}, \hat{t}_{od,k}\right\}$. In the counterfactual equilibrium, equations 3.3-3.6 can be written as

$$\chi_{j,g}\left(\left\{\hat{a}_{j,k}\hat{p}_{j,k}\varphi_{j,k}^{-1}\left(\left\{\xi_{j,k'}\right\}_{k'}, I_j\right)\right\}_k, \left\{\hat{v}_{j,k,\tau}\psi_{j,k,\omega}^{-1}\left(\left\{\pi_{j,k',\omega'}\right\}_{k',\omega'}\right)\right\}_{k\in\mathscr{K}_A}, \hat{I}_j I_j\right)$$
$$= p'_{j,g}\left(\sum_{o\in\mathscr{O}_g(j)} x'_{oj,g} - \sum_{d\in\mathscr{D}_g(j)} x'_{jd,g}\right)\forall j\in\mathscr{J}\setminus\{F\}, g, \tag{3.7}$$

$$\chi_{j,g}\left(\left\{\hat{p}_{j,g}\right\}_g\right) = p'_{j,g}\left(\sum_{o\in\mathscr{O}_g(j)} x'_{oj,g} - \sum_{d\in\mathscr{D}_g(j)} x'_{jd,g}\right)\forall j\in\{F\}, g, \tag{3.8}$$

$$\sum_{d\in\mathscr{D}_g(m)} x'_{md,g} = \sum_{o\in\mathscr{O}_g(m)} x'_{om,g}\forall m, g, \tag{3.9}$$

$$\tau'_{od,g}\left(p'_{o,g} + t'_{od,g}\right)p'_{o,g} \geq p'_{d,g} \perp x'_{od,g}\forall d\in\mathscr{D}_g(o), g, \tag{3.10}$$

where

$$\hat{v}_{i,k,\omega} = \left(\frac{\hat{b}_{i,k,\omega}\hat{p}_{i,k}}{\prod_n \hat{p}_{i,n}^{\alpha_{i,n,k,\omega}}}\right)^{\frac{1}{1-\sum_n \alpha_{i,n,k,\omega}}}\forall i\in\mathscr{J}, k\in\mathscr{K}_A, n\in\mathscr{N}, \omega\in\Omega. \tag{3.11}$$

The term on the LHS of Equation 3.7 is in terms of hat changes, as in exact-hat algebra, but the RHS of that equation as well as Equations 3.9 and 3.10 are in terms of counterfactual levels. This implies that in this system we have prices both in hat changes and counterfactual levels, $\left\{\hat{p}_{j,g}, \hat{p}_{m,g}\right\}$ and $\left\{p'_{j,g}, p'_{m,g}\right\}$, so we need the original prices $\left\{p_{j,g}, p_{m,g}\right\}$ to solve the system. We propose to recover these prices in a manner that is consistent with the model and the variables observed in microdata.

We observe expenditure shares for farmers and urban households, $\left\{\xi_{i,g}, \xi_{h,g}\right\}$, crop output levels for farmers $\left\{y_{i,k,\omega}\right\}$, output of manufacturing for urban households $\left\{y_{h,M}\right\}$, labor endowments of farmers, $\{L_i\}$, cost shares of farmers $\left\{\alpha_{i,n,k,\omega}\right\}$, and calibrated trade costs $\left\{t_{od,g}\right\}$. Let $\mathbb{D} \equiv \left\{\xi_{i,g}, \xi_{h,g}, y_{i,k,\omega}, y_{h,M}, L_i, \alpha_{i,n,k,\omega}, \tau_{od,g}, t_{od,g}, p^*_{F,g}\right\}$. First, we recast excess demand functions

$\chi_{j,g}(\bullet)$ as functions of data $\mathbb{D}$ and prices $\{p_{j,g}\}$ for farmers, urban households and Foreign (see Appendix 2 for expressions).

We then solve for prices $\{p_{j,g}, p_{m,g}\}$ in the initial equilibrium as a solution to the following system of equations in line with equations 3.3-3.6:

$$\chi_{j,g}\left(\{p_{j,g'}\}_{g'};\mathbb{D}\right) = p_{j,g}\left(\sum_{o\in\mathscr{O}_g(j)} x_{oj,g} - \sum_{d\in\mathscr{D}_g(j)} x_{jd,g}\right)\forall j,g, \tag{3.12}$$

$$\sum_{d\in\mathscr{D}_g(m)} x_{md,g} = \sum_{o\in\mathscr{O}_g(m)} x_{om,g}\forall m,g, \tag{3.13}$$

$$\tau_{od,g}\left(p_{o,g}+t_{od,g}\right) \geq p_{d,g} \perp x_{od,g}\forall d\in\mathscr{D}_g(o),g. \tag{3.14}$$

Because the price discovery step is tantamount to finding the equilibrium of an exchange economy, we can follow well-known methods for establishing uniqueness of equilibria in such an economy to uncover conditions under which price discovery yields a unique solution. As shown in Appendix 2, for a special case of our model with no additive trade costs and no trade with Foreign, we can show that, if there is a set of prices under which all agents are directly or indirectly connected through trade, then this is the unique set of prices that solves the price discovery step. We are currently working on extending this result to a less restrictive setting.

Using prices $\{p_{j,g}, p_{m,g}\}$ thus obtained, data $\mathbb{D}$ and shocks $\{\hat{a}_{j,k}, \hat{b}_{j,k,\omega}\}$ to the initial equilibrium, we evaluate the excess demand functions for all agents in the counterfactual equilibrium with the respective components computed as follows:

1. $\{I_j\}$ for farmers and urban households respectively as

$$I_i\left(\{p_{i,g}\}_g;\mathbb{D}\right) = \sum_{k,\omega}\left(1-\sum_n \alpha_{i,n,k,\omega}\right)p_{i,k}y_{i,k,\omega} + p_{i,L}L_i,$$

$$I_h\left(\{p_{h,g}\}_g;\mathbb{D}\right) = p_{h,M}y_{h,M},$$

2. $\{\pi_{i,k,\omega}\}$ as in $\pi_{i,k,\omega} = \frac{\left(1-\sum_n \alpha_{i,n,k,\omega}\right)p_{i,k}y_{i,k,\omega}}{\left(1-\lambda_{i,L}\right)I_i}$, where $\lambda_{i,L} = \frac{p_{i,L}L_i}{I_i}$ is the share of farmer's total income coming from wage income,

3. $\{\hat{I}_j\}$ for farmers and urban households respectively as

$$\hat{I}_i = (1-\lambda_{i,L})\sum_k \pi_{i,k}\hat{p}_{i,k} + \lambda_{i,L}\hat{p}_{i,L},$$

$$\hat{I}_h = \hat{p}_{h,M},$$

4. $\{\hat{v}_{i,k,\omega}\}_{k,\omega}$ as in eq. 3.11.

Finally, we can obtain counterfactual trade flows $\{x'_{od,g}\}$ and prices $\{p'_{j,g}, p'_{m,g}\}$ as a solution to the system of equations 3.7-3.10.

Parametrization

Motivated by the large differences across households in expenditure shares on food in the data, we assume non-homothetic preferences between food and manufacturing. In particular, we assume that upper tier preferences are Stone-Geary, so that households need to consume a minimum amount of the crop composite, $\bar{C}_A$. In turn, crops are aggregated into a CES composite with elasticity of substitution σ. The indirect utility function is then

$$V_j\left(\left\{a_{j,k}p_{j,k}\right\}_k, I_j\right) = \frac{I_j - P_{j,A}\bar{C}_A}{P_{j,A}^{\zeta}p_{j,M}^{1-\zeta}},$$

with

$$P_{j,A} = \left(\sum_{k\in\mathscr{K}_A}\left(a_{j,k}p_{j,k}\right)^{-(\sigma-1)}\right)^{-\frac{1}{\sigma-1}}.$$

This implies that

$$\xi_{j,k} = \varphi_{j,k}\left(\left\{a_{j,k}p_{j,k}\right\}_k, I_j\right) = \frac{\left(a_{j,k}p_{j,k}\right)^{-(\sigma-1)}}{P_{j,A}^{-(\sigma-1)}}\left(\zeta + (1-\zeta)\frac{P_{j,A}\bar{C}_A}{I_j}\right)$$

for $k \in \mathscr{K}_A$ and $\xi_{j,M} = (1-\zeta)\left(1 - \frac{P_{j,A}\bar{C}_A}{I_j}\right)$.

On the production side, we assume that

$$f_i(\left\{Z_{i,k,\tau}\right\}_{k,\tau}) = \gamma^{-1}\left(\sum_{k,\tau}Z_{i,k,\tau}^{\kappa/(\kappa-1)}\right)^{(\kappa-1)/\kappa}$$

with $\kappa > 1$, for some positive constant γ. It is easy to verify that this can be obtained from the Roy-Frechet microfoundations in Costinot & Donaldson (2016) and Sotelo (2017).[10] We then have

$$Y_i\left(\left\{v_{i,k,\tau}\right\}_{k,\tau}\right) = \gamma\left(\sum v_{i,k,\tau}^{\kappa}\right)^{1/\kappa}Z_i.$$

This implies that

$$\psi_{i,k,\tau}\left(\left\{v_{i,k',\tau'}\right\}_{k',\tau'}\right) = \frac{v_{i,k,\tau}^{\kappa}}{\sum_{k',\tau'}v_{i,k',\tau'}^{\kappa}}.$$

Given this setup, our framework allows for substituting into or out of traditional versus modern production techniques as long as some small amount of output is produced under both regimes (since zero production would imply a zero productivity draw in this setting). In our calibration in Section 3.4 we allow for this extensive margin by attributing 1 percent of total crop output to modern or traditional production regimes in cases where only one technique is observed for a given farmer in the microdata.[11]

[10] Such microfoundations would imply the need to restrict $\kappa > 1$, but this restriction is not necessary for the more general case of a PPF with a constant elasticity of transformation that we work with here.

[11] In ongoing work in progress, we explore the sensitivity of the counterfactuals to this ad-hoc choice and are working on alternative ways to dealing with the extensive margin.

3.4 Calibration

Building on the the results of the previous section, we calibrate the model to the Ugandan economy in two main steps. In the first step, we describe the calibration of trade frictions between individual households and their local markets (t_{img}) and across local markets (t_{odg}), and the calibration of the demand and supply parameters (ζ, σ, $\alpha_{k\omega}$, $\beta_{k\omega}$, $\gamma_{k\omega}$ and κ). In the second step, we use the survey data on household expenditure shares across crops and sectors and crop quantities produced (ξ_{ik} and $y_{ik\omega}$) from the UNPS panel data, and extrapolate this information to the Ugandan population at large. To this end, we use the microdata on household locations and their characteristics from the 100 percent sample of the Ugandan population census in 2002 described in Section 3.2.

Using the solution method of the previous section, this combination of parameter values and raw dis-aggregated information on pre-existing household consumption and production choices allows us to solve for unobserved farm-gate and market prices (p_{ig} and p_{mg}) and household revenue shares ($\pi_{ik\omega}$) for the whole of Uganda. This, in turn, allows us to use exact hat algebra to solve for GE counterfactuals in Section 3.5.

Trading Frictions

To calibrate trade frictions across local markets, we use results reported in recent work by Bergquist & McIntosh (2019) based on survey microdata that provide information on bilateral trade flows between Ugandan markets and origin and destination prices. Consistent with the stylized facts in Section 3.2, Bergquist & McIntosh (2019) estimate additive trade costs as a function of road distances between markets. Using their microdata, we revisit those results in our calibration. Using only bilateral price gaps from market pairs during months in which they observe positive trade flows between the pair, in addition to information on the road distance between the markets from the transportation network database, we estimate the following specification:

$$(t_{odkt}) = (p_{dkt} - p_{okt}) = \alpha + \beta\,(RoadDistance_{od}) + \varepsilon_{odkt}$$

where t indexes survey rounds and the error term ε_{odkt} is clustered at the level of bilateral pairs (od). $RoadDistance_{od}$ is measured in road kilometers traveled along the transportation network. As indicated, we estimate a single function of trade costs with respect to road distances across all goods, so $t_{odg} = t_{od}$. The estimated trade cost for an additional road kilometer traveled between two markets is 1.2 Ugandan shillings (standard error 0.29), which implies about one half a US Dollar cost per kilometer for one ton of shipments. To corroborate the plausibility of this result, we can also use additional survey data from Bergquist and McIntosh (2018) on the fuel cost on a given bilateral route, reported for a fully-loaded lorry truck (with capacity of about 5 tons in total). The point estimate per km of distance traveled between bilateral pairs (replacing price gaps by bilateral fuel costs in the specification above) is 1494 (standard error 122), which implies that fuel costs account for about one quarter of the total trade frictions. If we replace the specification above to be in logs on both left and right-hand sides, the distance elasiticity is 0.24 (standard error 0.008), which is close to existing recent evidence for within-country African trade flows by e.g. Atkin & Donaldson (2015).

To calibrate the trading frictions faced by farmers selling and buying to local markets, we implement a similar strategy, using gaps between selling farmers' farm-gate prices and local market prices. However, unlike the cross-market trade flow data, we do not have exact geo-locations for every household as part of the UNPS database, meaning that we cannot project trade frictions as a function of distance traveled to the local market. Instead of using distances, we project the observed price gaps for selling farmers on a number of socio-demographic characteristics that we observe in both the UNPS panel and 100 percent Census data.

$$t_{imkt} = p_{mkt} - p_{ikt} = \alpha + \beta\left(X'_{it}\right) + \varepsilon_{ikt}$$

where p_{mk} are prices at the local market and p_{ik} are farm-gate prices of households who sell to the local market. Again, we estimate a single function for all goods, so that $t_{img} = t_{im}$. All regressions involving the Ugandan household microdata include appropriate weights using survey weights. The estimated average farmer trade friction to their local markets is about 150 Ugandan shilling per kilogram, which amounts to roughly 25 percent of the average unit value for the lowest-price agricultural crop in our setting. We further discuss the household characteristics included in the vector X'_i as part of part of the next subsection (projection to population).

Parameter Estimation

We proceed using the Ugandan household panel microdata to calibrate ζ, σ, $\alpha_{ink\omega}$ and κ. To estimate the cost share parameters of the production function, $\alpha_{ink\omega}$, we take the median of the cost shares that we observe across households in the UNPS microdata by region of the country and appropriately weighted using sampling weights. Appendix Table 3.9 presents the cost shares observed in production across the 9 major crops and the two technology regimes averaged across Ugandan regions.

To estimate the key supply elasticity, κ, we derive the following estimation equation based on the previous Section 3.3:

$$\log\left(\frac{y_{ik\omega t}}{\prod_{n\in\mathcal{N}} q_{ink\omega t}^{\alpha_{ink\omega}}}\right) = -\left(\frac{1}{\kappa}-1\right)\alpha_{ik\omega}^{land}\log\pi_{ik\omega t} + \delta_{k\omega}\log\gamma + \frac{\delta_{k\omega}}{\kappa}\log B_{ik\omega t}, \tag{3.15}$$

where t is a year subscript capturing the panel nature of the microdata. The left-hand side of equation (3.15) is farmer i's harvest quantities for crop k grown under technology regime ω in survey year t (summed across both seasons) adjusted for the reported units of labor, modern intermediates and land used in production. For crops produced under the traditional technology regime $\omega = 0$, the cost share of modern inputs is equal to zero.

The first term on the right-hand side, $\alpha_{ik\omega}^{land}\log\pi_{ik\omega t}$, are land shares used in producing the harvests on the left-hand side multiplied by the cost of share of land in production. The final two terms capture both average and farmer-specific production shocks over time and across crops and technology regimes. In our regressions, we capture those shocks by including crop-by-year-by-technology regime fixed effects ($\theta_{k\omega t}$) and farmer-by-crop-by-technology regime fixed effects

($\phi_{ik\omega}$). The regression coefficient of interest, $\beta = 1 - \frac{1}{\kappa}$, is thus estimated using changes in land allocations within farmer-by-crop-by-technology cells controlling for average changes by crop-technology pairs across farmers over time. Alternatively, to allow for region-specific shocks across crops over time, we also replace $\theta_{k\omega t}$ with region-by-crop-by-year-by-technology regime fixed effects ($\theta_{rk\omega t}$).

The advantage of writing the estimation equation as in (3.15) is that each term on both the left and right-hand sides are observable to us using the rich Ugandan production microdata. In particular, using changes in land shares as the main regressor of interest on the right-hand side (instead of farm-gate prices with land shares on the left-hand side) provides us with a much more complete dataset for estimation given the somewhat scant nature of the available unit value information in the farm production microdata. Furthermore, while we observe changes in inputs to production by plot and farmer (and thus by crop and technology regimes over time), we do not observe changes of the input factor prices (e.g. locality-specific changes to the prices of modern intermediates).

To estimate κ convincingly, we require plausibly exogenous variation in land allocations ($\log \pi_{ik\omega t}$) 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 the unit cost nature of trade frictions documented in Section 3.2 implies that plausibly exogenous shocks to world market prices across crops k should propagate differentially across local markets in Uganda as a function of distances to the nearest border crossing post. In particular, a relative increase in the price of a crop k should lead to a larger reallocation of land shares toward that crop in closer proximity to the border compared to locations farther away (the percentage change in local producer prices is $\frac{\Delta p_{world}}{p_{world,t0}+bordercost_i}$). This relationship should bind for crops that are actively traded on world markets. Among the 9 main crops we study in Uganda, only coffee falls into this category: the share of exports to production for coffee exceeds 90 percent in all years of our sample, whereas the sum of exports plus imports over domestic production is close to zero (below 4 percent) for the other crops. We thus construct the instrument as the interaction of the log distance to the nearest border crossing for farmer i, a dummy for whether crop k is coffee or other and the log of the relative world price of coffee relative to the other 8 crops.

The $\phi_{ik\omega}$ fixed effects account for differences in productivity across farmers by crop and technology regime. The $\theta_{k\omega t}$ fixed effects account for productivity shocks across crops by technology regime over time (and over time by region when using $\theta_{rk\omega t}$). Note that these fixed effects absorb all but the triple interaction term we use in the IV estimation. The identifying assumption is thus that individual farmer productivity shocks in coffee production relative to other crops are not related to the direction of relative world price changes and distances of market places to the nearest border crossing.

As documented in appendix Figure 3.3, over our sample period 2005-2013 the relative world market price of coffee significantly dropped relative to the other major crops produced in Uganda. All else equal, the land shares used for coffee production should have thus fallen less strongly inland compared to regions closer to the border. Panel A of Table 3.10 documents that this is indeed the case. Panel A presents the first-stage regressions of our IV estimation strategy, with $\log \pi_{ik\omega t}$ on the left-hand side and the IV on the right in addition to the various fixed effects. The negative point estimate on the triple interaction term that is our instrument implies that positive (negative) world price changes in coffee relative to other crops increase (decrease) land allocations to coffee significantly more so in closer proximity to the border compared to inland. This relationship holds both before and after including region-by-crop-by-technology-by-time fixed effects, and when using all years of data (2005, 2009, 2010, 2011 and 2013) or just using long changes 2005-2013.

Panels B and C of Table 3.10 proceed to the OLS and IV estimation of equation (3.15). In Panel B, we report estimation results before adjusting farmer harvests ($y_{ik\omega t}$) by inputs used in production in the denominator of the left-hand side. In Panel C, we then report estimation results where point estimates capture $\beta = 1 - \frac{1}{\kappa}$.

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

Moving on the the kappa estimation in Panel C, we find statistically significant point estimates in the range of 0.45-0.66 that imply κ estimates in the range of 1.8-2.9. Reassuringly, these are close to existing estimates of this parameter reported in Sotelo (2017) ($\kappa = 1.7$). Using the results in Panel C, we pick the low estimate of $\kappa = 1.8$ as our baseline calibration (which is conservative in terms of welfare impacts, and in terms of the difference between local-vs-at-scale effects). We also report estimation results across a range of alternative parameter assumptions in a number of additional robustness checks.

We now turn to the two missing parameter estimates on the demand side of the economy, σ and ζ in Section 3.3. Unfortunately, we cannot use the same identification strategy for the demand elasticity that we have used for the supply side above. The reason is that coffee is close to a pure "cash crop" for exports, and rarely consumed locally in Uganda (on average slightly less than 6 percent of households report any consumption of coffee in the consumption microdata). This is work in progress at the time of this writing, and for the moment we rely on existing estimates of the elasticity of demand across agricultural crops reported in a similar setting using Peruvian data by Sotelo (2017). In particular, we use $\sigma = 2.6$ as our baseline calibration, and report counterfactual results across alternative parameter assumptions in a number of robustness checks.

Finally, to calibrate the demand parameter, ζ, we use the following relationship that holds subject to utility maximization under Stone-Geary:

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

where the left-hand side is the share of household income spent on subsistence, and ξ_{iA} is the observed share spent on total food consumption. We set ζ equal to 0.1, consistent with a share spent on subsistence that is on average about 38 percent across Ugandan households. This calibration of ζ is also consistent with an alternative approach that uses the average share of expenditure, ξ_{iA}, among the richest Ugandan households in our survey data for whom $\frac{P_{iA}\bar{C}_A}{I_i}$ approaches zero (which is close to $\xi_{iA} = 0.1$ among the richest 5 percent of households).

Extrapolation from Survey Data to Population

To conduct a meaningful aggregation exercise of the impact of a policy shock at scale, we need to calibrate the model to the full set of local markets populating Uganda. The challenge we face is that the required household-level information on pre-existing production quantities and expenditure shares across crops and sectors is generally not available in microdata covering the whole population.

Instead, we use the fact that the UNPS –which includes such detailed household-level information for a nationally representative sample of Ugandan households– and the 100 percent sample microdata from the 2002 population census –which provides information on all household locations– share a number of household characteristics that are observed in both datasets.

In particular, we estimate a series of regression equations in the UNPS sample data, with outcomes to be projected to the population on the left-hand side and household and location characterstics observed in both datasets on the right-hand side. For each of these predictions, the commonly observed household and location characteristcs that we project crop production and expenditure shares in each local market of Uganda are as follows: cubic of age of head of household, cubic of education, cubic of latitude and longitude, series of dummies for household asset ownership and the potential yield of a given location in the FAO/GAEZ database.

Using these covariates, the average R-squared that we obtain within the UNPS sample is above 40 percent, providing some reassurance that our extrapolation makes use of highly relevant location and household characteristics. [In work in progress, we plan to improve the the precision of this extrapolation exercise by departing from a simple linear prediction framework and implement less parametric approaches using recent machine learning tools.]

3.5 Counterfactual Analysis and Robustness

Using the model, solution method and calibration described in the previous sections, this section presents the counterfactual analysis. We first present counterfactual results on the main questions we discuss in the introduction. In the final subsection, we present additional robustness checks to both investigate the sensitivity of our findings across a range of alternative parameters, and to validate the structure of the model.

Local Effects vs Scaling Up

To fix ideas, we focus on the effects of a subsidy for modern inputs (chemical fertilizers and hybrid seed varieties) on average household welfare, the distributional implications and the underlying mechanisms. We investigate an intervention that gives a 90 percent cost subsidy for these inputs across all crops. Using the production parameterization of the model, this intervention is akin to a positive productivity shock to producing crop k under modern production technology $\omega = 1$ as follows: $\hat{B}_{ik1} = .1^{-\kappa \frac{\sum_{n \in \mathcal{N}_1} \alpha_{i,n,k,1}}{1 - \sum_{n \in \mathcal{N}_1} \alpha_{i,n,k,1}}}$.

To simplify the exercise, we focus on the local vs scaled effects of this shock to agricultural production, and leave aside for the moment the public finance dimension of the subsidy (e.g. financed by a lump-sum tax), and instead solve for counterfactuals after directly shocking the model with the implicit productivity shock under the modern technology regime across the crops.[12]

We implement this intervention using the calibrated model in two different ways. In the local intervention, we randomly select a nationally representative sample of ten thousand rural households (roughly 0.2% of Ugandan households). To do so, we stratify the random sample across 141 counties and four quartiles of food expenditure shares within those counties. We then draw a random 0.2% of the households from each of these 564 bins in Uganda. In the local intervention, we then shock these ten thousand rural households with the subsidy for modern inputs. For the intervention at scale, we then offer the subsidy to all farming households in the economy. In both counterfactuals, we solve for hat changes in household-level outcomes across all ≈ 5 million Ugandan households. As depicted in Figure 3.4, households are located in roughly 5000 rural parish markets and 70 urban centers. We then compare the changes in economic outcomes among the same ten thousand representative sub-sample of Ugandan households in both the local intervention and the at-scale intervention.

In addition to the "local" vs "at scale" counterfactuals, we also implement a number of additional model-based counterfactuals to answer the main questions we discuss in the introduction. To answer the second question, we increase the national saturation rates of farmers receiving the subsidy from the initial ten thousand households to 100% of rural households in steps of 10% that are randomly chosen in each step subject to the same stratification procedure outlined above.[13] We then track the welfare effect on ten thousand rural households and see how they evolve as the national saturation rate increases to 100%.

To answer the third question, we re-write and calibrate the model to allow for different types of trading frictions in the economy. Instead of our baseline specification in terms of additive unit costs, we then conduct the counterfactual analysis after assuming trading frictions are of the iceberg type (ad valorem), or alternatively after assuming there are no trading frictions within the Ugandan economy, but keeping the costs of border crossing to the rest of the world.

To answer the fourth question, we follow a more standard case and implement the counterfactual analysis at the level of regions instead of households. To this end, we aggregate households, including our ten thousand farmer sample, into 51 Ugandan districts plus 70 urban centers. We then implement the intervention at scale by treating each of the 51 representative rural regional agents with the subsidy and solve for counterfactual hat changes across all 121 regions. We then assign counterfactual welfare changes to the identical sample of ten thousand farmers (based on initial consumption and production choices and regional counterfactual changes in prices and wages), and compare the average and distributional effects among this group of farmers across the two levels of aggregation.

In the following figures and tables, all effects are expressed as hat changes (ratios of outcomes

[12] Note that given the assumption that intermediates are not produced domestically and only imported, this simplification should not omit potentially important GE effects in this setting.

[13] The first step adds 9.98% to the 0.2% already treated in the local intervention.

after the policy shock relative to their baseline levels). To investigate distributional effects, we depict changes in outcomes as a function of years of education of the household head. Since this variable is directly used in the extrapolation of household production and consumption choices from the UNPS survey data to the entire population in Section 3.4, this ensures that heterogeneity in e.g. expenditure shares or cropping choices is well-captured along this dimension.[14]

Q1: Local vs Scaling Up: Average Effect on Household Welfare

Figure 3.5 and Table 3.11 present the average effect on household welfare in the local intervention compared to scaling it up to the national level. Table 3.12 presents additional results on the underlying channels.

Q1: Local vs Scaling Up: Distributional Effect on Household Welfare

Figures 3.6 and 3.7 present the distributional implications of the local intervention compared to scaling up. Table 3.13 and Figures 3.8, 3.9 and 3.10 present additional results on the underlying channels.

Q2: GE Forces as a Function of Scale

Figure 3.11 plots the direction and magnitude of GE forces as a function of the scale of the policy intervention in the Ugandan economy. Figure 3.12 plots the distributional implications as a function of the national saturation rate.

Q3: The Role of Trading Frictions for Aggregation

Figure 3.13 presents the difference in the welfare effects across alternative model assumptions about the nature of trade costs linking markets and agents in the economy. Figure 3.14 presents the difference in the welfare effects as a function of pre-existing export shares.

Q4: The Role of Household Aggregation

Figure 3.15 presents the difference in the welfare effect (local vs at scale) for our baseline household-level modeling compared to modeling representative agents across regions in Uganda.

Robustness and Model Validation

In the final section, we explore the robustness of our baseline results across combinations of alternative parameter ranges. We also assess the validity of the model using both cross-sectional moments in the microdata, that we do not use as inputs in the calibration, and evidence from natural experiments that have affected crops and markets differently over time in Uganda.

[14]Note that household incomes are not observed in both UNPS and the 2002 Census microdata. We also confirm that e.g. asset ownership or calibrated household incomes are monotonically related to years of education.

Results Across Alternative Parameters

Figures 3.16 and 3.17 present the counterfactual results for the intervention at scale under alternative parameter assumptions on the supply side (κ) and the demand side (σ).

Moments Not Used in the Calibration

Our calibration makes use of raw household data on pre-existing expenditure shares in consumption and production quantities across crops. As described in Section 3.3, we use this information to set up and solve the price discovery problem such that the observed consumption and production choices are rationalized given the matrix of bilateral trading costs on the transportation graph.

In the cross-section, we can validate this part of our solution method by comparing model-implied prices and and active trading pairs with sub-samples for which we can observe market prices and trading activity from the trader survey data we have obtain from Bergquist and McIntosh (2018). To do this convincingly, we also exclude the part of the price microdata in this validation exercise that we had used for the trade cost calibration across markets in Section 3.4 (i.e. we exclude the market-by-crop observations for which we observed active trading pairs to other markets in the trader survey sample).

We regress deviations in log market crop prices (relative to crop-by-year fixed effects) from the model-based price discovery solution on those same deviations for the same crops and markets in the trader survey microdata. The model only generates one cross-seciton of market prices, whereas we have multiple rounds (years) of trader survey data. We thus stack the repeated cross-sections in the observed data (hence the crop-by-year fixed effects).

There are several reasons why the relationship between deviations in log model-based crop prices on the left-hand side to deviations in reported market prices on the right could be distorted. The reported price data could be subject to measurement error, unobserved variation in crop quality, as well as temporary variation on the day that information was collected across different market places in Uganda. In addition, parish markets in the model are based on centroids, whereas real-world market places that are assigned to the same parish identifier do not necessarily coincide geographically. All of these factors would imply "noise" from the perspective of the model, and would lead to attenuation bias in the regression coefficient, even if the model were 100 percent true in terms of structure.

With these caveats in mind, Figure 3.18 presents the estimation results. We find that the model-based price discovery appears to be strongly positively related with observed variation in crop prices across markets in Uganda. The regression coefficient is .56 with a t-statistic of 3.5 (standard errors clustered at the level of parish markets). These results provide some reassurance that the price discovery step of the model solution provides the model economy with meaningful variation in market prices across crops faced by households.

Additional Evidence from Natural Experiments

We are also working on validating our theoretical framework by comparing the effects of shocks observed in the microdata over time in Uganda, compared to the effects of those same shocks after solving for the counterfactuals in our model.

Road Building First, we study road additions to Uganda's road network over our sample period (we have survey microdata between 2005 and 2013). We obtain data on Uganda's road network, and how it has evolved over time from recent work by Jedwab and Storeygard (2018) who kindly agreed to share their Ugandan database with us. Between 2000-2013 there has been one added road segment, connecting the cities of Lira and Kachung in the center of Uganda. Figure 3.19 depicts this empirical setting. The new road segment was completed and open to traffic in 2003, roughly 1.5 years before our first wave of household microdata from the UNPS in 2005.

We use this instance as a case study for model validation. To investigate the effect of this road addition on household outcomes in the microdata, we test for the effect of the new road on changes in household outcomes in our first two waves in the UNPS: from 2005 to 2009. To do this, we run regressions of the following form:

$$\log\left(y_{ik\omega t}\right) = \beta \log MarketAccess_{ot} + \phi_{ik\omega} + \theta_{k\omega t} + \varepsilon_{ik\omega t}, \tag{3.16}$$

where the left-hand side are log outcomes for farmer i residing in parish market o, producing crop k with technology ω in year t. $\phi_{ik\omega}$ are household-by-crop-by-technology fixed effects and $\theta_{k\omega t}$ are crop-by-technolgy-by-time fixed effects. Alternatively, we also replace $\theta_{k\omega t}$ with crop-by-district-by-technology-by-time fixed effects ($\theta_{km\omega t}$), further restricting the identifying variation to only local comparisons. The explanatory variable $\log MarketAccess_{ot}$ is affected by new road additions as follows:

$$MarketAccess_{ot} = \sum_{d} \left(\frac{Population_{d}^{2002}}{TravelTime_{odt}^{\gamma}} \right).$$

We compute the travel time-weighed access to other markets in Uganda for each parish market o using the road network before and after the new road addition depicted in Figure Figure 3.19. To estimate travel times on the road network, we use information on average speed per km attached to different types of roads in Uganda (using 75 km/h for major paved roads, 60 km/h for other paved roads and 20 km/h for unpaved dirt roads). We follow recent work by Jedwab and Storeygard (2018) on African road market access by setting the elasticity of trading costs with respect to travel time to $\gamma = 3.8$. Note that this common formulation of regional market access (e.g. Donaldson and Hornbeck, 2016) does not have a structural interpretation through the lens of our model. Our model does not feature "structural gravity" (Head and Mayer, 2013), so the estimation results from specification (3.16) are purely reduced-form. Instead, we set out to compare the effect of changes in this measure of market access on local farmer outcomes in the model simulation and the actual microdata for model validation.

To do this convincingly, it is important that variation in $\log MarketAccess_{ot}$ is causally identified when estimating the reduced-form results in the microdata. The identifying assumption in (3.16) is that locations more or less affected by the new road addition did not experience other shocks over the period 2005-09 that affected household outcomes differently (given the inclusion of household fixed effects and conditional on crop-by-technology-by-time fixed effects). The main concern is that the road investment may have been targeted at particular locations that were expected to grow more or less quickly over this period. Though the pervasiveness of long delays between planning and construction periods for public investments in Uganda may limit this concern, we further address this issue in two ways when estimating (3.16) above. First, we exclude household observations from the targeted urban nodes of the new road segment (in fact we exclude all urban households in the estimation, given our interest in agricultural outcomes). Second, we exclude rural markets that directly traversed on the way between the urban nodes in the estimation. We thus estimate the effect of changes in road market access among rural communities that were more or less affected by the new road, but excluding the potentially directly targeted locations. [15]

Using this strategy, Table 3.14 presents the reduced-form estimation results. We focus on two main outcomes in the microdata: changes in total household per capita consumption expenditures and changes in land allocations across different crops. When having outcomes with variation only at the household-by-time level ($\log y_{it}$) with expenditures, we include household and time fixed effects (ϕ_i and θ_t), or alternatively household and district-by-time fixed effects (θ_{mt}). When using land allocation with variation by household, crop, technology and time, we use the fixed effects specifications as written and described in (3.16).

We find that increases in market access due to the new road (opened in 2003) have a positive and significant effect on changes in household total outlays over the following years (2005-2009). The estimated elasticity is about 0.5 with household and time fixed effects, and increases to roughly 0.75 when restricting the variation to local comparisons with additional district-by-time fixed effects. As noted above, these effects do not have a structural interpretation through the lens of our model. On land allocations, we find a pattern of both positive and negative statistically significant effects. The effect appears to be negative for land allocations to beans and sweet potatoes and positive for millet in this setting.

To compare these reduced-form results with changes implied by the calibrated model, we feed in the percentage reduction in bilateral trade costs across market places, Δt_{od} implied by the percentage reductions in bilateral travel times on the road network. At the time of this draft, solving for this counterfactual and running back-to-back regressions with model-simulated vs actual microdata are work in progress.

Weather Shocks The second natural experiment we are working on for model validation is based on temperature and precipitation shocks across regions and time over the sample period 2005-2013. Following recent work by e.g. Aufhammer et al. (2013) we record the total amount of rainfall and count the number of days with a maximum daily temperature at 29 degrees Celsius or above across all growing seasons per year for each parish. To do so, we use daily rainfall and temperature information at the level of 0.1 degree grids for Sub-Saharan Africa from the Famine Early Warning Systems Network (FEWS NET) Land Data Assimilation System (FLDAS) (McNally et al., 2017).

[15] This is akin the identificaiton strategy in e.g. Donaldson and Hornbeck (2016) who use changes in market access due to railway additions made outside the direct vicinity of the origin location.

Using these two variables ($RainFall_{mt}$ and $HotDays_{mt}$), we run regressions of the following form:

$$\log\left(y_{ik\omega t}\right) = \beta_1 HotDays_{mt} + \beta_2 \log RainFall_{mt} + \phi_{ik\omega} + \theta_{k\omega t} + \varepsilon_{ik\omega t}, \tag{3.17}$$

where notation and fixed effects are as discussed above. We use these regression to estimate reduced-form effects of weather shocks on farmer-level (e.g. total expenditures) and farmer-crop-technology-level outcomes (e.g. harvests) in the data over this period.

Table 3.15 presents estimation results for specification (3.17). Using farmer-level log consumption expenditures per capita as the outcome on the left-hand side, we find that larger amounts of precipitation during growing seasons significantly increase total consumption expenditures. We find no significant effect of a larger number of days above 29 degrees Celsius on average in this setting. When using farmer-by-technology-by-crop level observations, with log harvests or log land allocations on the left-hand side, we find significant and heterogenous effects of precipitation and temperature variation across different classes of crops: cereals (maize, millet and sorghum), starchy tubers (cassava and sweet potato) and other (banana, beans, coffee, groundnuts).

In order to be able to compare these observed average effects of weather shocks in Uganda to the predicted changes in our calibrated model, we also need to know what the underlying productivity changes, $\hat{B}_{ik\omega}$, of $HotDays_{mt}$ and $RainDays_{mt}$ are (so that we can solve for counterfactuals after feeding those shocks into the model). To quantify these shocks across crops, technology regimes and regions, we also estimate (3.17) with harvests on the left-hand side and restricting the estimation sample to observations from plots for which we know that the land allocation and other inputs have not changed over time. For this, we make use of the UNPS microdata that report harvests, land allocations and input usage across all plots that a farmer owns (see Section 3.2). For example, if a farmer i plants 100 percent of given plot with crop k under technology ω, and this did not change between two survey rounds, then the observed effect of a weather shock on yields can be used to identify the underlying effect on $\hat{B}_{ik\omega}$. At the time of this writing, this is work in progress.

3.6 Conclusion

Policy interventions aimed at increasing agricultural productivity in developing countries have been a centerpiece in the global fight against poverty. Much of the recent evidence in this space has been based on field experiments and natural experiments, with some well-known limitations that variation from local shocks often do not speak to the GE implications that would unfold if the policy were to be scaled up to the regional or national levels.

In this paper, we develop a rich but tractable quantitative GE model of farmer-level production and trading. To capture a number of salient features that we document in the agricultural empirical context, the model departs from the workhorse "gravity" structure in international trade and economic geography in several dimensions. We propose a new solution approach that allows us to study GE counterfactuals in this rich environment. We then bring to bear administrative microdata on household locations, production, consumption and the transportation network within and across local markets to calibrate the model to the roughly 5 million households populating Uganda in 2002.

In a first set of preliminary results, we find that the average effect of a subsidy for modern fertilizer on rural household real incomes can differ substantially when implemented at scale compared to results from a local intervention that leaves output and factor prices largely unaffected. We show that this difference extends to the policy's distributional implications, which are strongly regressive according to results from the local intervention, but much less so when implemented at scale. We also find that the direction and magnitude of the GE forces appear to be non-linear and non-monotonic as a function of changes in the global saturation rate (moving from 0 to 100 percent of the rural Ugandan population). We also investigate the importance of modeling realistic trading frictions between agents in the economy and the role of household aggregation when solving for the average and distributional implications of shocks at the aggregate level.

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

3.7 Appendix

Model and Solution Method

We first present the excess demand functions $\chi_{j,g}(\bullet)$ used in the text to define the equilibrium, and then we present the excess demand functions for the price discovery step. In the final part, we develop the proof for uniqueness in the price discovery (work in progress).

Excess Demand Functions

The excess demand function for farmers $\chi_{i,g}(\bullet)$ are given by

$$\begin{aligned}
&\chi_{i,k}\left(\left\{a_{i,k'}p_{i,k'}\right\}_{k'},\left\{v_{i,k',\omega'}\right\}_{k',\omega'},I_i\right)\\
&=\varphi_{i,k}\left(\left\{a_{i,k'}p_{i,k'}\right\}_{k'},I_i\right)I_i-\tilde{\psi}_{i,k}\left(\left\{v_{i,k',\omega'}\right\}_{k',\omega'}\right)Y_i\left(\left\{v_{i,k',\omega'}\right\}_{k',\omega'}\right) && \forall k\in\mathscr{K}_A,\\
&\chi_{i,k}\left(\left\{a_{i,k'}p_{i,k'}\right\}_{k'},\left\{v_{i,k',\omega'}\right\}_{k',\omega'},I_i\right)\\
&=\varphi_{i,k}\left(\left\{a_{i,k'}p_{i,k'}\right\}_{k'},I_i\right)I_i && \forall k\in\{M\},\\
&\chi_{i,n}\left(\left\{a_{i,k'}p_{i,k'}\right\}_{k'},\left\{v_{i,k',\omega'}\right\}_{k',\omega'},I_i\right)\\
&=\phi_{i,n}\left(\left\{v_{i,k',\omega'}\right\}_{k',\omega'}\right)Y_i\left(\left\{v_{i,k',\omega'}\right\}_{k',\omega'}\right) && \forall n\in\mathscr{N}_I,\\
&\chi_{i,n}\left(\left\{a_{i,k'}p_{i,k'}\right\}_{k'},\left\{v_{i,k',\omega'}\right\}_{k',\omega'},I_i\right)\\
&=\left[1+\phi_{i,n}\left(\left\{v_{i,k',\omega'}\right\}_{k',\omega'}\right)\right]Y_i\left(\left\{v_{i,k',\omega'}\right\}_{k',\omega'}\right)-I_i && \forall n\in\{L\},
\end{aligned}$$

where $I_i=Y_i\left(\left\{v_{i,k',\omega'}\right\}_{k',\omega'}\right)+p_{i,L}L_i$ and where $\psi_{i,k,\omega}\left(\left\{v_{i,k',\omega'}\right\}_{k',\omega'}\right)\equiv\sum_\omega\frac{\psi_{i,k,\omega}\left(\left\{v_{i,k',\omega'}\right\}_{k',\omega'}\right)}{1-\sum_n\alpha_{i,n,k,\omega}}$.

Similarly, the excess demand for urban households $\chi_{h,g}(\bullet)$ are given by

$$\begin{aligned}
&\chi_{h,k}\left(\left\{a_{h,k'}p_{h,k'}\right\}_{k'},\left\{v_{h,k',\omega'}\right\}_{k',\omega'},I_h\right)=\varphi_{h,k}\left(\left\{a_{h,k'}p_{h,k'}\right\}_{k'},I_h\right)I_h && \forall k\in\mathscr{K}_A,\\
&\chi_{h,k}\left(\left\{a_{h,k'}p_{h,k'}\right\}_{k'},\left\{v_{h,k',\omega'}\right\}_{k',\omega'},I_h\right)=\left[\varphi_{h,k}\left(\left\{a_{h,k'}p_{h,k'}\right\}_{k'},I_h\right)-1\right]I_h && \forall k\in\{M\},\\
&\chi_{h,k}\left(\left\{a_{h,k'}p_{h,k'}\right\}_{k'},\left\{v_{h,k',\omega'}\right\}_{k',\omega'},I_h\right)=0 && \forall n\in\mathscr{N}_I.
\end{aligned}$$

where $I_h = b_{h,M} p_{h,M} L_h$.[16],[17] Finally, for Foreign we have

$$\chi_{F,g} = \begin{cases} -\infty & \text{if} \quad p_{F,g} < p^*_{F,g} \\ ]-\infty,\infty[& \text{if} \quad p_{F,g} = p^*_{F,g} \\ \infty & \text{if} \quad p_{F,g} > p^*_{F,g} \end{cases}.$$

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

$$\begin{aligned}
\chi_{i,k}\left(\left\{p_{i,g}\right\}_g;\mathbb{D}\right) &= \xi_{i,k} I_i - \sum_{\omega} p_{i,k} y_{i,k,\omega} && \forall k \in \mathscr{K}_A, \\
\chi_{i,k}\left(\left\{p_{i,g}\right\}_g;\mathbb{D}\right) &= \xi_{i,k} I_i && \forall k \in \{M\}, \\
\chi_{i,n}\left(\left\{p_{i,g}\right\}_g;\mathbb{D}\right) &= \sum_{k,\omega} \alpha_{i,n,k,\omega} p_{i,k} y_{i,k,\omega} && \forall n \in \mathscr{N}_I, \\
\chi_{i,n}\left(\left\{p_{i,g}\right\}_g;\mathbb{D}\right) &= \sum_{k,\omega} \alpha_{i,n,k,\omega} p_{i,k} y_{i,k,\omega} - p_{i,n} L_i && \forall n \in \{L\}, \\
\chi_{h,k}\left(\left\{p_{h,g}\right\}_g;\mathbb{D}\right) &= \xi_{h,k} I_h && \forall k \in \mathscr{K}_A, \\
\chi_{h,k}\left(\left\{p_{h,g}\right\}_g;\mathbb{D}\right) &= \left(\xi_{h,k} - 1\right) I_h && \forall k \in \{M\}, \\
\chi_{h,n}\left(\left\{p_{h,g}\right\}_g;\mathbb{D}\right) &= 0 && \forall n \in \mathscr{N}_I, \\
\chi_{F,g}\left(\left\{p_{j,g}\right\}_g;\mathbb{D}\right) &= \begin{cases} -\infty & \text{if} \quad p_{F,g} < p^*_{F,g} \\ ]-\infty,\infty[& \text{if} \quad p_{F,g} = p^*_{F,g} \\ \infty & \text{if} \quad p_{F,g} > p^*_{F,g} \end{cases} && \forall g,
\end{aligned}$$

where

$$\begin{aligned}
I_i\left(\left\{p_{i,g}\right\}_g;\mathbb{D}\right) &= \sum_{k,\omega}\left(1 - \sum_n \alpha_{i,n,k,\omega}\right) p_{i,k} y_{i,k,\omega} + p_{i,L} L_i, \\
I_h\left(\left\{p_{h,g}\right\}_g;\mathbb{D}\right) &= p_{h,M} y_{h,M}.
\end{aligned}$$

Price Discovery

In this subsection we show that, in the case with only iceberg trade costs (i.e., $t_{od,g} = 0$ for all o,d,g), no inputs and no trade with Foreign, the price discovery step described in the previous

[16]We include $\{v_{h,k',\omega'}\}_{k',\omega'}$ as an argument in $\chi_{h,k}(\bullet)$ so that $\chi_{j,g}\left(\left\{a_{j,k} p_{j,k}\right\}_k, \left\{v_{j,k,\omega}\right\}_{k,\omega}, I_j\right)$ capture urban households – since the function does not depend on these arguments, there is no need to define them.

[17]In parallel to our treatment of land for farmers, we assume that there is no market for household labor in urban areas, and hence the equilibrium system does not have to determine the price of this good.

section is well defined in the sense that there is a unique set of prices $\{p_{j,g}, p_{m,g}\}$ that solves the system of equations 12-14. 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

$$\mathscr{V} \equiv \{(o,g) \in \mathscr{J} \times \mathscr{K} \mid y_{o,g} > 0\}.$$

A variety of good g produced by agent o is indexed by $(o,g) \in \mathscr{J} \times \mathscr{K}$. Agent $o's$ endowment of (o,g) is $\tilde{y}_{o,g}$. Naturally, no other agent $o' \neq o$ has a positive endowment of (o,g) and so $y_{o,g}$ is also the total endowment of variety (o,g) in the economy.

Next, we let $\tilde{\tau}_{od,g}$ be the minimum cost at which variety (o,g) can be transported from its origin to destination d.[18] Letting $p_{o,g}$ denote the price of variety $(o,g) \in \mathscr{V}$, we know that in a competitive equilibrium the price of a variety must be equal to its cost, and hence the price at which agent d has access to variety (o,g) is $\tilde{\tau}_{od,g} p_{o,g}$.

We let $\xi_{d,g} \in [0,1]$ denote the expenditure share of gross income of agent d (i.e., $\sum_{g \in \mathscr{K}} p_{d,g} y_{d,g}$) on good g. The excess demand function (in quantities) for a variety $(o,g) \in \mathscr{V}$ is given by

$$\bar{\chi}_{o,g}(p) = \sum_d \bar{\chi}_{d,o,g}(p) - y_{o,g},$$

where $p \equiv \{p_{o,g}\}_{(o,g) \in \mathscr{V}}$ and $\bar{\chi}_{d,o,g}(\bullet)$ is the demand function of agent d for variety (o,g) and is given by

$$\bar{\chi}_{d,o,g}(p) \in \begin{cases} \left[0, \frac{\xi_{d,g}}{p_{o,g} \tilde{\tau}_{od,g}} I_d\right] & \text{if } o \in \arg\min_{o' \in \mathscr{J}} p_{o',g} \tilde{\tau}_{o'd,g} \\ \{0\} & \text{if } o \notin \arg\min_{o' \in \mathscr{J}} p_{o',g} \tilde{\tau}_{o'd,g} \end{cases},$$

$$I_d = \sum_g p_{d,g} y_{d,g}.$$

In what follows, we follow the convention that $y_{o,g} = 0 \implies p_{o,g} = \infty$.

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

$$\bar{\chi}_{o,g}(p) = 0 \forall (o,g) \in \mathscr{V}. \tag{3.18}$$

Assumption 1. *(Endowments and demand.)*

1. $\sum_{g \in \mathscr{K}} y_{o,g} > 0 \forall o \in \mathscr{J}$.

---

[18] Formally,

$$\tilde{\tau}_{od,g} \equiv \min_{\phi \in \Phi_{od,g}} \prod_{(o',d') \in \phi} \tau_{o'd',g},$$

where $\Phi_{od,g}$ is the set of directed walks from o to d for good g and each element $\phi \in \Phi_{od,g}$ is a sequence of edges $\{(o_n, d_n)\}_{n=1}^N$ such that $o_1 = o$, $d_N = d$, and $o_{n+1} = d_n$.

2. $y_{d,g} > 0 \implies \xi_{d,g} > 0 \forall d \in \mathscr{J}, g \in \mathscr{K}$.

Definition 1. *A price vector p is strongly connected if there is no partition $\{\mathscr{J}_0, \mathscr{J}_1\}$ of $\mathscr{J}$ such that for all $g \in \mathscr{K}$, $\bar{\chi}_{d,o,g}(p) = \bar{\chi}_{o,d,g}(p) = 0 \forall o \in \mathscr{J}_0, d \in \mathscr{J}_1$.*

Proposition 3. *Given Assumption 1, there can be at most one strongly connected price vector p – up to the choice of numeraire – that solves the system of equations 3.18.*

Assume by contradiction that there are two strongly connected price vectors $p \neq p'$ such that both solve Equation 3.18 with some price in p' and p being the same (to rule out the case in which $p' = \kappa p$ for some positive κ). Since $\bar{\chi}_{o,g}(p)$ is homogeneous of degree zero, we can assume without loss of generality that there is a partition $\{\mathscr{M}_0, \mathscr{M}_1\}$ of $\mathscr{V}$ such that

$$
\begin{aligned}
p'_{o,g} &= p_{o,g} \forall (o,g) \in \mathscr{M}_0, \\
p'_{o,g} &> p_{o,g} \forall (o,g) \in \mathscr{M}_1,
\end{aligned}
$$

where $\mathscr{M}_0 \neq \emptyset$ and $\mathscr{M}_1 \neq \emptyset$. Focusing on such prices (p, p'), we show a contradiction by way of five claims.

Before stating the claims, we introduce some additional definitions and notation.

Definition 2. *Given (p, p'), consider the set of partitions $\left\{\mathscr{O}^*_{-1,g}, \mathscr{O}^*_{0,g}, \mathscr{O}^*_{1,g}\right\}_{g \in \mathscr{K}}$ of $\mathscr{J}$ such that*

$$
\begin{aligned}
y_{o,g} &= 0 \forall o \in \mathscr{O}^*_{-1,g}, \\
p'_{o,g} &= p_{o,g} \forall o \in \mathscr{O}^*_{0,g}, \\
p'_{o,g} &> p_{o,g} \forall o \in \mathscr{O}^*_{1,g},
\end{aligned}
$$

the set of partitions $\left\{\mathscr{G}_{-1,d}, \mathscr{G}_{0,d}, \mathscr{G}_{1,d}\right\}_{d \in \mathscr{J}}$ of $\mathscr{K}$ such that

$$
\begin{aligned}
\xi_{d,g} &= 0 \forall g \in \mathscr{G}_{-1,d}, \\
\xi_{d,g} > 0 \wedge \left\{\arg\min_{o \in \mathscr{J}} p_{o,g} \tilde{\tau}_{od,g}\right\} \cap \mathscr{O}^*_{0,g} &\neq \emptyset \forall g \in \mathscr{G}_{0,d}, \\
\xi_{d,g} > 0 \wedge \left\{\arg\min_{o \in \mathscr{J}} p_{o,g} \tilde{\tau}_{od,g}\right\} \cap \mathscr{O}^*_{0,g} &= \emptyset \forall g \in \mathscr{G}_{1,d},
\end{aligned}
$$

the set of partitions $\left\{\mathscr{D}^*_{-1,g}, \mathscr{D}^*_{0,g}, \mathscr{D}^*_{1,g}\right\}_{g \in \mathscr{K}}$ of $\mathscr{J}$ such that

$$
\begin{aligned}
g &\in \mathscr{G}_{-1,d} \forall d \in \mathscr{D}^*_{-1,g}, \\
g &\in \mathscr{G}_{0,d} \forall d \in \mathscr{D}^*_{0,g}, \\
g &\in \mathscr{G}_{1,d} \forall d \in \mathscr{D}^*_{1,g},
\end{aligned}
$$

and finally the partition $\{\mathscr{J}_-, \mathscr{J}_+, \mathscr{J}_\pm\}$ *of* $\mathscr{J}$ *such that*

$$\begin{aligned}\mathscr{G}_{1,d} &= \emptyset \forall d \in \mathscr{J}_-,\\ \mathscr{G}_{0,d} &= \emptyset \forall d \in \mathscr{J}_+,\\ \mathscr{G}_{0,d} \neq \emptyset \wedge \mathscr{G}_{1,d} &\neq \emptyset \forall d \in \mathscr{J}_\pm.\end{aligned}$$

In words, $\mathscr{O}^*_{-1,g}$ is the set of agents who have a zero endowment of good g while $\mathscr{O}^*_{0,g}$ ($\mathscr{O}^*_{1,g}$) is the set of agents o for which their variety of good g does not change (increases in) price; $\mathscr{G}_{-1,d}$ is the set of goods for which expenditure is zero for agent d while $\mathscr{G}_{0,d}$ ($\mathscr{G}_{1,d}$) is the set of goods whose price (inclusive of trade cost) does not increase (increases) for agent d; $\mathscr{D}^*_{-1,g}$ is the set of agents who have zero expenditure on good g while $\mathscr{D}^*_{0,g}$ ($\mathscr{D}^*_{1,g}$) is the set of agents who do not (do) face an increase in price of good g; and $\mathscr{J}_-$ ($\mathscr{J}_+$) is the set of agents for whom the price of all goods remains the same (increases) and $\mathscr{J}_\pm$ is the set of agents that for whom some goods increase in price while others do not. Note that we could have equivalently defined $\mathscr{J}_- \equiv \cap_g \mathscr{D}^*_{0,g}$ and $\mathscr{J}_+ \equiv \cap_g \mathscr{D}^*_{1,g}$,with $\mathscr{J}_\pm \equiv \mathscr{J} \setminus (\mathscr{J}_- \cup \mathscr{J}_+)$.

Claim 1. *For all* $g \in \mathscr{K}$*,*

$$\begin{aligned}\bar{\chi}_{d,o,g}(p) &= 0 \forall o \in \mathscr{J}_-, d \in \mathscr{J}_+,\\ \bar{\chi}_{d,o,g}(p) &= 0 \forall o \in \mathscr{J}_+, d \in \mathscr{J}_-.\end{aligned}$$

Proof. For any $d \in \mathscr{J}_+$, $\mathscr{G}_{0,d} = \emptyset$. Hence, for all $d \in \mathscr{J}_+$ and $g \in \mathscr{K} \setminus \mathscr{G}_{-1,d}$, it must be the case that $\exists o \in \mathscr{O}^*_{1,g}$ such that $\bar{\chi}_{d,o,g}(p) > 0$ while $\bar{\chi}_{d,o,g}(p) = 0 \forall o \notin \mathscr{O}^*_{1,g}$. Further, it follows from the definition of $\mathscr{J}_-$ and the fact that in equilibrium an agent always consumes a positive amount of its own varieties (from Assumption 1, $y_{d,g} > 0 \implies \xi_{d,g} > 0$) that $\mathscr{J}_- \cap \left(\cup_{g \in \mathscr{K}} \mathscr{O}^*_{1,g}\right) = \emptyset$. This implies that agents in $\mathscr{J}_+$ purchase no positive value from agents in $\mathscr{J}_-$ under equilibrium prices p, that is, for all $g \in \mathscr{K}$, $\bar{\chi}_{d,o,g}(p) = 0 \forall o \in \mathscr{J}_-, d \in \mathscr{J}_+$.

Similarly, for any $d \in \mathscr{J}_-$, $\mathscr{G}_{1,d} = \emptyset$. Hence, for all $d \in \mathscr{J}_-$ and $g \in \mathscr{K} \setminus \mathscr{G}_{-1,d}$, it must be the case that $\exists o \in \mathscr{O}^*_{0,g}$ such that $\bar{\chi}_{d,o,g}(p) > 0$ while $\bar{\chi}_{d,o,g}(p) = 0 \forall o \notin \mathscr{O}^*_{0,g}$. Since $\mathscr{J}_+ \cap \left(\cup_{g \in \mathscr{K}} \mathscr{O}^*_{0,g}\right) = \emptyset$, this implies that agents in $\mathscr{J}_-$ purchase no positive value from agents in $\mathscr{J}_+$ under equilibrium prices p, that is, for all $g \in \mathscr{K}$, $\bar{\chi}_{d,o,g}(p) = 0 \forall o \in \mathscr{J}_+, d \in \mathscr{J}_-$. □

Claim 2. $\mathscr{J}_\pm \neq \emptyset$.

Proof. Suppose $\mathscr{J}_\pm = \emptyset$. Then $\{\mathscr{J}_-, \mathscr{J}_+\}$ form a partition of $\mathscr{J}$ and then Claim 1 implies that p is not strongly connected, leading to a contradiction. □

Claim 3. $\mathscr{J}_\pm \cap \left(\cup_{g \in \mathscr{K}} \mathscr{O}^*_{1,g}\right) \neq \emptyset$.

In words, there is a least one agent for whom the prices of some goods increase and the prices of others do not while at the same time experiencing an increase in the price of it's variety of at least one good.

Proof. We consider two cases.

1. $\mathscr{J}_+ = \emptyset$.
   Since $\mathscr{M}_1 \neq \emptyset$, then $\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g} \neq \emptyset$, and so $\mathscr{J} \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right) \neq \emptyset$. Then,

$$\begin{aligned} &\mathscr{J} \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right) \neq \emptyset \\ \Longrightarrow\ &\left(\mathscr{J}_- \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right)\right) \cup \left(\mathscr{J}_\pm \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right)\right) \neq \emptyset \\ \Longrightarrow\ &\mathscr{J}_\pm \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right) \neq \emptyset, \end{aligned}$$

   where the second line follows from the fact that, since $\mathscr{J}_+ = \emptyset$, then $\mathscr{J}_- \cup \mathscr{J}_\pm = \mathscr{J}$, and the last line follows from the fact that $\mathscr{J}_- \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right) = \emptyset$.

2. $\mathscr{J}_+ \neq \emptyset$.

   We proceed by contradiction. Suppose $\mathscr{J}_\pm \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right) = \emptyset$. From Claim 2, we know that $\mathscr{J}_\pm \neq \emptyset$. For any $d \in \mathscr{J}_\pm$, $\mathscr{G}_{1,d} \neq \emptyset$. Hence, for all $d \in \mathscr{J}_\pm$, $\exists\, g \in \mathscr{K}$ and $o \in \mathscr{O}^*_{1,g}$ such that $\bar{\chi}_{d,o,g}(p) > 0$. In other words, for a good whose price increased for $d \in \mathscr{J}_\pm$, it must come (in the equilibrium with p) from an agent for whom the price of that variety increased. Further, note that for all $g \in \mathscr{K}$, $\mathscr{O}^*_{1,g} \subseteq \mathscr{D}^*_{1,g}$ (from part 2 of Assumption 1). Since $\mathscr{J}_\pm \cup \mathscr{J}_+ = \cup_{g\in\mathscr{K}} \mathscr{D}^*_{1,g}$ and $\mathscr{J}_\pm \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right) = \emptyset$, it then follows that $\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g} \subseteq \mathscr{J}_+$. This implies that if an agent $o \in \mathscr{O}^*_{1,g}$ for any $g \in \mathscr{K}$, it must also be the case that $o \in \mathscr{J}_+$. Hence, for all agents $d \in \mathscr{J}_\pm$, $\exists g \in \mathscr{K}$ and $o \in \mathscr{J}_+$ such that $\bar{\chi}_{d,o,g}(p) > 0$. This implies that agents in $\mathscr{J}_\pm$ must purchase a positive value from agents in $\mathscr{J}_+$ under equilibrium prices p, that is,

$$\sum_{o\in\mathscr{J}_+} \sum_{d\in\mathscr{J}_\pm} \sum_{g\in\mathscr{K}} p_{o,g} \tilde{\tau}_{od,g} \bar{\chi}_{d,o,g}(p) > 0.$$

   For any $d \in \mathscr{J}_+$, $\mathscr{G}_{0,d} = \emptyset$. Hence, for all $d \in \mathscr{J}_+$ and $g \in \mathscr{K} \backslash \mathscr{G}_{-1,d}$, it must be the case that $\exists o \in \mathscr{O}^*_{1,g}$ such that $\bar{\chi}_{d,o,g}(p) > 0$ and for all $o \notin \mathscr{O}^*_{1,g}$, $\bar{\chi}_{d,o,g}(p) = 0$. Since $\mathscr{J}_\pm \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right) = \emptyset$, this implies that agents in $\mathscr{J}_+$ purchase no positive value from agents in $\mathscr{J}_\pm$ under equilibrium prices p, that is,

$$\sum_{o\in\mathscr{J}_\pm} \sum_{d\in\mathscr{J}_+} \sum_{g\in\mathscr{K}} p_{o,g} \tilde{\tau}_{od,g} \bar{\chi}_{d,o,g}(p) = 0.$$

   From Claim 1, it also follows that

$$\begin{aligned} \sum_{o\in\mathscr{J}_-} \sum_{d\in\mathscr{J}_+} \sum_{g\in\mathscr{K}} p_{o,g} \tilde{\tau}_{od,g} \bar{\chi}_{d,o,g}(p) = 0, \\ \sum_{o\in\mathscr{J}_+} \sum_{d\in\mathscr{J}_-} \sum_{g\in\mathscr{K}} p_{o,g} \tilde{\tau}_{od,g} \bar{\chi}_{d,o,g}(p) = 0. \end{aligned}$$

Since the budget constraint of an agent $o \in \mathscr{J}_+$ is satisfied with equality under equilibrium prices p, we have

$$\sum_{g\in\mathscr{K}}\sum_{d\in\mathscr{J}} p_{o,g}\tilde{\tau}_{od,g}\bar{\chi}_{d,o,g}(p) = \sum_{g\in\mathscr{K}}\sum_{o'\in\mathscr{J}} p_{o',g}\tilde{\tau}_{o'o,g}\bar{\chi}_{o,o',g}(p)\,\forall o \in \mathscr{J}_+$$

and hence

$$\begin{aligned}
&\sum_{g\in\mathscr{K}}\sum_{d\in\mathscr{J}_\pm} p_{o,g}\tilde{\tau}_{od,g}\bar{\chi}_{d,o,g}(p) \\
&+\sum_{g\in\mathscr{K}}\sum_{d\in\mathscr{J}_-} p_{o,g}\tilde{\tau}_{od,g}\bar{\chi}_{d,o,g}(p) \\
&+\sum_{g\in\mathscr{K}}\sum_{d\in\mathscr{J}_+} p_{o,g}\tilde{\tau}_{od,g}\bar{\chi}_{d,o,g}(p) = \sum_{g\in\mathscr{K}}\sum_{o'\in\mathscr{J}_\pm} p_{o',g}\tilde{\tau}_{o'o,g}\bar{\chi}_{o,o',g}(p) \\
&\qquad+\sum_{g\in\mathscr{K}}\sum_{o'\in\mathscr{J}_-} p_{o',g}\tilde{\tau}_{o'o,g}\bar{\chi}_{o,o',g}(p) \\
&\qquad+\sum_{g\in\mathscr{K}}\sum_{o'\in\mathscr{J}_+} p_{o',g}\tilde{\tau}_{o'o,g}\bar{\chi}_{o,o',g}(p)\,\forall o \in \mathscr{J}_+ \\
\implies &\sum_{o\in\mathscr{J}_+}\sum_{g\in\mathscr{K}}\sum_{d\in\mathscr{J}_\pm} p_{o,g}\tilde{\tau}_{od,g}\bar{\chi}_{d,o,g}(p) \\
&+\sum_{o\in\mathscr{J}_+}\sum_{g\in\mathscr{K}}\sum_{d\in\mathscr{J}_-} p_{o,g}\tilde{\tau}_{od,g}\bar{\chi}_{d,o,g}(p) \\
&+\sum_{o\in\mathscr{J}_+}\sum_{g\in\mathscr{K}}\sum_{d\in\mathscr{J}_+} p_{o,g}\tilde{\tau}_{od,g}\bar{\chi}_{d,o,g}(p) = \sum_{d\in\mathscr{J}_+}\sum_{g\in\mathscr{K}}\sum_{o\in\mathscr{J}_\pm} p_{o,g}\tilde{\tau}_{od,g}\bar{\chi}_{d,o,g}(p) \\
&\qquad+\sum_{d\in\mathscr{J}_+}\sum_{g\in\mathscr{K}}\sum_{o\in\mathscr{J}_-} p_{o,g}\tilde{\tau}_{od,g}\bar{\chi}_{d,o,g}(p) \\
&\qquad+\sum_{d\in\mathscr{J}_+}\sum_{g\in\mathscr{K}}\sum_{o\in\mathscr{J}_+} p_{o,g}\tilde{\tau}_{od,g}\bar{\chi}_{d,o,g}(p) \\
\implies &\underbrace{\sum_{o\in\mathscr{J}_+}\sum_{g\in\mathscr{K}}\sum_{d\in\mathscr{J}_\pm} p_{o,g}\tilde{\tau}_{od,g}\bar{\chi}_{d,o,g}(p)}_{>0} \\
&+\underbrace{\sum_{o\in\mathscr{J}_+}\sum_{g\in\mathscr{K}}\sum_{d\in\mathscr{J}_-} p_{o,g}\tilde{\tau}_{od,g}\bar{\chi}_{d,o,g}(p)}_{=0} = \underbrace{\sum_{d\in\mathscr{J}_+}\sum_{g\in\mathscr{K}}\sum_{o\in\mathscr{J}_\pm} p_{o,g}\tilde{\tau}_{od,g}\bar{\chi}_{d,o,g}(p)}_{=0} \\
&\qquad+\underbrace{\sum_{d\in\mathscr{J}_+}\sum_{g\in\mathscr{K}}\sum_{o\in\mathscr{J}_-} p_{o,g}\tilde{\tau}_{od,g}\bar{\chi}_{d,o,g}(p)}_{=0}
\end{aligned}$$

Clearly, we have reached a contradiction, hence it must be the case that $\mathscr{J}_\pm \cap \left(\cup_{g\in\mathscr{K}}\mathscr{O}^*_{1,g}\right) \neq \emptyset$. □

Claim 4. *Let $\mathscr{M}_{0,d} \equiv \left\{(o,g) \in \mathscr{M}_0 \mid g \in \mathscr{G}_{0,d} \wedge o \in \left\{\arg\min_{o'} p_{o'g}\tau_{o'd,g}\right\}\right\}$ be the set of varieties consumed by agent d in the equilibrium with prices p which do not change in price under p'. Then $d \in \mathscr{J}_+ \iff \mathscr{M}_{0,d} = \emptyset$.*

In words, an agent experiences an increase in the price of all goods it consumes if and only if this agent was not consuming any variety in the equilibrium with prices p whose price remained unchanged as we move from p to p'.

Proof. Consider an agent $d \notin \mathscr{J}_+$. Then $\mathscr{G}_{0,d} \neq \emptyset$ and $\forall g \in \mathscr{G}_{0,d}$ d purchased good g from an agent o such that $(o,g) \in \mathscr{M}_0$. We then have $\mathscr{M}_{0,d} \neq \emptyset$. The converse follows trivially. Hence $d \in \mathscr{J}_+ \iff \mathscr{M}_{0,d} = \emptyset$. □

Claim 5. *If $\bar{\chi}_{o,g}(p) = 0 \forall (o,g) \in \mathscr{V}$, then $\exists (o,g) \in \mathscr{V}$ such that $\bar{\chi}_{o,g}(p') > 0$.*

Proof. Note that at p (similarly at p')

$$\bar{\chi}_{d,o,g}(p) = \frac{\xi_{d,g}}{p_{o,g}\tilde{\tau}_{od,g}} I_d, \forall (o,g) \in \mathscr{M}_{0,d},$$
$$I_d = \sum_{g \in \mathscr{K}} p_{d,g} y_{d,g}.$$

Since $\mathscr{M}_{0,d} \subseteq \mathscr{M}_0 \forall d \in \mathscr{J}$, $p'_{o,g} = p_{o,g}$ for all $(o,g) \in \mathscr{M}_{0,d}$. It is then clear that, $\forall (o,g) \in \mathscr{M}_{0,d}$,

$$I'_d > I_d \iff \bar{\chi}_{d,o,g}(p') > \bar{\chi}_{d,o,g}(p),$$
$$I'_d = I_d \iff \bar{\chi}_{d,o,g}(p') = \bar{\chi}_{d,o,g}(p).$$

Further, note that $I'_d > I_d$ if $\exists g \in \mathscr{K}$ such that $d \in \mathscr{O}^*_{1,g}$ and $I'_d = I_d$ if $d \notin \mathscr{O}^*_{1,g}$ $\forall g \in \mathscr{K}$. That is,

$$d \in \cup_{g \in \mathscr{K}} \mathscr{O}^*_{1,g} \implies I'_d > I_d,$$
$$d \notin \cup_{g \in \mathscr{K}} \mathscr{O}^*_{1,g} \implies I'_d = I_d.$$

Consequently, we have $\forall (o,g) \in \mathscr{M}_{0,d}$,

$$d \in \cup_{g \in \mathscr{K}} \mathscr{O}^*_{1,g} \implies \bar{\chi}_{d,o,g}(p') > \bar{\chi}_{d,o,g}(p),$$
$$d \notin \cup_{g \in \mathscr{K}} \mathscr{O}^*_{1,g} \implies \bar{\chi}_{d,o,g}(p') = \bar{\chi}_{d,o,g}(p).$$

Note that for any $d \in \mathscr{J}$ we have $\bar{\chi}_{d,o,g}(p') \geq \bar{\chi}_{d,o,g}(p) = 0 \forall (o,g) \in \mathscr{M}_0 \setminus \mathscr{M}_{0,d}$.

Next, consider an agent $d \notin \mathscr{J}_\pm \cap \left(\cup_{g \in \mathscr{K}} \mathscr{O}^*_{1,g}\right)$. Then, $d \in \mathscr{J}_+ \cup \mathscr{J}_- \cup \left(\mathscr{J}_\pm \setminus \left(\cup_{g \in \mathscr{K}} \mathscr{O}^*_{1,g}\right)\right)$. We have three possibilities:

1. Suppose $d \in \mathscr{J}_+$. From Claim 4, $d \in \mathscr{J}_+ \implies \mathscr{M}_{0,d} = \emptyset$. Therefore, for $d \in \mathscr{J}_+$, $\bar{\chi}_{d,o,g}(p') = \bar{\chi}_{d,o,g}(p) \forall (o,g) \in \mathscr{M}_{0,d}$ vacuously.

2. Suppose $d \in \mathscr{J}_-$. Since $\mathscr{J}_- \cap \left(\cup_{g\in\mathscr{G}} \mathscr{O}^*_{1,g}\right) = \emptyset$, $d \notin \cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}$. Therefore, for $d \in \mathscr{J}_-$, $\bar{\chi}_{d,o,g}(p') = \bar{\chi}_{d,o,g}(p) \forall (o,g) \in \mathscr{M}_{0,d}$.
3. Suppose $d \in \mathscr{J}_\pm \backslash \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right)$. Clearly $d \notin \cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}$.
   Therefore, $\bar{\chi}_{d,o,g}(p') = \bar{\chi}_{d,o,g}(p) \forall (o,g) \in \mathscr{M}_{0,d}, d \in \mathscr{J}_\pm \backslash \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right)$.

Therefore, for $d \notin \mathscr{J}_\pm \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right)$, $\bar{\chi}_{d,o,g}(p') = \bar{\chi}_{d,o,g}(p) \forall (o,g) \in \mathscr{M}_{0,d}$.

Finally, consider an agent $d \in \mathscr{J}_\pm \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right)$ (such an agent exists thanks to Claim 3). Clearly $d \in \cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}$. Therefore, for $d \in \mathscr{J}_\pm \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right)$, $\bar{\chi}_{d,o,g}(p') > \bar{\chi}_{d,o,g}(p) \forall (o,g) \in \mathscr{M}_{0,d}$ ($\mathscr{M}_{0,d} \neq \emptyset$ by Claim 4 since $d \in \mathscr{J}_\pm$ implies $d \notin \mathscr{J}_+$).

Putting these cases together we have

$$\begin{aligned}
&\bar{\chi}_{d,o,g}(p') \geq \bar{\chi}_{d,o,g}(p) \forall (o,g) \in \mathscr{M}_0 \backslash \mathscr{M}_{0,d}, d \in \mathscr{J},\\
&\bar{\chi}_{d,o,g}(p') = \bar{\chi}_{d,o,g}(p) \forall (o,g) \in \mathscr{M}_{0,d}, d \notin \mathscr{J}_\pm \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right),\\
&\bar{\chi}_{d,o,g}(p') > \bar{\chi}_{d,o,g}(p) \forall (o,g) \in \mathscr{M}_{0,d}, d \in \mathscr{J}_\pm \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right).
\end{aligned}$$

Adding across all agents $d \in \mathscr{J}$ we obtain

$$\bar{\chi}_{o,g}(p') > \bar{\chi}_{o,g}(p) \forall (o,g) \in \cup_{d\in \mathscr{J}_\pm \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right)} \mathscr{M}_{0,d}.$$

Since $\bar{\chi}_{o,g}(p) = 0 \forall (o,g) \in \mathscr{V}$, it follows that

$$\bar{\chi}_{o,g}(p') > 0 \forall (o,g) \in \cup_{d\in \mathscr{J}_\pm \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right)} \mathscr{M}_{0,d}.$$

Since $\mathscr{J}_\pm \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right) \neq \emptyset$ (from Claim 3) and $\mathscr{M}_{0,d} \neq \emptyset \forall d \in \mathscr{J}_\pm \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right)$ (from Claim 4) imply $\cup_{d\in \mathscr{J}_\pm \cap \left(\cup_{g\in\mathscr{K}} \mathscr{O}^*_{1,g}\right)} \mathscr{M}_{0,d} \neq \emptyset$, $\exists (o,g) \in \mathscr{V}$ such that $\bar{\chi}_{o,g}(p') > 0$. $\square$

3.8 Figures and Tables

Figure 3.1: Regional Specialization

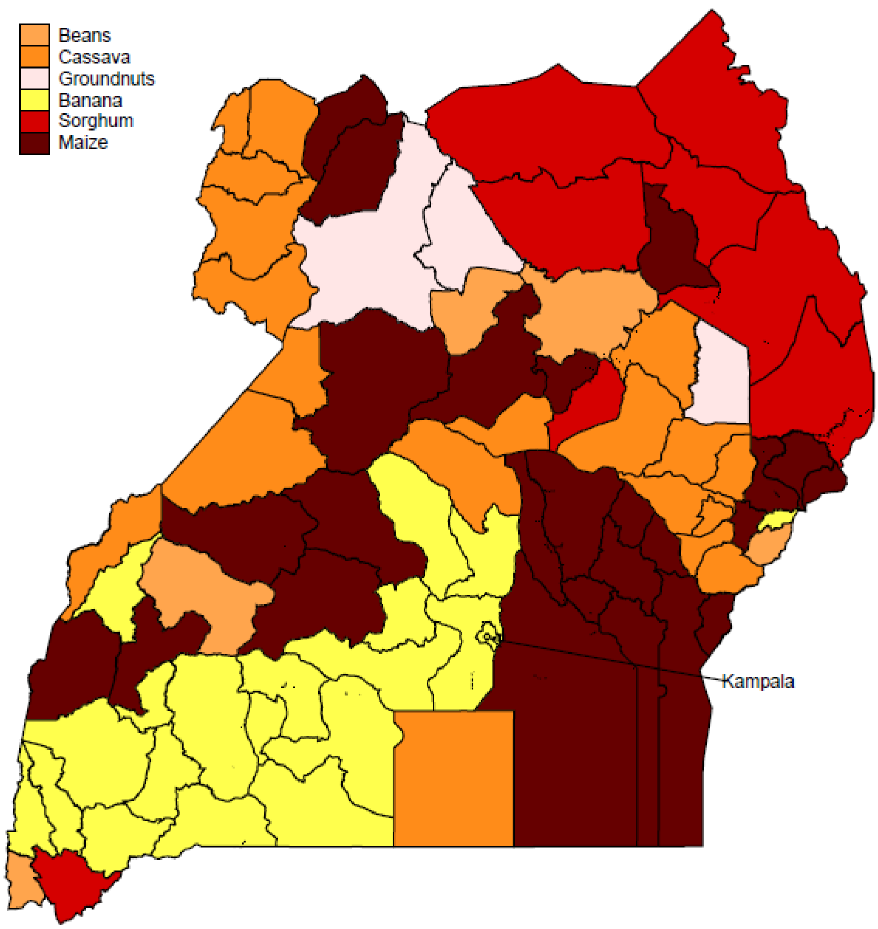

The figure displays the crop with the highest land allocation in each Ugandan district. We use the UNPS data to compute the mean of each crop's land shares across 4 rounds of data. See Section 3.2 for discussion of the data.

Figure 3.2: Household Preferences (Non-Homotheticity)

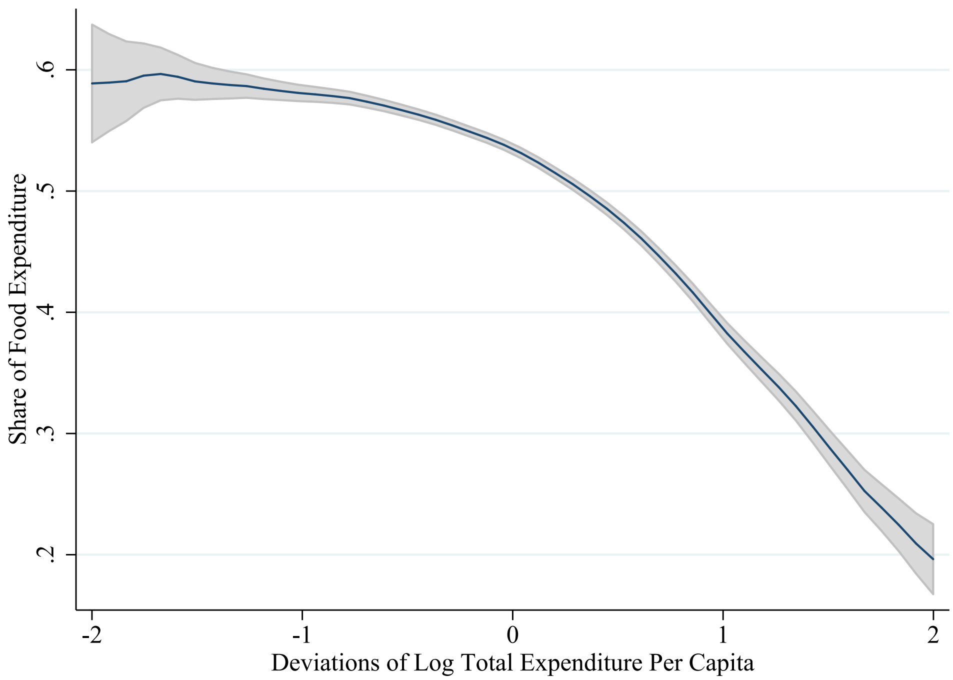

See Section 3.2 for discussion of the data.

Figure 3.3: Relative World Price Changes Over the Sample Period

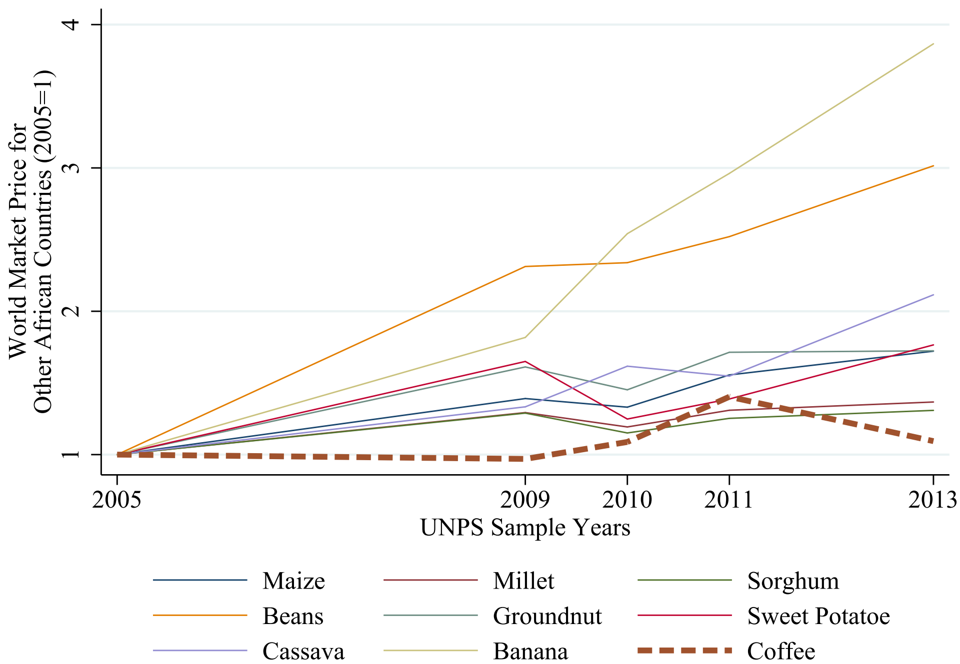

See Section 3.4 for discussion of the data.

Figure 3.4: Ugandan Markets and Transportation Network

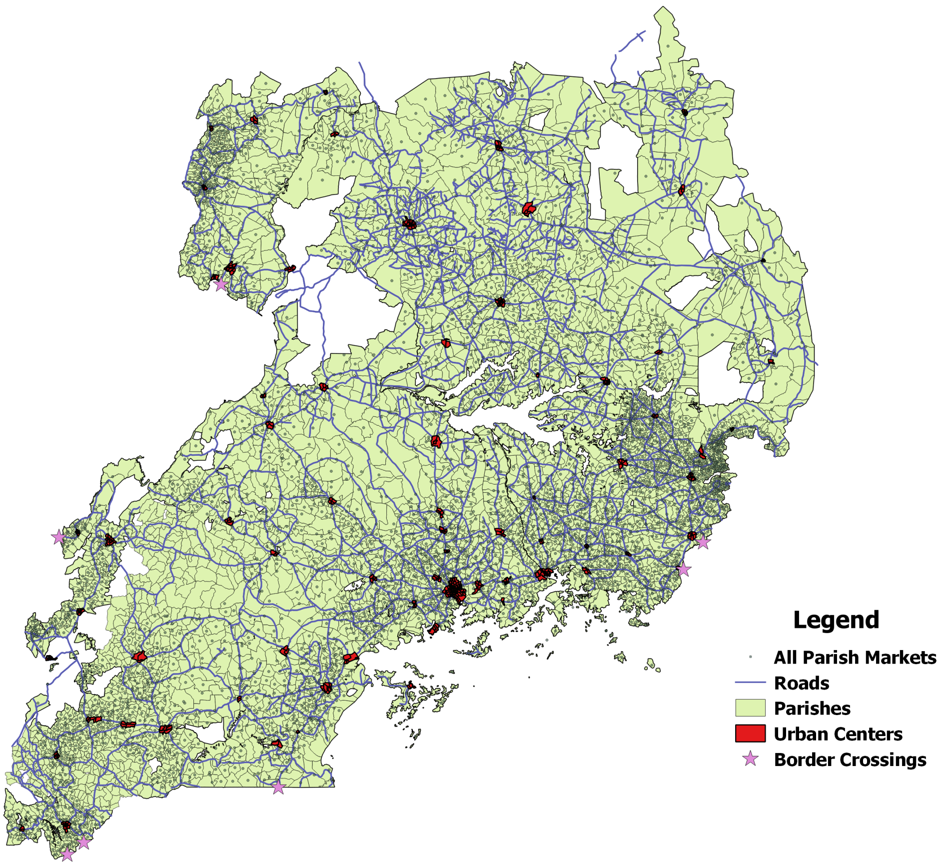

The figure displays the location of local parish markets, urban markets, border crossings and the road network in Uganda. See Section 3.2 for discussion of the data and Section 3.5 for the counterfactual analysis based this geography.

Figure 3.5: Difference in Average Effect of Local Intervention vs At Scale

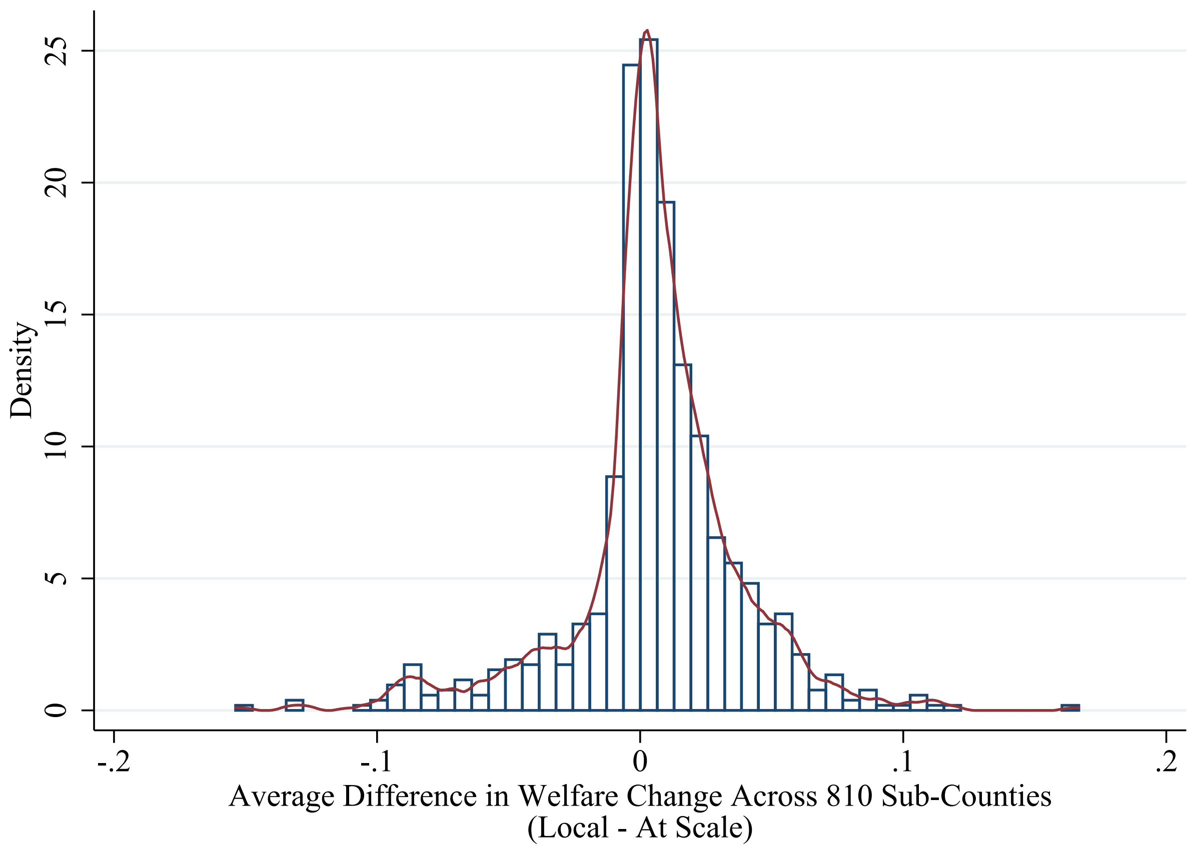

The figure plots household averages for the difference in the welfare change between the local intervention vs at scale across 810 Ugandan sub-counties. These estimates are based on a sample of 10k nationally representative rural households. See Section 3.5 for discussion.

Figure 3.6: Distributional Implications

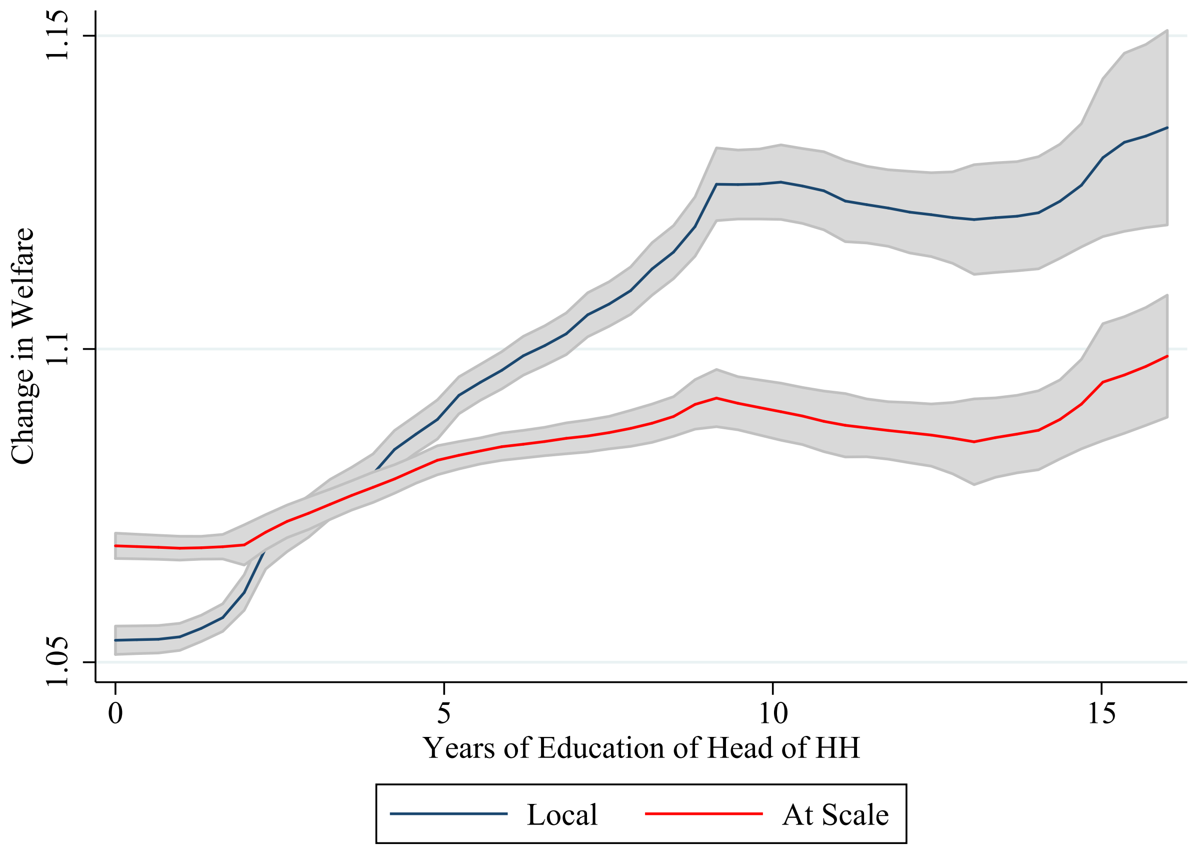

The figure plots estimates from local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 3.5 for discussion.

Figure 3.7: Comparing 10k Farmers to All Rural Households

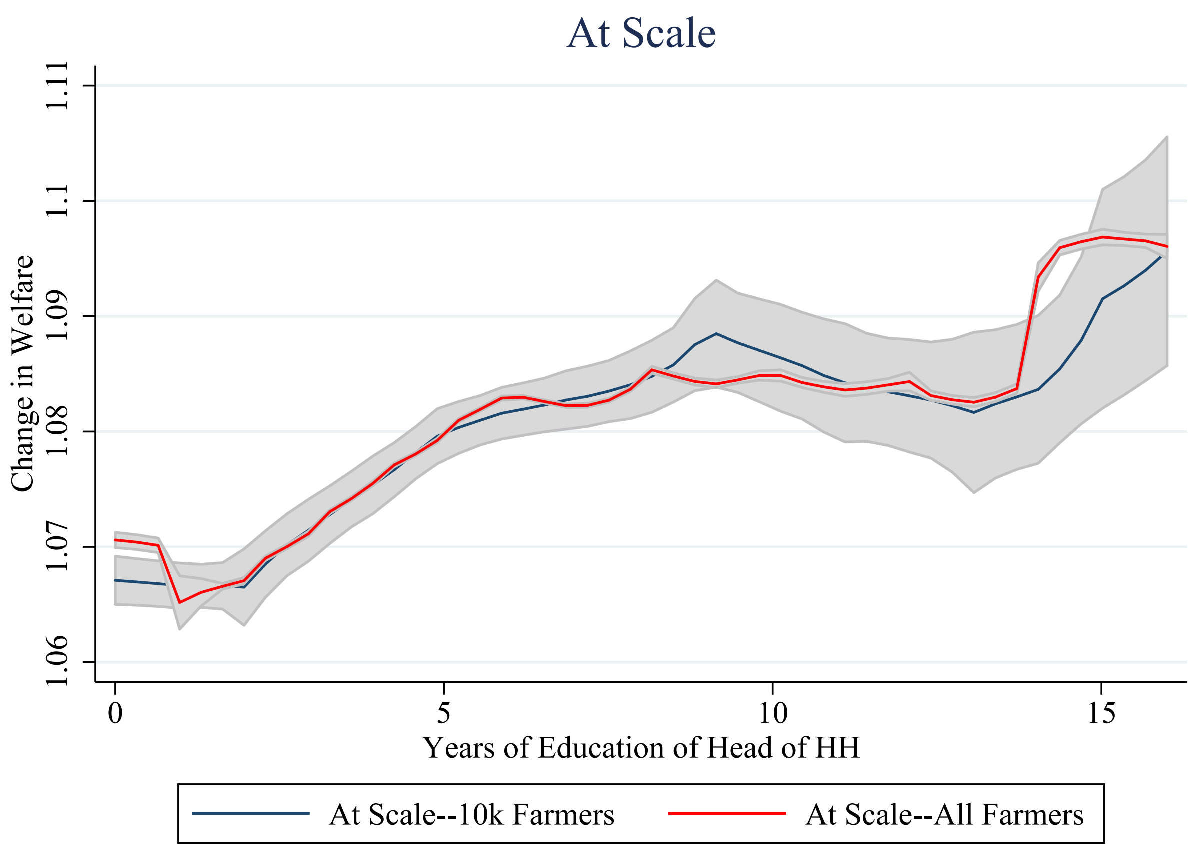

The figure plots estimates from local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 3.5 for discussion.

Figure 3.8: Pre-Existing Use of Modern Inputs

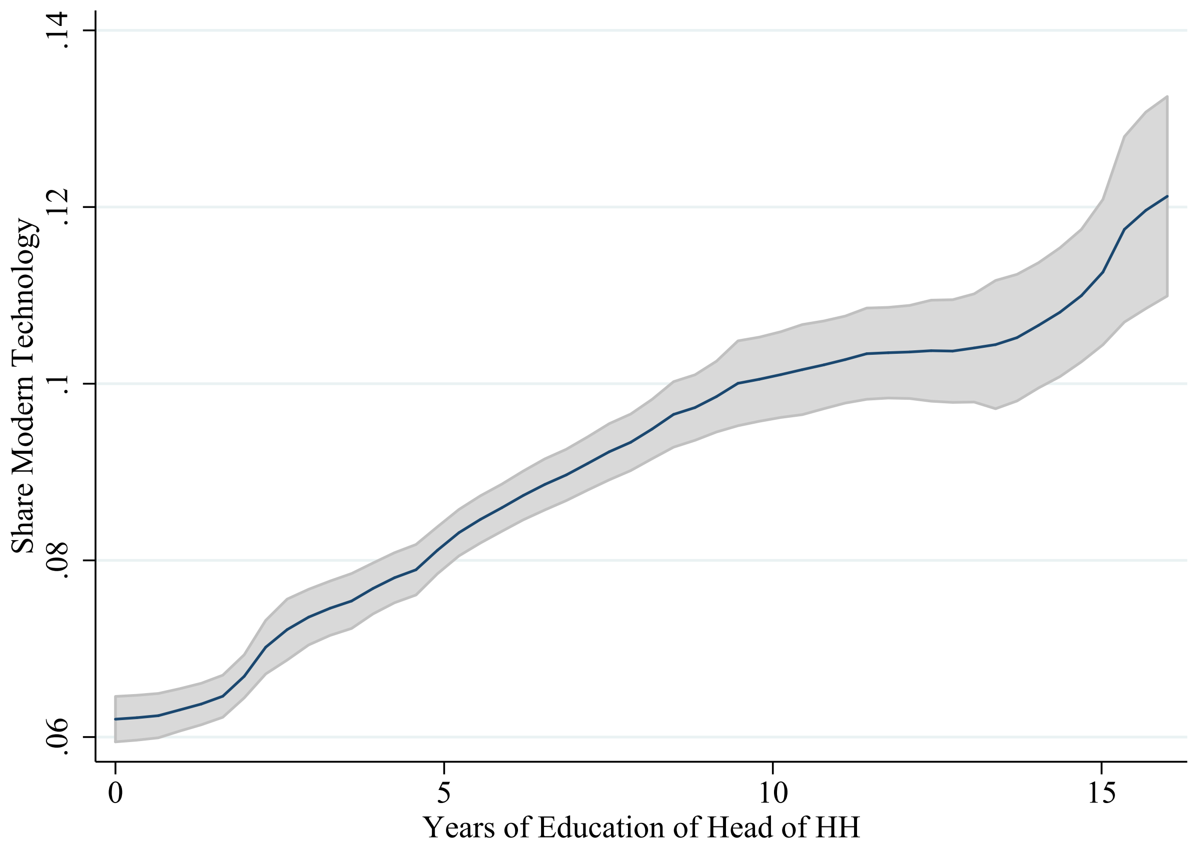

The figure plots estimates from local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 3.5 for discussion.

Figure 3.9: Pre-Existing Shares of Labor Income

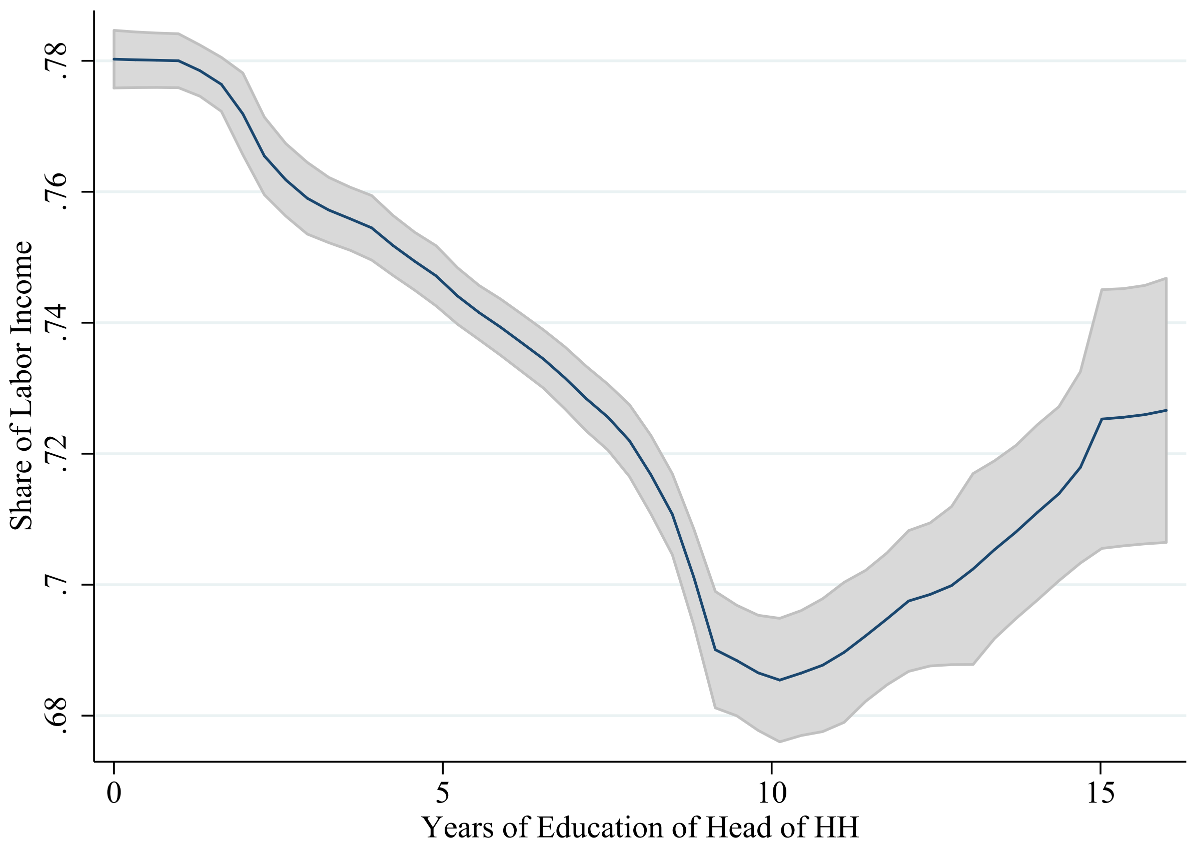

The figure plots estimates from local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 3.5 for discussion.

Figure 3.10: Pre-Existing Expenditure Shares on Agriculture

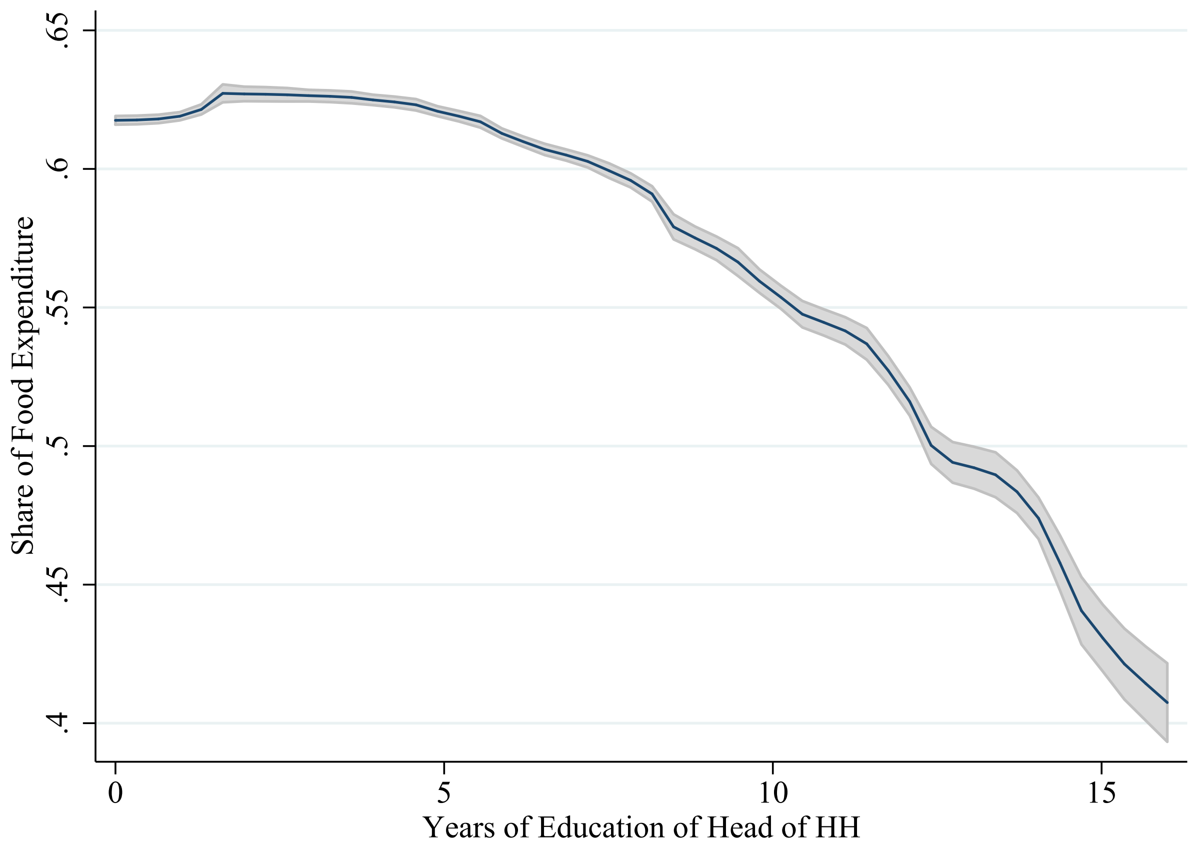

The figure plots estimates from local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 3.5 for discussion.

Figure 3.11: GE Forces as a Function of Scale: Average Effect

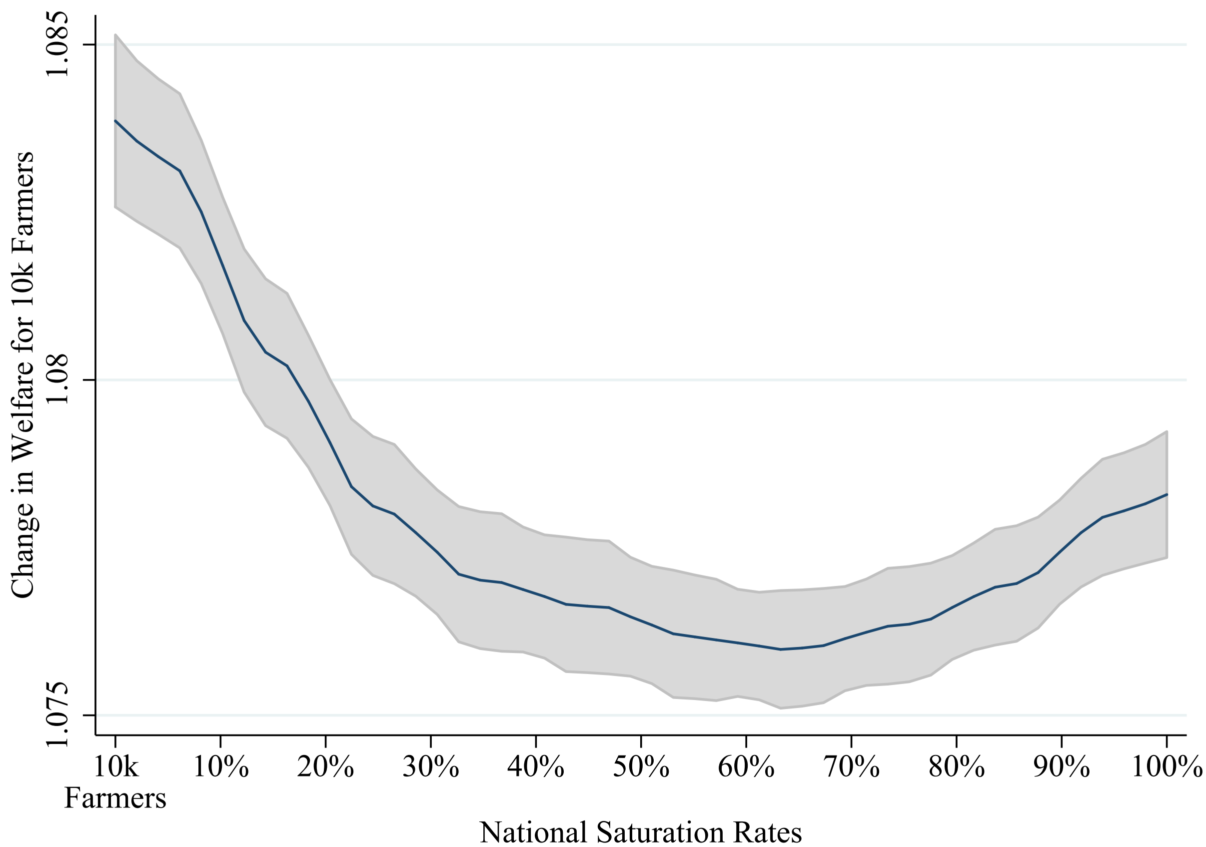

The figure plots estimates from local polynomial regressions. The outcomes are welfare changes among the 10k rural households who initially receive the local intervention. Shaded areas indicate 95 percent confidence intervals. See Section 3.5 for discussion.

Figure 3.12: GE Forces as a Function of Scale: Distribution

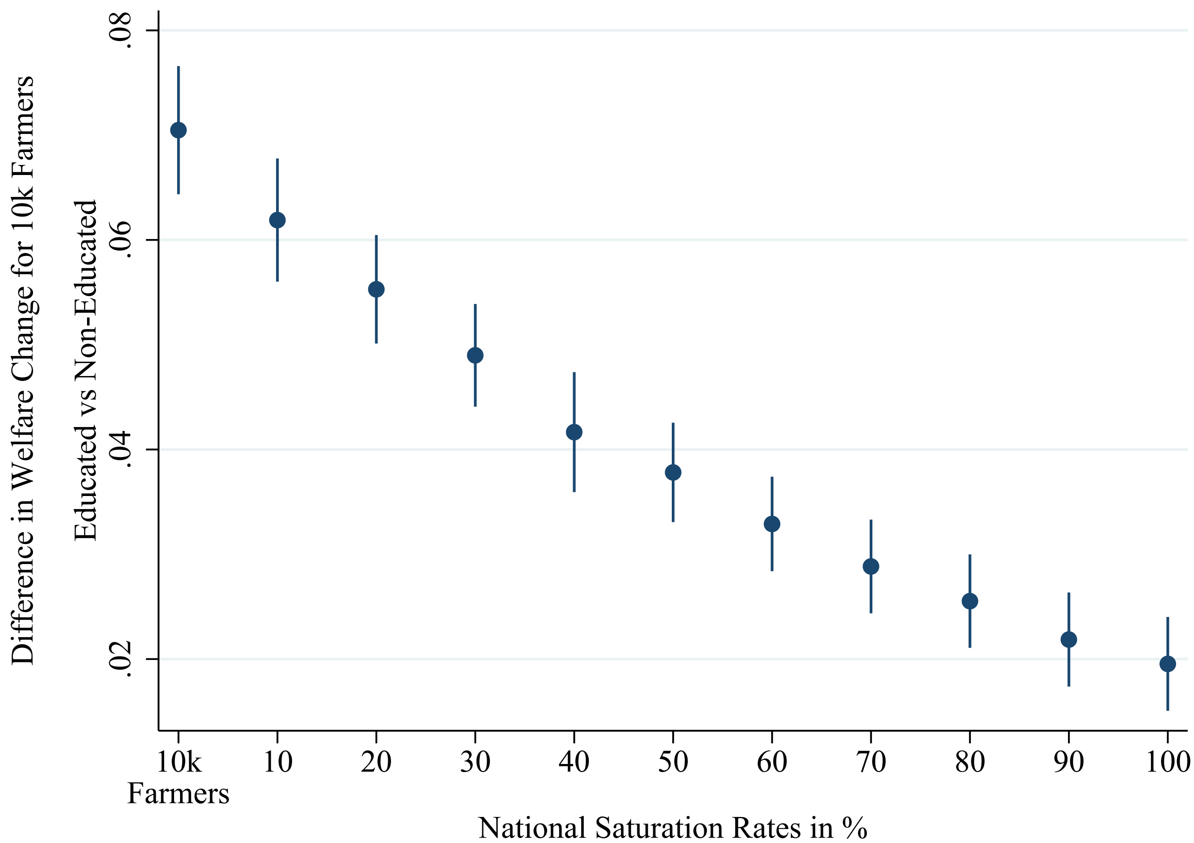

The figure plots point estimates from regressions of welfare changes (LHS) on a dummy for heads of households with 10 or more years of education (RHS) among the 10k of rural households who initially receive the local intervention. The reference category are households without any years of education (roughly 25% of rural households). Point estimates are estimated in separate regressions with counterfactual welfare changes solved for different levels of national saturation. Shaded areas indicate 95 percent confidence intervals. See Section 3.5 for discussion.

Figure 3.13: Difference in Welfare Effects Across Alternative Assumptions About Trading Frictions

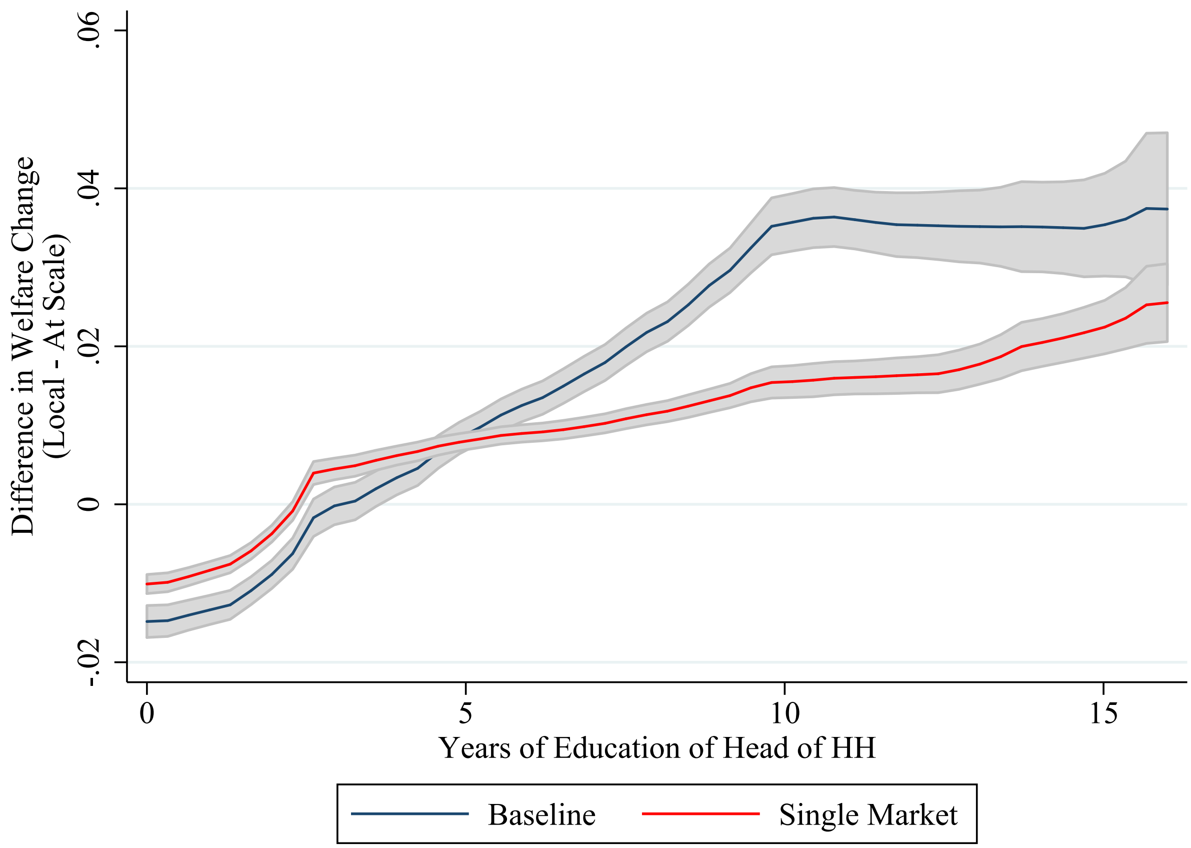

The figure plots estimates from local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 3.5 for discussion.

Figure 3.14: Differences in Welfare Effects as a Function of Pre-Existing Exporting

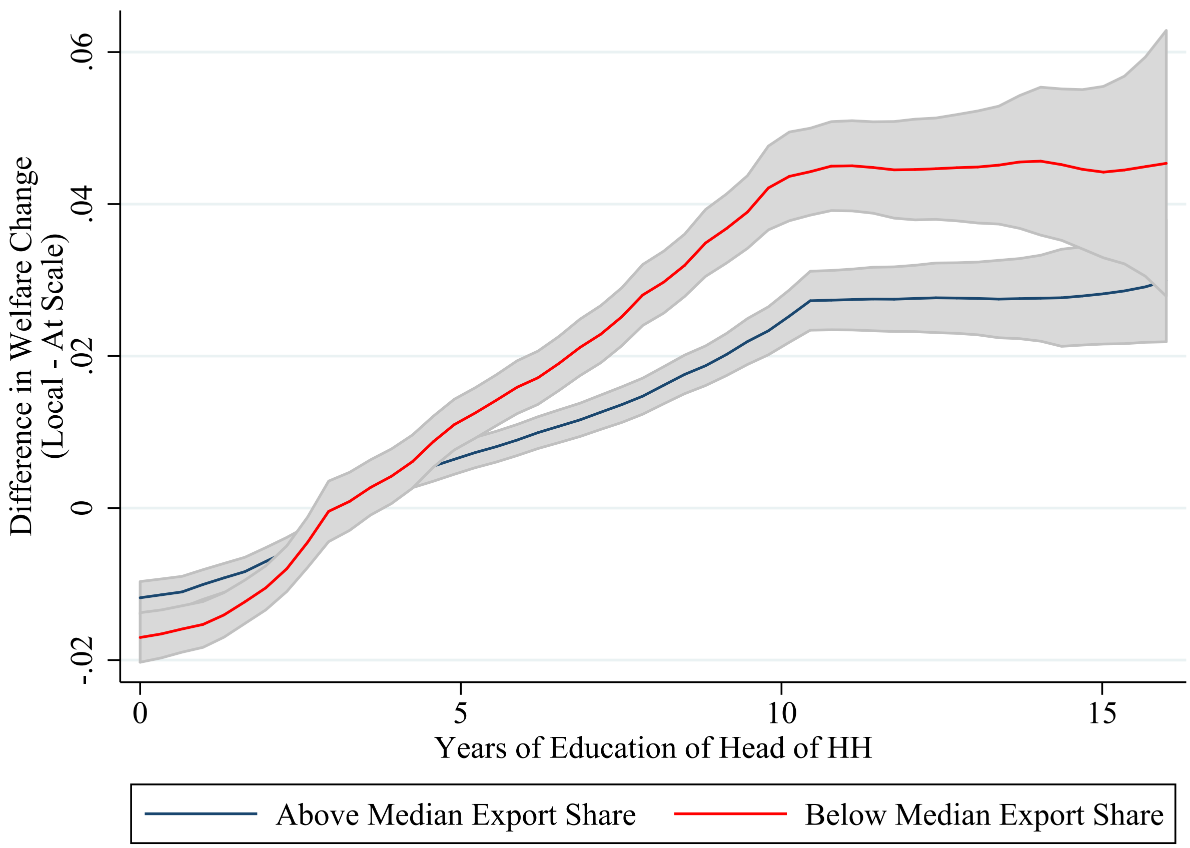

The figure plots estimates from local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 3.5 for discussion.

Figure 3.15: Difference in Welfare Effects Across Alternative Levels of Aggregation

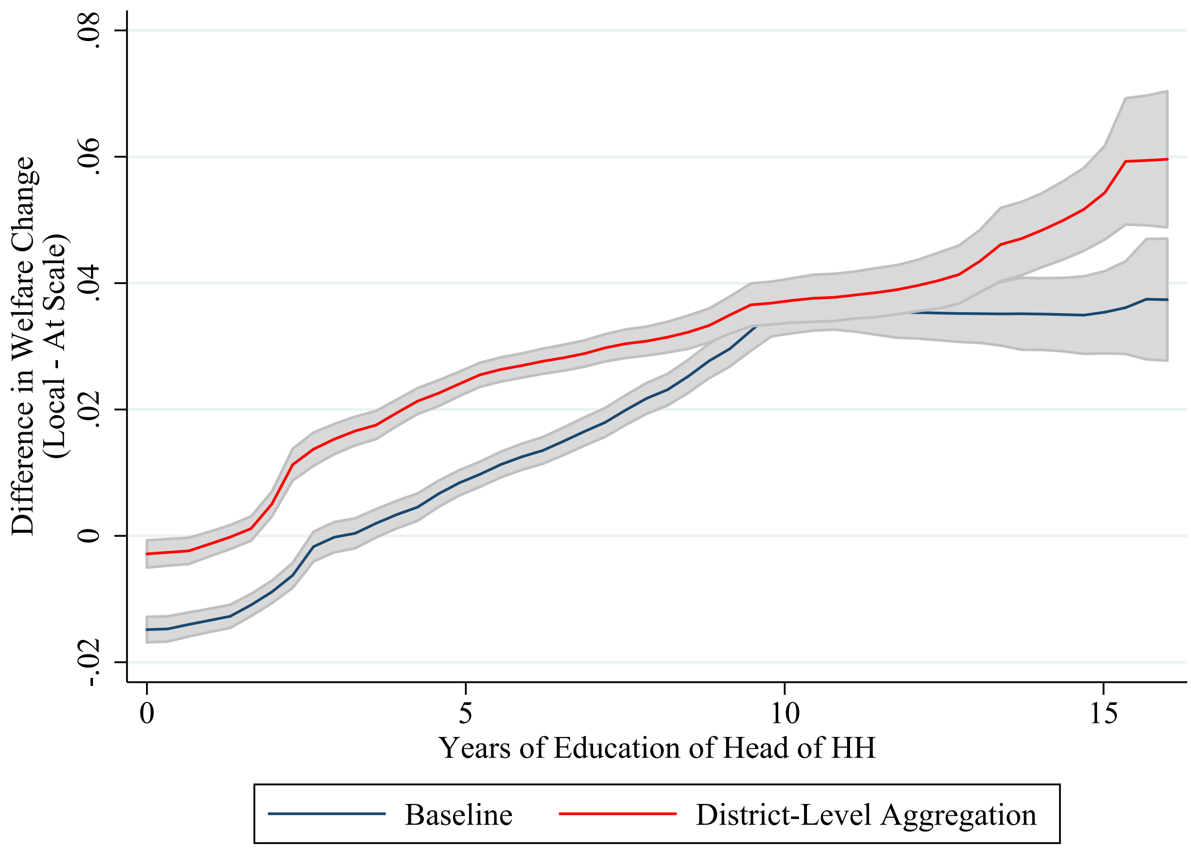

The figure plots estimates from local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 3.5 for discussion.

Figure 3.16: Results Across Alternative Parameter Assumptions: Kappas

Panel A

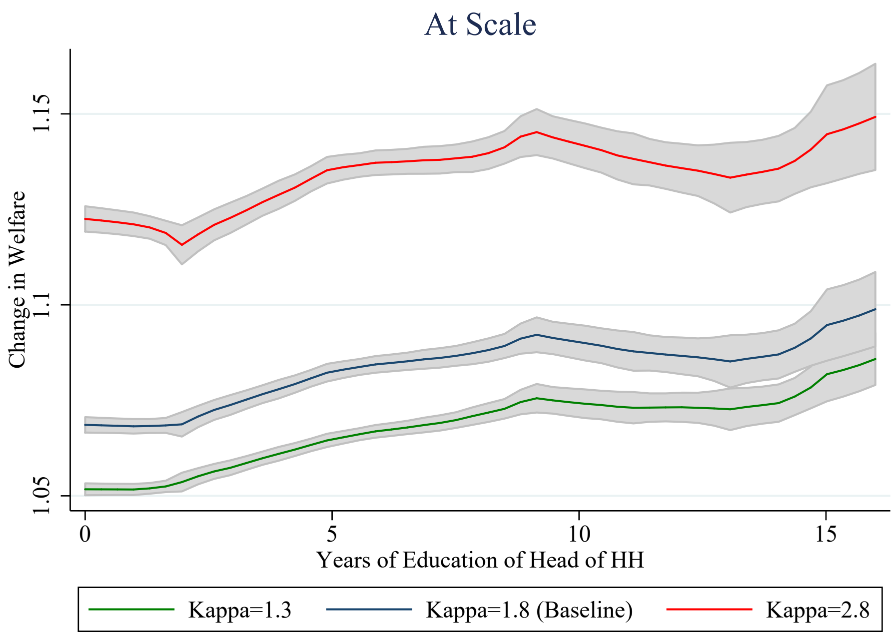

Panel B: Differences in Welfare Effect (Local-At Scale)

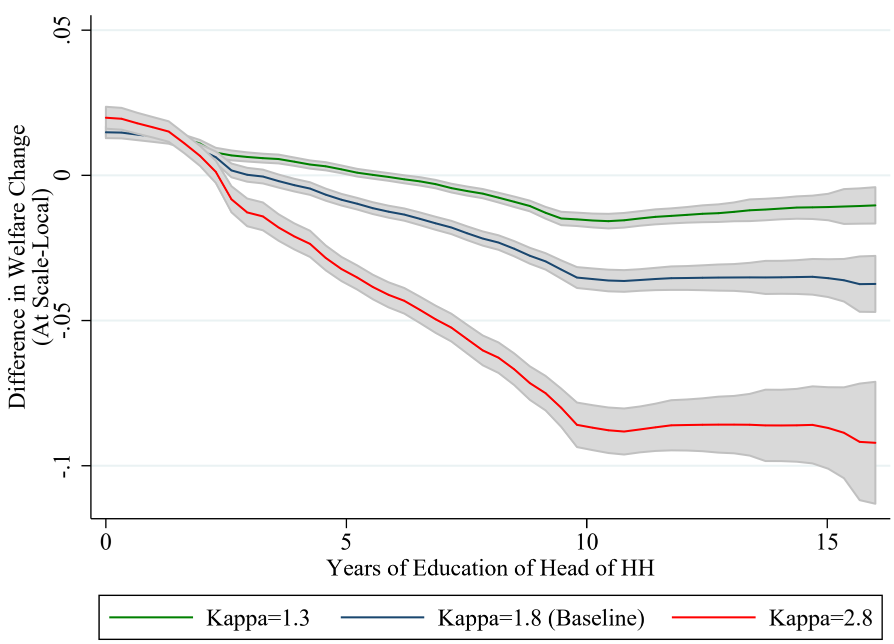

The figure plots estimates from local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 3.5 for discussion.

Figure 3.17: Results Across Alternative Parameter Assumptions: Sigmas

Panel A

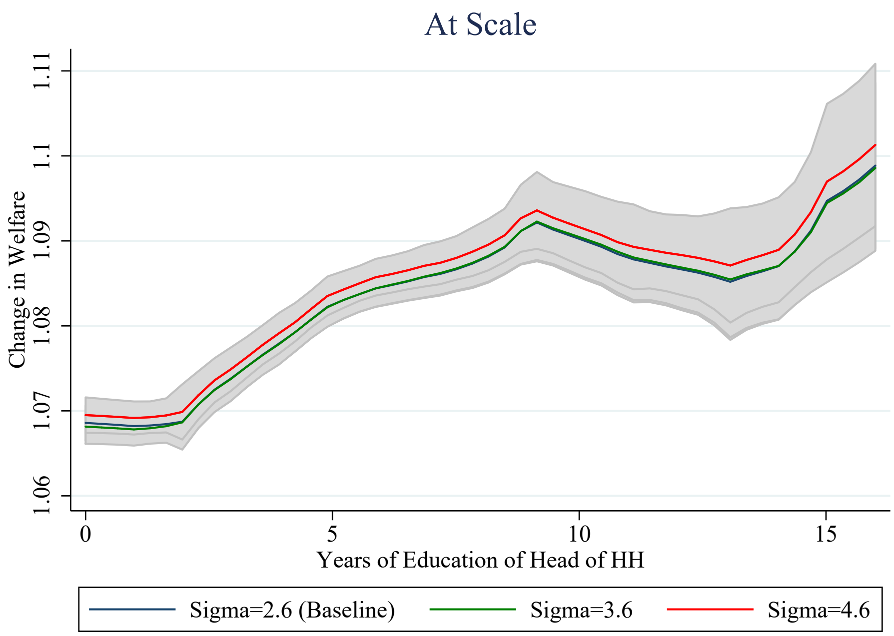

Panel B: Differences in Welfare Effect (Local - At Scale)

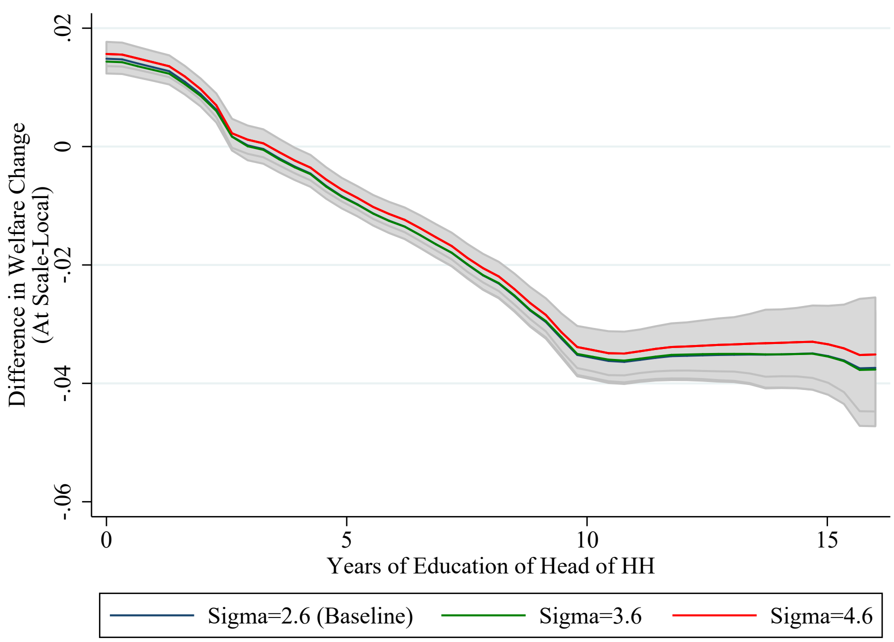

The figure plots estimates from local polynomial regressions. Shaded areas indicate 95 percent confidence intervals. See Section 3.5 for discussion.

Figure 3.18: Model Validation Using Cross-Sectional Moments: Market Prices

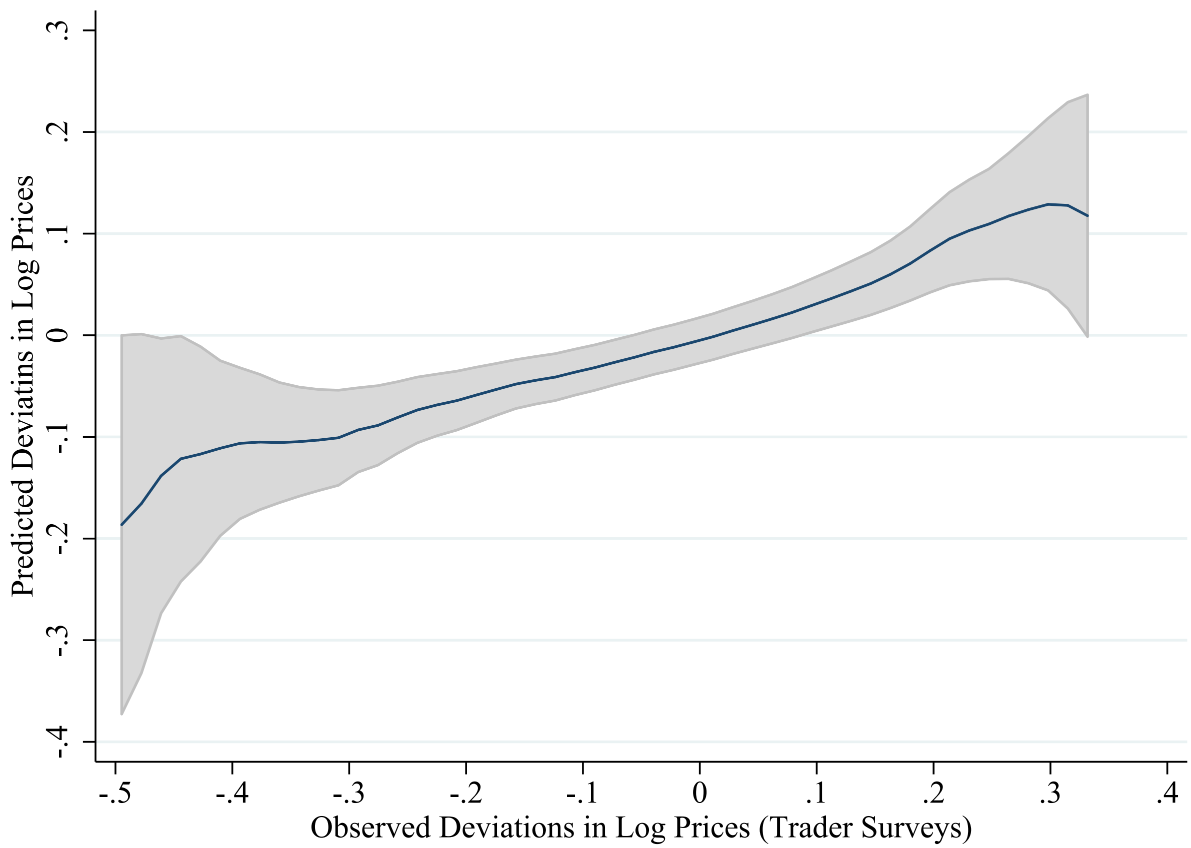

The figure plots estimates from local polynomial regressions. Both y and x-axes are deviations relative to crop-by-time fixed effects. Shaded areas indicate 95 percent confidence intervals. See Section 3.5 for discussion.

Figure 3.19: Road Addition in 2003 as a Natural Experiment for Model Validation

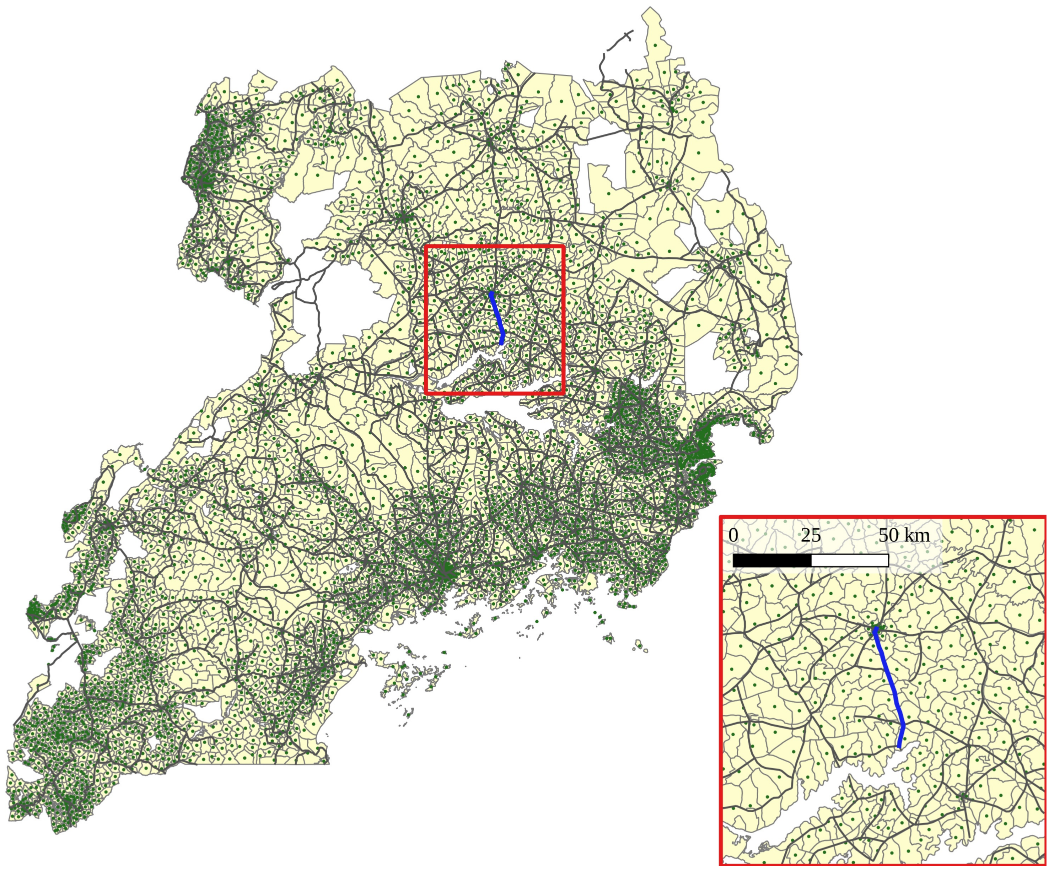

The map displays road additions over the period 2000-2013. See Section 3.5 for discussion.

Table 3.1: Main Crops

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

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

Aggregate and median shares for each of the 9 crops are computed for each of four years of data from the UNPS. The table reports the means and standard deviations across the 4 rounds of data. See Section 3.2 for discussion of the data.

Table 3.2: Regional Price Gaps

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

See Section 3.2 for discussion of the data.

Table 3.3: Farmer Trading vs Subsistence

Panel A

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

Panel B

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

See Section 3.2 for discussion of the data.

Table 3.4: Farmers Sell Their Crops to Local Markets

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

See Section 3.2 for discussion of the data.

Table 3.5: Farmers Re-Allocate Their Land Across Crops Over Time

Panel A

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

Panel B

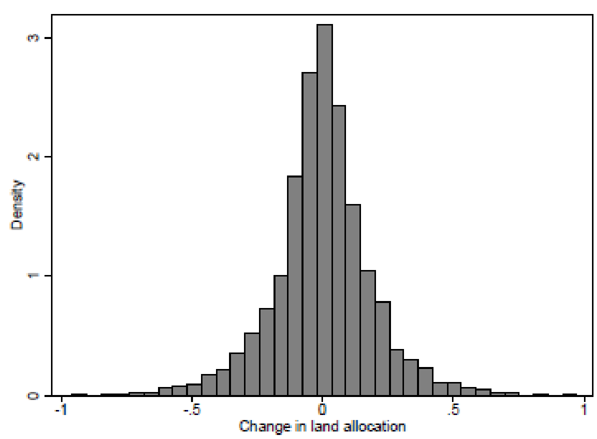

See Section 3.2 for discussion of the data.

Table 3.6: Product Differentiation (Missing Trade Flows

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

*** p<0.01, ** p<0.05, * p<0.1
See Section 3.2 for discussion of the data.

Table 3.7: Nature of Trade Costs

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

Standard errors clusterd at level of bilateral pairs.
*** p<0.01, ** p<0.05, * p<0.1
See Section 3.2 for discussion of the data.

Table 3.8: Technology Adoption and Production Cost Shares

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

Standard errors clusterd at level of farmers.
*** p<0.01, ** p<0.05, * p<0.1
See Section 3.2 for discussion of the data.

Table 3.9: Calibrated Cost Shares in Production

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

See Section 3.4 for discussion and Section 3.2 for description of the data.

Table 3.10: Kappa Estimation

Panel A: First Stage Regressions with $log(\pi_{ik\omega t})$ on Left-Hand Side

VARIABLES	(1) $\log(\pi_{ik\omega t})$ All Years	(2) $\log(\pi_{ik\omega t})$ All Years	(3) $\log(\pi_{ik\omega t})$ 2005-13	(4) $\log(\pi_{ik\omega t})$ 2005-13
$\log(DistBorderPost_i) * Exported_k * \log(WorldP_{kt})$	-2.0073***	-1.8547**	-3.8138***	-5.6151***
	(0.5527)	(0.8636)	(0.6060)	(1.4508)
Observations	27,541	27,520	4,096	4,090
HH-Crop-Tech FX	yes	yes	yes	yes
Crop-Tech-Year FX	yes	.	yes	.
Region-Crop-Tech-Year FX	no	yes	no	yes
Number of clusters	135	135	95	95

Panel B: Log Harvest $(log(y_{ik\omega t}))$ on Left-Hand Side

VARIABLES	(1) OLS All Years	(2) IV All Years	(3) OLS All Years	(4) IV All Years	(5) OLS 2005-13	(6) IV 2005-13	(7) OLS 2005-13	(8) IV 2005-13
$\alpha^{land}_{ik\omega}$*Log(LandShare)	0.6681***	1.1167	0.6695***	0.7074	0.7525***	1.4252**	0.7817***	1.4215***
	(0.0297)	(0.8285)	(0.0292)	(0.8418)	(0.0662)	(0.6678)	(0.0660)	(0.4059)
Observations	27,861	27,541	27,840	27,520	4,296	4,096	4,290	4,090
HH-Crop-Tech FX	yes	yes	yes	yes	yes	yes	yes	yes
Crop-Tech-Year FX	yes	yes	.	.	yes	yes	.	.
Region-Crop-Tech-Year FX	no	no	yes	yes	no	no	yes	yes
Number of clusters	135	135	135	135	95	95	95	95
1st Stage F-Stat		13.09		4.805		41.22		14.80

Panel C: Log Adjusted Output $\log\left(\frac{y_{ik\omega t}}{\prod_{n\in\mathcal{N}} q_{ink\omega t}^{\alpha_{ink\omega}}}\right)$ on Left-Hand Side

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
VARIABLES	All Years	All Years	All Years	All Years	2005-13	2005-13	2005-13	2005-13
$\alpha_{ik\omega}^{land}$*Log(LandShare)	0.4188***	0.1977	0.4157***	0.5966	0.4382***	0.3385	0.4431***	0.6034*
	(0.0349)	(0.6056)	(0.0350)	(0.7303)	(0.0586)	(0.5474)	(0.0635)	(0.3338)
Observations	27,861	27,541	27,840	27,520	4,296	4,096	4,290	4,090
HH-Crop-Tech FX	yes	yes	yes	yes	yes	yes	yes	yes
Crop-Tech-Year FX	yes	yes	.	.	yes	yes	.	.
Region-Crop-Tech-Year FX	no	no	yes	yes	no	no	yes	yes
Number of clusters	135	135	135	135	95	95	95	95
1st Stage F-Stat		13.09		4.805		41.22		14.80

See Section 3.4 for discussion. Standard errors clustered at level of counties. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 3.11: Effect on Household Welfare

VARIABLES	(1) Welfare Local	(2) Welfare At Scale
Hat Change	1.0847***	1.0787***
	(0.0011)	(0.0009)
Observations	10,000	10,000
No Clusters	3731	3731

Standard errors clustered at market-level.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

The table presents effects from the local and from the intervention at scale for the identical representative sample of 10k randomly selected rural households. See Section 3.5 for discussion.

Table 3.12: Effects on Incomes and Prices

Panel A

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Income	Wage	P_manu	P_banana	P_bean	P_cassava	P_coffee	P_groundnut	P_maize	P_millet	P_sorghum	P_sweetpot
VARIABLES	Local	Local	Local	Local	Local	Local	Local	Local	Local	Local	Local	Local
Effect	1.0790***	1.0117***	1.0000	1.0003***	0.9933***	0.9909***	1.0000	0.9993***	0.9906***	1.0001***	0.9994***	0.9992***
	(0.0011)	(0.0004)	(0.0000)	(0.0000)	(0.0002)	(0.0003)	(0.0000)	(0.0001)	(0.0003)	(0.0000)	(0.0001)	(0.0001)
Observations	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000
No Clusters	3731	3731	3731	3731	3731	3731	3731	3731	3731	3731	3731	3731

Panel B

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Income	Wage	P_manu	P_banana	P_bean	P_cassava	P_coffee	P_groundnut	P_maize	P_millet	P_sorghum	P_sweetpot
VARIABLES	At Scale	At Scale	At Scale	At Scale	At Scale	At Scale	At Scale	At Scale	At Scale	At Scale	At Scale	At Scale
Effect	1.0683***	1.0671***	1.0000	1.0031***	0.9568***	0.9794***	0.9985***	0.9750***	0.9719***	1.0282***	0.9970***	1.0087***
	(0.0010)	(0.0012)	(0.0000)	(0.0001)	(0.0004)	(0.0007)	(0.0001)	(0.0004)	(0.0004)	(0.0003)	(0.0001)	(0.0004)
Observations	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000
No Clusters	3731	3731	3731	3731	3731	3731	3731	3731	3731	3731	3731	3731

Standard errors clustered at market-level.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

The table presents effects from the local and from the intervention at scale for the identical representative sample of 10k rural households. See Section 3.5 for discussion.

Table 3.13: Effect on Rural vs Urban Households

VARIABLES	(1) Welfare Local 10k Farmers	(2) Welfare Local Urban HHs	(3) Welfare At Scale 10k Farmers	(4) Welfare At Scale Urban HHs
Hat Change	1.0847***	1.0000	1.0787***	1.0082***
	(0.0011)	(0.0000)	(0.0009)	(0.0005)
Observations	10,000	70	10,000	70
No Clusters	3731		3731	

Standard errors clustered at market-level.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$
Urban households are aggregated as representative agents across 70 cities. See Section 3.5 for discussion.

Table 3.14: Reduced-Form Evidence from Road Building

	(1)	(2)	(3)	(4)
VARIABLES	Log Total Outlays	Log Total Outlays	Log Land	Log Land
Beans*logMA			-1.3318***	-1.3049**
			(0.4534)	(0.5429)
Cassava*logMA			0.9325***	0.2430
			(0.2526)	(0.4371)
Millet*logMA			1.4664***	1.5100**
			(0.5048)	(0.7126)
Maize*logMA			-0.2183	-0.5949
			(0.7370)	(0.5601)
Sorghum*logMA			0.9958	-1.2796***
			(1.0425)	(0.4813)
SweetPot*logMA			-2.8033***	-2.5430***
			(0.6914)	(0.5206)
log(Market Access)	0.5893***	0.9053***		
	(0.1804)	(0.1681)		
Observations	3,434	3,432	6,958	6,922
R-squared	0.7605	0.7892	0.6675	0.7134
HH FX	yes	yes	.	.
Year FX	yes	.	.	.
District-Year FX	no	yes	.	.
HH-Crop-Tech FX	.	.	yes	yes
Crop-Tech-Year FX	.	.	yes	.
District-Crop-Tech-Year FX	.	.	no	yes
Number of clusters	256	255	250	250

See Section 3.5 for discussion. Log Total Outlays are based on annualized household consumption expenditure per capita. Log Land is land area in acres used for the production of each crop. The regression uses the UNPS microdata for 2005 and 2009 to test for the effects of a road addition in 2003 (see map in Figure 3.19). The regression sample excludes parishes in the UNPS data directly affected by the new road. Banana, coffee and groundnuts are excluded as those crops are not grown in the region of the road addition. Standard errors are clustered at the level of parishes.

Table 3.15: Reduced-Form Evidence from Weather Shocks

VARIABLES	(1) Log Total Outlays	(2) Log Harvest	(3) Log Land
Cereals*HotDays		-0.0167	-0.0149*
		(0.0111)	(0.0083)
Tubers*HotDays		-0.0056	-0.0073
		(0.0072)	(0.0081)
Other*HotDays		-0.0056	-0.0131*
		(0.0083)	(0.0070)
Cereals*logRain		0.3856***	0.2551***
		(0.0903)	(0.0803)
Tubers*logRain		-0.7561***	-0.1996***
		(0.0461)	(0.0491)
Other*logRain		0.3954***	0.2557***
		(0.0369)	(0.0386)
HotDays	0.0030		
	(0.0034)		
log(RainFall)	0.6258***		
	(0.0174)		
Observations	8,988	31,478	31,478
R-squared	0.7124	0.7245	0.6045
HH FX	yes	.	.
Year FX	yes	.	.
HH-Crop-Tech FX	.	yes	yes
Crop-Tech-Year FX	.	yes	yes
Number of clusters	381	380	380

See Section 3.5 for discussion. Log Total Outlays are based on annualized household consumption expenditure per capita. Log Land is land area in acres used for the production of each crop. Log Harvest is the amount of crop harvest measured in kg. HotDays and logRainfall are the number of days above 29 degrees Celsius and the total amount of rainfall recorded during growing seasons per year. Standard errors are clustered at the level of parishes.

Chapter 4

Dissertation Conclusion

This dissertation studies how the average and distributional implications of policies are impacted by spatial linkages that operate in general equilibrium. These linkages can take the form of varying access to service firms as a function of the local neighborhood composition as in chapter 2 or price effects in markets for agricultural crops and local labor markets as in chapter 3. Both chapters provide important conclusions for the design of policies and contribute to the discussions on urban inequality in chapter 2 and the fight against poverty in developing countries in chapter 3.

From a policy maker's perspective, *Two-Sided Sorting and Spatial Inequality in Cities* tells two stories. On the one hand, it is a cautionary tale in the sense that policy makers need to take into account how the local access to services changes in response to place-based policies and consequently how heterogeneous households value impacted locations. For example, local group-specific price indices of consumption respond to policies targeting the location of firms differently for high and low-skilled households, hence, amplifying household mobility differently. Furthermore, modelling the endogenous nature of amenities as reduced-form spillovers as a proxy for the price index channel cannot capture these effects qualitatively and quantitatively. On the other hand, the forces outlined in chapter 2 can be used as a policy tool. For example, policies aimed at attracting high value services into a neighborhood such as a boutique shopping district can induce neighborhood change by attracting rich households and unleash the endogenous amplification outlined in this chapter.

Scaling Agricultural Policy Interventions: Theory and Evidence from Uganda in chapter 2 similarly has important policy implications. In particular, implementing policies aimed at increasing agricultural productivity at the national scale can benefit low income households in developing countries, the poorest of the poor, more than we would expect from local interventions. The chapter, therefore, provides further support for the hypothesis that the key to development lies in improving productivity in agriculture.

Bibliography

Adao, Rodrigo, Costinot, Arnaud, & Donaldson, Dave. 2017. Nonparametric counterfactual predictions in neoclassical models of international trade. *American Economic Review*, 107(3), 633–89.

Adao, Rodrigo, Arkolakis, Costas, & Esposito, Federico. 2018. Spatial Linkages, Global Shocks, and Local Labor Markets: Theory and Evidence. *Working Paper*.

Agarwal, Sumit, Jensen, J Bradford, & Monte, Ferdinando. 2018. The Geography of Consumption.

Aguiar, Mark, & Bils, Mark. 2015. Has consumption inequality mirrored income inequality? *American Economic Review*, 105(9), 2725–56.

Ahlfeldt, Gabriel M, Redding, Stephen J, Sturm, Daniel M, & Wolf, Nikolaus. 2015. The economics of density: Evidence from the Berlin Wall. *Econometrica*, 83(6), 2127–2189.

Albouy, David, Ehrlich, Gabriel, & Liu, Yingyi. 2016. Housing demand and expenditures: How rising rent levels affect behavior and cost-of-living over space and time.

Allen, Treb, & Arkolakis, Costas. 2014. Trade and the Topography of the Spatial Economy. *Quarterly Journal of Economics*, 129(3), 1085–1140.

Allen, Treb, Arkolakis, Costas, & Li, Xiangliang. 2015. Optimal city structure. *Yale University, mimeograph*.

Athey, Susan. 2002. Monotone comparative statics under uncertainty. *The Quarterly Journal of Economics*, 117(1), 187–223.

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

Atkin, David, Faber, Benjamin, & Gonzalez-Navarro, Marco. 2018. Retail globalization and household welfare: Evidence from mexico. *Journal of Political Economy*, 126(1), 1–73.

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

Baqaee, David, & Farhi, Emmanuel. 2018. Productivity and Misallocation in General Equilibrium. *Working Paper*.

Barnatchez, Keith, Crane, Leland Dod, & Decker, Ryan. 2017. An assessment of the national establishment time series (nets) database.

Baum-Snow, Nathaniel, & Hartley, Daniel A. 2016. Accounting for central neighborhood change, 1980-2010.

Behrens, Kristian, Duranton, Gilles, & Robert-Nicoud, Frédéric. 2014. Productive cities: Sorting, selection, and agglomeration. *Journal of Political Economy*, 122(3), 507–553.

Bergquist, Lauren Falcao. 2017. Pass-Through, Competition, and Entry in Agricultural Markets: Experimental Evidence from Kenya. *Working Paper*.

Bergquist, Lauren Falcao, & McIntosh, Craig. 2019. Market Linkages for Smallholder Farmers in Uganda. *Working Paper*.

Bergquist, Lauren Falcao, Faber, Benjamin, Fally, Thibault, Hoelzlein, Matthias, Miguel, Edward, & Rodriguez-Clare, Andres. 2019. *Scaling Agricultural Policy Interventions: Theory and Evidence from Uganda*. Tech. rept. Working Paper.

Borusyak, Kirill, & Jaravel, Xavier. 2018. *The Distributional Effects of Trade: Theory and Evidence from the United States*.

Borusyak, Kirill, Hull, Peter, & Jaravel, Xavier. 2018. *Quasi-experimental shift-share research designs*. Tech. rept. National Bureau of Economic Research.

Brinkman, Jeffrey, Coen-Pirani, Daniele, & Sieg, Holger. 2015. Firm dynamics in an urban economy. *International Economic Review*, 56(4), 1135–1164.

Brueckner, Jan K, Thisse, Jacques-Francois, & Zenou, Yves. 1999. Why is central Paris rich and downtown Detroit poor?: An amenity-based theory. *European economic review*, 43(1), 91–107.

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

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

Busso, Matias, Gregory, Jesse, & Kline, Patrick. 2013. Assessing the incidence and efficiency of a prominent place based policy. *American Economic Review*, 103(2), 897–947.

Carter, Michael, Laajaj, Rachid, & Yang, Dean. 2014. Subsidies and the Persistence of Technology Adoption: Field Experimental Evidence from Mozambique. *NBER Working Paper*.

Comin, Diego, Lashkari, Danial, & Mestieri, Martı. 2018. Structural Change with Long-run Income and Price Effects.

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

Couture, Victor. 2016. Valuing the consumption benefits of urban density. *University of California, Berkeley. Processed*.

Couture, Victor, & Handbury, Jessie. 2017. *Urban revival in America, 2000 to 2010*. Tech. rept. National Bureau of Economic Research.

Couture, Victor, Gaubert, Cecile, Handbury, Jessie, & Hurst, Erik. 2019. Income Growth and the Distributional Effects of Urban Spatial Sorting.

Davis, Donald R, Dingel, Jonathan I, Monras, Joan, & Morales, Eduardo. 2019. How segregated is urban consumption? *Journal of Political Economy*, 127(4), 000–000.

Davis, Morris A, Gregory, Jesse, & Hartley, Daniel A. 2018. The Long Run Effects of Low Income Housing on Neighborhood Composition.

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

Deaton, Angus. 2010. Instruments, Randomization, and Learning about Development. *Journal of Economic Literature*, 48(2), 424–55.

DellaVigna, Stefano, & Gentzkow, Matthew. 2019. Uniform pricing in us retail chains. *The Quarterly Journal of Economics*, 134(4), 2011–2084.

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

Diamond, Rebecca, & McQuade, Tim. 2019. Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development. *Journal of Political Economy*, 127(3), 000–000.

Diamond, Rebecca, McQuade, Timothy, & Qian, Franklin. 2018. *The effects of rent control expansion on tenants, landlords, and inequality: Evidence from San Francisco*. Tech. rept. National Bureau of Economic Research.

Dolfen, Paul, Einav, Liran, Klenow, Peter J, Klopack, Benjamin, Levin, Jonathan D, Levin, Laurence, & Best, Wayne. 2019. *Assessing the Gains from E-commerce*. Tech. rept. National Bureau of Economic Research.

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

Eaton, Jonathan, & Kortum, Samuel. 2002. Technology, Geography, and Trade. *Econometrica*, 70(5), 1741–79.

Egger, Dennis, Haushofer, Johannes, Miguel, Edward, Neihaus, Paul, & Walker, Michael. 2019. General Equilibrium Effects of Cash Transfers. *Working Paper*.

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

Fajgelbaum, Pablo, & Gaubert, Cecile. 2019. *Optimal spatial policies, geography and sorting*. Tech. rept. National Bureau of Economic Research.

Fajgelbaum, Paublo D, & Khandelwal, Amit K. 2016. Measuring the Unequal Gains from Trade. *Quarterly Journal of Economics*, 131, 1113–80.

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

Gaubert, Cecile. 2018. Firm sorting and agglomeration. *American Economic Review*, 108(11), 3117–53.

Glaeser, Edward L, Kim, Hyunjin, & Luca, Michael. 2018. Nowcasting gentrification: using yelp data to quantify neighborhood change. *Pages 77–82 of: AEA Papers and Proceedings*, vol. 108.

Guerrieri, Veronica, Hartley, Daniel, & Hurst, Erik. 2013. Endogenous gentrification and housing price dynamics. *Journal of Public Economics*, 100, 45–60.

Handbury, Jessie. 2013. Are poor cities cheap for everyone? non-homotheticity and the cost of living across us cities. *The wharton school research paper*, 1–054.

Hanoch, Giora. 1975. Production and demand models with direct or indirect implicit additivity. *Econometrica (pre-1986)*, 43(3), 395.

Heckman, James, & Smith, Jeffrey. 1995. Assessing the Case for Social Experiments. *The Journal of Economic Perspectives*, 9(2), 85–110.

Helliwell, John F, & Verdier, Genevieve. 2001. Measuring internal trade distances: a new method applied to estimate provincial border effects in Canada. *Canadian Journal of Economics*, 1024–1041.

Hottman, Colin, & Monarch, Ryan. 2018. Estimating Unequal Gains across US Consumers with Supplier Trade Data. *FRB International Finance Discussion Paper*.

Hubmer, Joachim. 2018. *The Race Between Preferences and Technology*. Tech. rept.

Krugman, Paul. 1991. Increasing returns and economic geography. *Journal of political economy*, 99(3), 483–499.

Lee, Sanghoon, & Lin, Jeffrey. 2017. Natural amenities, neighbourhood dynamics, and persistence in the spatial distribution of income. *The Review of Economic Studies*, 85(1), 663–694.

Magruder, Jeremy. 2018. An Assessment of Experimental Evidence on Agricultural Technology Adoption in Developing Countries. *Annual Review of Resource Economics*, 10, 299–316.

Matsuyama, Kiminori. 2019. Engel's Law in the Global Economy: Demand-Induced Patterns of Structural Change, Innovation, and Trade. *Econometrica*, 87(2), 497–528.

Moffitt, Robert. 2009. Estimating Marginal Treatment Effects in Heterogeneous Populations. *Annales d'Economie et de Statistique*, 239–261.

Monte, Ferdinando, Redding, Stephen J, & Rossi-Hansberg, Esteban. 2018. Commuting, migration, and local employment elasticities. *American Economic Review*, 108(12), 3855–90.

Neumark, David, Zhang, Junfu, & Wall, Brandon. 2005. Employment Dynamics and Business Relocation: New Evidence from the National Establishment Time Series.

Ravallion, Martin. 2009. Should the Randomistas Rule? *The Economists' Voice*, 6(2).

Ravallion, Martin. 2018. Should the Randomistas (Continue to) Rule? *Center for Global Development Working Paper*.

Redding, Stephen. 2016. Goods Trade, Factor Mobility, and Welfare. *Journal of International Economics*, 101, 148–67.

Redding, Stephen, & Weinstein, David. 2017. *Aggregating from micro to macro patterns of trade*. Tech. rept. National Bureau of Economic Research.

Redding, Stephen J, & Weinstein, David E. 2019. Measuring Aggregate Price Indexes with Taste Shocks: Theory and Evidence for CES Preferences.

Reynolds, C Lockwood, & Rohlin, Shawn M. 2015. The effects of location-based tax policies on the distribution of household income: evidence from the federal Empowerment Zone program. *Journal of Urban Economics*, 88, 1–15.

Rodrik, Dani. 2008. The New Development Economics: We Shall Experiment, but How Shall we Learn? *Working Paper Series, Harvard University, John F. Kennedy School of Government*.

Schiff, Nathan. 2014. Cities and product variety: evidence from restaurants. *Journal of Economic Geography*, 15(6), 1085–1123.

Sotelo, Sebastian. 2017. Domestic Trade Frictions and Agriculture. *Working Paper*.

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

Su, Yichen. 2018a. Measuring the Value of Urban Consumption Amenities: A Time-Use Approach.

Su, Yichen. 2018b. The Rising Value of Time and the Origin of Urban Gentrification.

Svensson, Jakob, & Yanagizawa-Drott, David. 2012. Estimating Impact in Partial vs. General Equilibrium: A Cautionary Tale from a Natural Experiment in Uganda. *Working Paper*.

Theodos, Brett, Meixell, Brady, & Hedman, Carl. 2018. Did States Maximize Their Opportunity Zone Selections? *Retrieved from Urban Institute: https://www. urban. org/sites/default/files/publication/98445/did_states_maximize_their_ opportunity_zone_selections_1. pdf*.

Tsivanidis, Nick. 2018. The Aggregate and Distributional Effects of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio.

Waldfogel, Joel. 2008. The median voter and the median consumer: Local private goods and population composition. *Journal of Urban Economics*, 63(2), 567–582.

Ziv, Oren. Productivity, Density, and Sorting.