Essays on Macroeconomic Development

by

Alberto Rivera Padilla

A Dissertation Presented in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Approved April 2019 by the
Graduate Supervisory Committee:

Todd Schoellman, Co-Chair
Berthold Herrendorf, Co-Chair
Domenico Ferraro

ARIZONA STATE UNIVERSITY

May 2019

ABSTRACT

This dissertation consists of two essays with a macroeconomic approach to economic development. These essays explore specific barriers that prevent economic agents from exploiting opportunities across regions or sectors in developing countries, and to what extent the observed allocations are inefficient outcomes or just an efficient response to economic fundamentals and technological constraints.

The first chapter is motivated by the fact that a prominent feature of cities in developing countries is the existence of slums: locations with low housing-quality and informal property rights. This paper focuses on the allocation of land across slums and formal housing, and emphasizes the role of living in central urban areas for the formation of slums. I build a quantitative spatial general equilibrium model to study the aggregate effects of anti-slum policies and use microdata from India for the quantitative implementation. According to my findings, demolishing slums in central urban areas leads to a decrease in welfare, aggregate labor productivity, and urban population. In contrast, decreasing formal housing distortions in India to the U.S. level increases the urban population share by 20% and labor productivity by 2.4%, and reduces the share of the urban population living in slums by 19%.

The second chapter is motivated by the fact that labor productivity gaps between rich and poor countries are much larger for agriculture than for non-agriculture. Using detailed data from Mexican farms, this paper shows that value added per worker is frequently over two times larger in cash crops than in staple crops, yet most farmers choose to produce staples. These findings imply that the agricultural productivity gap is actually a staple productivity gap and understanding production decisions of farmers is crucial to explain why labor productivity is so low in poor countries. This paper develops a general equilibrium framework in which subsistence consumption and interregional trade costs determine the efficient selection of farmers into types of crops. The quantitative results of the model imply that decreasing trade costs in Mexico to the U.S. level reduces the ratio of employment in staple to cash crops by 17% and increases agricultural labor productivity by 14%.

To my parents and sister, for their unconditional support.

To Karolina, for her patience and encouragement throughout these years.

ACKNOWLEDGMENTS

I am extremely grateful to Todd Schoellman, Berthold Herrendorf, and Domenico Ferraro for their guidance and support. The chapters of this dissertation have benefited from comments and insights by seminar participants at Arizona State, Bank of Mexico, Cal State Fullerton, CIDE, and Econometric Society NASM 2018, especially those from Alexander Bick, Alvin Murphy, Bart Hobijn, Galina Vereshchagina, Gustavo Ventura, and Richard Rogerson.

I am thankful to Edward Prescott for his mentorship during my time as his research assistant and for inviting me to spend a summer working with him at the Minneapolis Fed. I am also grateful to other ASU faculty members for their time and advice, including B. Ravikumar, Edward Schlee, Daniel Silverman, Manjira Datta, Natalia Kovrijnykh, and Stephie Fried. Finally, I thank my fellow students at ASU for helpful conversations and the staff in the economics department for all their help.

TABLE OF CONTENTS

LIST OF TABLES

Table Page

LIST OF FIGURES

Chapter 1

SLUMS, ALLOCATION OF TALENT, AND BARRIERS TO URBANIZATION

1.1 Introduction

More than half of the total world population lives in urban regions today. Particularly, urbanization has increased dramatically in poor countries. In 1960, the urban population share was 17% and 15% in South Asia and Sub-Saharan Africa, respectively; these shares were 34% and 40% in 2017. A prevalent feature of urban regions in poor countries is the large population share living in slums: locations characterized by low-quality housing, lack of public services, and informal property rights. Understanding the formation of slums as part of the urbanization process is crucial because the emergence and growth of cities are closely related to economic development.[1]

The existence of slums implies a trade-off for policymakers. On one hand, slums are locations with poorly built houses often occupying valuable land that could otherwise be used for commercial purposes or construction of formal housing.[2] On the other hand, slums are dense areas where dwellers, especially rural migrants, can gain access to labor market opportunities in urban centers. This trade-off is intensified by the presence of burdensome regulations that limit the supply of formal housing, which further raises prices in this sector and increases the demand for and value of residential space in slums.

The goal of this paper is to measure the aggregate implications of anti-slum policies in a context of inefficient housing markets. I build a general equilibrium model

[1]See Kuznets (1968) and Glaeser (2011).

[2]Henderson *et al.* (2016) show evidence of the inefficient land-use generated by slums in Nairobi.

that features spatial differences in human capital, housing rents, and productivity. Urban locations offer better economic opportunities in terms of income and returns to education, but have a higher cost of living than rural regions. Housing rents in urban locations depend on the supply of residential land, which can be used for producing two types of housing: formal or slums. The demand for land is higher in central urban areas, so the marginal value of land occupied by slums is high; nevertheless, such slums are formed because workers want to live close to the most productive areas of the city and formal housing is costly due, in part, to regulations.

The model takes into account the general equilibrium effects of counterfactual policies through changes in prices and sorting of talent. In particular, housing prices act as a congestion cost in urban regions due to the limited supply of land, while selection has two effects: there are potential gains of sorting based on comparative advantage across locations, but the marginal resident in each location has a relatively low productivity in that place. Furthermore, the local supply of skills is a function of returns to education in each location and wages depend on the spatial allocation of heterogeneous workers.

I focus on the case of India for the empirical motivation and quantitative implementation of the model. India is widely considered a country with inefficient urban policies, governments have opted to sanitize cities by demolishing slums in urban areas where land is most valuable, and its population size (around 1.3 billion) makes it a particularly important case among developing countries.[3] Moreover, Figure 1.1 shows that, in comparison to countries with similar and lower levels of development, India has a small share of urban population, as well as a low share of the urban population living in slums. This paper explores the importance of anti-slum policies

[3]See Bertaud (2002), Brueckner and Sridhar (2012), Bertaud and Brueckner (2005), Diwakar and Peter (2016), Kumar (2010), Dupont (2008), and Bhan (2009).

Figure 1.1: Urban and Slum Population in Poor Countries

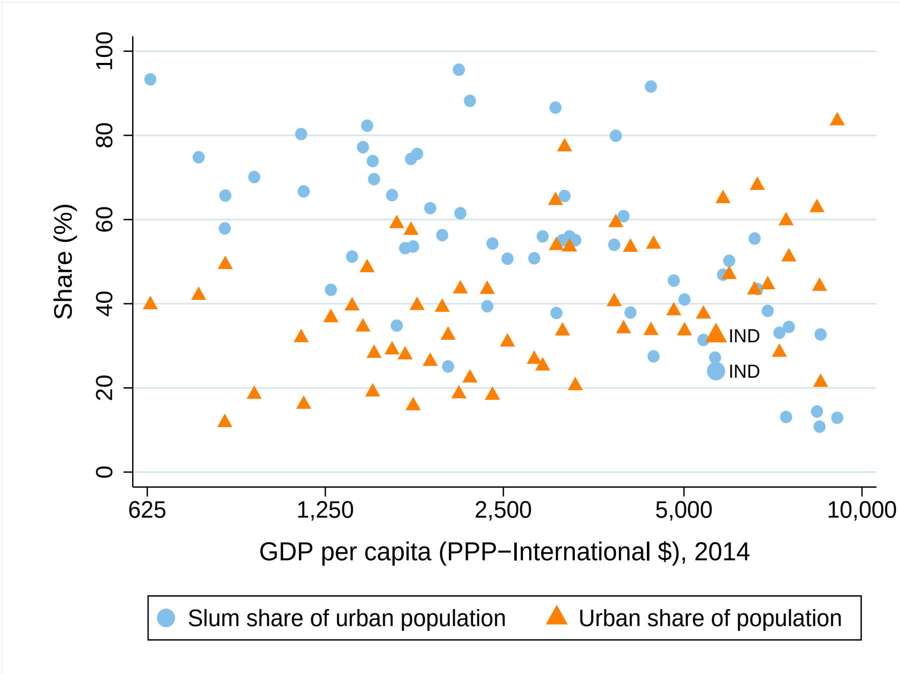

Source: UN Habitat, World Development Indicators.

and restrictive building regulations in explaining such facts.

To discipline the model quantitatively, I exploit rich microdata from India. Then, I use the estimated model to evaluate the aggregate effects of two different policies that can reduce the slum population: evicting slum dwellers from central urban areas, which is a common policy in poor countries, and reducing distortions in formal housing. The quantitative results show that evicting half of the population from central slums generates a fall in total output of 1.6% and decreases urban population by 11%. Among individuals evicted from central slums, 76% move to rural regions and 17% move to formal housing in the central urban region. These results imply that destroying central slums may hinder the urbanization process by pushing people to rural regions and deterring urban migration.[4] Given the high price of formal housing

[4]In a recent paper, Jedwab and Vollrath (2019) also find that restricting urban migration into

in central urban areas, slums are the only option to gain access to the most productive urban jobs and the periphery is not an attractive alternative.

In contrast, reducing formal housing distortions in India to the U.S level increases the urban population share by 20% and aggregate labor productivity by 2.4%, and reduces the share of the urban population living in slums by 19%. Furthermore, I find that formal housing distortions substantially amplify the negative effects of destroying central slums, especially in terms of rural-urban migration, by limiting the capacity of the formal sector to accommodate more population with the available supply of land. The quantitative results imply that relaxing distortionary regulations is a more effective policy to increase the urbanization rate and reduce the urban population living in slums.

This paper is related to a small literature that has considered the economics of slums and their relationship to economic development.[5] In particular, Monge-Naranjo *et al.* (2018) document similar findings on economic opportunity gaps between rural and urban regions in Brazil, including slums. They stress the role of intergenerational links in explaining the emergence of slums, while this paper focuses on the allocation of land across types of housing and the importance of living in central urban areas for the formation of slums. In a related paper, Cavalcanti *et al.* (2018) examine the formation of slums in Brazil and analyze the impact of barriers in formal markets. Their focus is on housing choices within a city, whereas this paper considers the effects of housing and slums policies on the decision to live in the countryside or one of multiple urban locations.[6] Another difference with these papers is that I study a

informal urban areas reduces welfare by keeping people in rural areas.

[5]See Marx *et al.* (2013) for a review of slums in developing countries, and Brueckner and Lall (2015) for a survey on urbanization and housing in the same class of countries.

[6]Alves (2018) also looks at the formation of slums in Brazil and stresses the importance of differences in housing supply elasticity between slums and non-slums in cities.

country where most people still live in rural areas and barriers to urbanization might be a key obstacle for development.

This paper also contributes to the literature on India's urbanization. Munshi and Rosenzweig (2016) document large wage gaps between urban and rural regions in India, and argue that caste-based insurance networks act as a barrier to rural-urban migration.[7] I document evidence on wage differences between slums and formal housing in urban regions, and focus on housing regulations as possible barriers to urbanization in India. Additionally, this paper relates to the large macroeconomics literature on urban-rural development. Specifically, Gollin *et al.* (2014) and Herrendorf and Schoellman (2018) look at human capital gaps as a possible explanation for the large productivity gaps between non-agriculture and agriculture. I build on their ideas to measure spatial differences in human capital across locations.

Finally, this paper relates to the literature on quantitative spatial economics. I build on the framework of Ahlfeldt *et al.* (2015) and Bryan and Morten (2018), who draw on results from Eaton and Kortum (2002) to characterize the allocation of workers across locations as gravity equations that depend on congestion and agglomeration forces.[8] In a recent paper, Gechter and Tsivanidis (2017) use a quantitative general equilibrium model to assess the impact of land redevelopment based on a natural experiment in Mumbai, taking into account the effects on slums. This paper focuses on the aggregate effects of restrictive building regulations and anti-slum policies in India, taking into consideration differences in labor market opportunities between urban and rural regions.[9]

[7]Hnatkovska and Lahiri (2014) find that urban-rural wage gaps in India have decreased due to urban sprawl (rural areas reclassified as urban), and urban migration accounts for only a small fraction of the decline.

[8]Similar frameworks include Monte *et al.* (2018) and Rossi-Hansberg *et al.* (2017).

[9]In a related paper, Hsieh and Moretti (2018) show that housing constraints are important for

The rest of the paper is organized as follows. Section 1.2 presents the main data used in the paper and empirical evidence from India. Section 1.3 describes policies related to housing development and slums in this country. Next, Section 2.3 introduces a spatial model of talent allocation that is consistent with the trade-offs motivated by the empirical evidence. Then, Section 1.5 provides a description of the estimation and presents the quantitative results of counterfactual policies. Section 1.6 considers extensions and robustness checks, and Section 2.5 concludes.

1.2 Empirical Evidence

In this section I document empirical evidence from India. I use survey data to present facts of rural and urban regions, splitting the latter into formal locations and slums. The data come from the India Human Development Survey-II, 2011-2012 (IHDS II). This is a nationally representative survey of 42,152 households in 1,420 villages and 1,042 urban neighborhoods. The information in these surveys includes individual data on education, labor income, and hours worked; as well as household data on expenditures and dwelling characteristics. These data have the advantage of allowing me to identify slums in urban regions and use a rich set of individual information to compare relevant outcomes that characterize the main trade-offs between rural and urban areas.

The classification of households as urban or rural in IHDS II is based on the Census of India 2011. According to the latter, urban areas are towns where local government bodies are situated, and places with: (i) a minimum population of 5,000; (ii) at least 75% of male working population engaged in non-agricultural activities; and (iii) a population density of at least 400 persons per sq.km. Therefore, urban regions include administrative centers of local governments, such as municipal corporations, and areas

the spatial allocation of labor and aggregate labor productivity in the US.

with a high population density where non-agricultural activities predominate. Rural villages are areas that do not satisfy the previous definition. In 2011, urban population in India represented 31% of total population.

I use IHDS II data to identify slums based on the definition of the United Nations (UN).[10] I define slums as urban households that: (i) do not have access to piped water or protected source; (ii) do not have a private toilet or dispose waste in open fields; (iii) have roof/walls made of thatch, mud or grass; or (iv) do not have proof of residence such as an electric bill or rent agreement.[11] Formal urban houses are dwellings located in urban areas that are not classified as slums. Adapting the UN definition to the available data is a standard practice in papers about slums or informal urban settlements. Note that slums are defined at the household level and this type of dwellings tend to be clustered in neighborhoods within a city.

1.2.1 Education and Labor Income

In this section I document economic opportunity gaps across location types: rural, slums, and formal urban housing. The evidence presented has two patterns: individuals living in slums have lower outcomes in terms education and labor income within urban regions, but they are more educated and have higher income than those in rural regions.

First, Figure 1.2 shows differences in educational attainment for men that are

[10] An urban household lives in a slum if it *lacks at least one* of the following: improved source of water, improved sanitation facilities, housing durability, tenure security, or sufficient living area.

[11] IHDS II do not provide good data to measure sufficient living space. Using this definition, I obtain a share of slums population in urban regions equal to 36%. This share is higher than the one observed in the closest year to the surveys in UN data: 29% in 2009. If I do not take into account tenure security, I obtain a similar share to the UN; however, the empirical evidence is not very sensitive to this distinction and the lack of property rights is a key feature of slums.

Figure 1.2: Educational attainment across locations

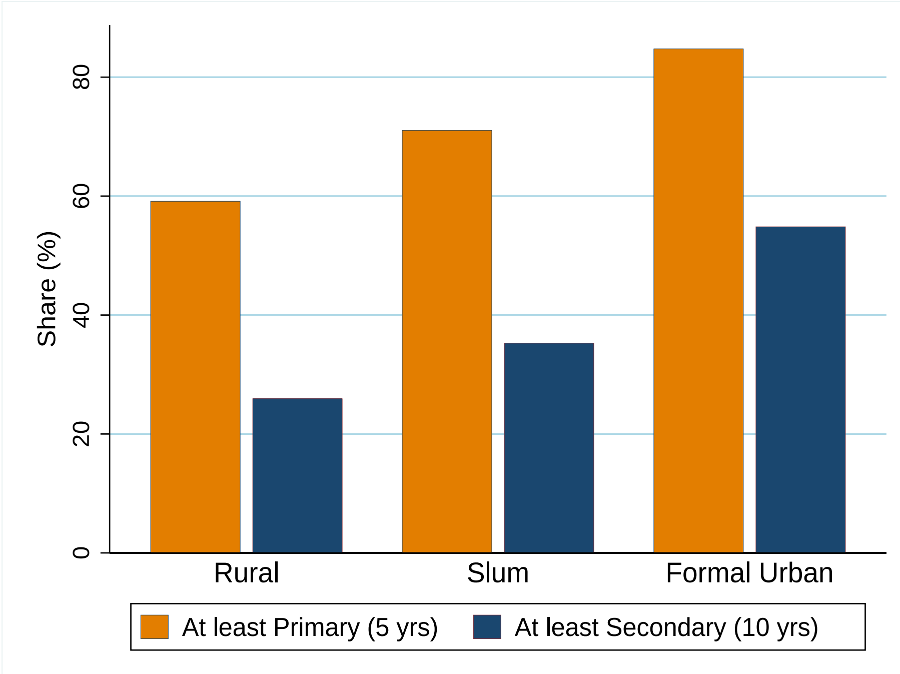

Notes: Shares of male population 25 years or older. Source: Author estimates using IHDS II 2012.

25 years or older. I make this sample restriction to focus on individuals who have likely finished their education. The results show that attainment is higher in both types of urban locations. The share of individuals with at least primary school is 11 percentage points higher in slums than rural regions, and 24 percentage points higher in formal urban areas. Similarly, the share of individuals with at least secondary school is higher in slums than in rural regions, although the difference is smaller in this case. These results show that educational gaps between individuals living in slums and those living in formal urban housing are especially large for higher levels of education, which is consistent with the notion that low-skilled individuals tend to live in slums.

Another key difference between rural and urban regions are labor market outcomes. Table 1.1 presents the gaps in hourly wages between urban locations and

Table 1.1: Hourly Wage Relative to Rural Regions

	Raw Gap	Adjusted Gap for Education
Slum	1.41	1.32
Formal urban	1.79	1.51
Observations	34,337	34,299

Notes: Based on working age male population (15-65 years). Controlling for age in both cases.

Source: Author estimates using IHDS II 2012.

rural regions for working age (15-65 years) men who worked at least part time. These gaps are obtained by running a regression of individual log-wages on location type, choosing rural areas as the omitted group. The gaps reported are equal to $\exp(\beta_\ell)$, where β_ℓ is the regression coefficient of each location type. The raw wage gap between slums and rural regions is equal to 1.4, which means that hourly wages are 40% higher in slums than in rural regions, while the raw gap between formal urban areas and rural regions is equal to 1.8. These large wage gaps are usually interpreted as suggestive evidence of barriers to labor mobility, or as a reflection of skill sorting across locations. To account for selection of skilled individuals, I estimate the wage gaps controlling for educational attainment. The adjusted gaps are smaller but still large and, while there can be selection based on unobservable human capital, they suggest that there are barriers preventing individuals from exploiting opportunities across space.

Next, I relate the findings on education levels and wages by exploiting the micro data to compare returns to education by location type. To do so, I estimate the following log-wage regression:

$$\log(W_{i\ell}) = \beta_\ell D_\ell + \eta_j e_i + \alpha f(exp_{i\ell}) + \epsilon_{i\ell}. \tag{1.1}$$

Where $W_{i\ell}$ is the hourly wage of individual i living in location ℓ, D_ℓ is a location dummy, $f(exp_i)$ is a quartic in experience at the place of residence, e_i are education years completed by individual i, and $\epsilon_{i\ell}$ is an i.i.d. error with zero mean. In the same spirit as Lagakos *et al.* (2018b), I define potential experience as $exp_{i\ell} = \min\{\text{age} - e_i - 6, \text{ age} - 16, \text{ years living in } \ell\}$.[12] This definition implies that individuals start working in their current location at age 16, when they finish school, or when they move to this place, whichever comes last. The results presented in Table 1.2 show that returns to education are higher in urban regions, especially in formal urban areas. Thus, not only are wages higher in urban regions for a given level of education, but individuals with high educational attainment gain more from working in cities. One interpretation of these results is that occupations in cities are more skill-intensive in comparison to those in rural areas.[13]

The evidence on education levels and labor income implies that, despite the unhygienic conditions and lack of basic public services, slums tend to have residents with a higher level of education than those in rural regions, and they give access to labor markets with higher returns. That said, these results also stress that outcomes are worse in slums within urban regions. One of the reasons previous papers have argued that slums have characteristics of poverty traps is due to the fact that human capital is low in those places. Having access to formal housing is crucial to fully exploit labor market opportunities in urban regions and, as explained below, regulations limiting the supply of formal housing might be keeping too many people living in slums in

[12] The results are very similar if I use total experience instead of experience at the place of residence.

[13] Young (2013) formalizes this argument to explain urban-rural wage gaps.

Table 1.2: Returns to Education by Location

Location	Mincer return
Rural	4.3%
Slum	5.1%
Formal urban	6.5%
Observations	34,074

Notes: Based on working age male population (15-65 years). Source: Author estimates using IHDS II 2012.

India.

1.2.2 Housing Rents

The previous section focused on the benefits of urban regions. However, living in cities generally increases the costs of living for rural migrants; in particular, housing costs are higher in urban regions. Based on reports from National Sample Surveys (NSS) in India - Housing Conditions Round 58th in 2002, Table 1.3 presents the gaps in monthly house rents for different levels of floor area.[14] According to these data, the average monthly rent in slums is 1.2 times higher than in rural regions, while the average monthly rent of a formal urban house is 2.4 times higher than in the countryside. Furthermore, monthly rents are higher in urban locations for every level

[14]Compared with more recent rounds of National Sample Surveys, round 58th reports monthly rents of slums/squatters in urban regions. In these surveys, 92% of rural households and 60% of urban households reported owning their house, so the rents reported include imputed rents to non-rental households.

Table 1.3: Monthly House Rent Relative to Rural Regions

	Floor area (sq.m.)				
Location	20-30	30-40	40-50	75-100	All
Slum	1.6	1.3	1.9	1.5	1.2
Formal urban	2.0	2.1	2.4	4.2	2.4

Source: National Sample Survey 58th Round, 2002. Housing Condition in India.

of floor area. In comparison to the average dwelling in rural regions, which has a floor area of 30 to 40 sq.m., housing rents are 30% higher in slums and more than two times higher in formal urban areas. These housing rent gaps are larger for middle size dwellings and extremely large for big houses. The average floor area in slums is almost half of the average floor area of a dwelling in both rural regions and formal urban areas, which explains the smaller size of the average rent gap compared to the gaps conditional on floor area.

I use data from NSS to compare housing rents because these surveys were made with the goal of comparing housing conditions in India, and the estimates of monthly rents include imputed rents on non-rental households based on similar dwellings in a given locality. That said, using data from IHDS II and focusing on rental households that report positive rents, I find similar results for the gap in average monthly rent between formal urban areas and rural regions, 2.4, and a moderately larger gap between slums and rural regions, 1.9.[15]

[15] In IHDS II data, 99% of rural households reported owning their house, as opposed to renting or being an accommodation provided by an employer. This share is equal to 84% in urban regions.

A possible concern is that many slum dwellers in India do not have tenure security, which is evident from the common forced evictions. The latter means that payments of rents are not necessarily enforced through legal rental agreements. If many slums dwellers are occupying land illegally, then their relevant housing cost might not only be the payments done in the informal rental market, but also other types of protection payments. Overall, the evidence shows that housing costs are higher in urban regions, even in slums. Thus, individuals face a trade-off between better labor market opportunities and higher costs of living in cities.

1.3 Urban Policies in India

This section provides a description of urban policies in India. This country is considered a notable case of inefficient urban regulations. Examples of such policies include building-height limits which restrict the available floor space per capita; rent controls that keep the rental market underdeveloped; urban land ceilings which constrain the holdings of land in the private sector; and regulations that restrict the ability to transfer property. These regulations limit the supply of formal urban housing in India and raise the price of houses.

A common policy to control density growth in cities by restricting building heights are legal limits to the floor area ratio (FAR). The FAR is defined as the total area of floor space contained in a building divided by the area of its lot. According to a study from the World Bank in 2013, the FAR in the main cities of India is very low by international standards. For example, Delhi and Mumbai have a FAR of approximately 1.3 in the city center; by contrast, San Francisco and New York have a FAR of 9 and 15, respectively. These limits are not equal to the total number of floors that a building can have. For example, buildings can cover only a section of their lot, while the remaining lot area is used for open spaces or plazas. However, a low

FAR constrains the supply of urban formal housing and the capacity to accommodate population with the available supply of land.

In addition, approval processes to build new houses are lengthy and costly in India. According to 2018 data from the Word Bank *Doing Business* reports, construction permits necessary to build a warehouse cost 23% of the value of the building. The latter compares with 8% in China, 10% in South Africa, and 1% in the United States.[16] A study from KPMG in 2014 shows that the approval process for housing development in India takes between two and three years and increases construction costs by 20 to 30 per cent. These regulations increase the incentives to build high-value real estate to cover approvals costs, instead of affordable housing. Furthermore, taxes and fees in India account for 30 to 35 per cent of housing development cost. This means that approvals costs and regulation fees together account for more than 50% of housing construction costs. As a comparison, in the US government regulations at different levels account for 24% of the final price of a new house according to Emrath (2016).

With respect to slums, governments in India have carried on massive forced demolitions in recent decades. For example, between 2004 and 2005, at least 90,000 slums houses were demolished in Mumbai, which represented around 8% of the slum population in the city; similarly, not less than 45,000 slums houses were demolished in Delhi from 2004 to 2007. Governments have tried to recover valuable land occupied by slums in the main cities of the country, and a significant amount of slums located in central urban areas have been demolished. In most cases, only a small fraction of evicted dwellers are offered subsidized relocation, mainly to peripheral areas of the city: 20 to 40 kilometers from the original location.[17] To put this into

[16]These costs include fees associated with obtaining land use approvals, construction inspections, utility connections, and registration of the warehouse.

[17]See Bhan (2009), Kumar (2010), Dupont (2008) and Diwakar and Peter (2016).

Figure 1.3: Share of Population Living in Slums in Mumbai (%)

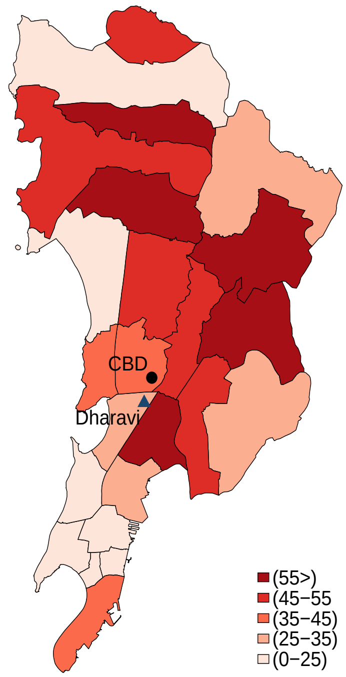

Notes: Division by Ward (subdistrict). CBD: Central Business District (Bandra Kundra Complex). Dharavi is one of the biggest slums in the world.

Source: Development Plan for Greater Mumbai 2014-2034. Municipal Corporation of Greater Mumbai.

context, according to the 2011 Census, 78% of urban commuters in India work within 10 kilometers from their home. In Mumbai, the average one-way commute is just 5.3 kilometers for all workers and 3.9 kilometers for the poor (Baker *et al.* (2005)).

Therefore, it is not surprising that slums dwellers displaced to outskirts of the cities tend to lose their livelihoods. A study from slums evictions and resettlements in the city of Chennai revealed that 80% of evicted dwellers lost their employment after being relocated. The forced eviction of slum dwellers in central urban areas fails to recognize that such slums are formed because households want to live close to the most productive areas of the city. To see this, Figure 1.3 shows the distribution of slum population across Mumbai subdistricts. Slums are located in almost every area of the city and, in particular, they are located close to the business districts in the center and south of the city. In fact, one of the biggest slums in the world, Dharavi, is located just 5 kilometers away from the main Central Business District (CBD) of Mumbai. This case exemplifies the fact that slums occupy valuable land in the most attractive urban locations, which partially explains the incentives of governments to evict their dwellers and allocate that land to the formal sector.

1.4 Model

This section introduces a static general equilibrium model with spatial sorting. There is a mass one of individuals in the economy and three regions indexed by $x \in \{c, p, r\}$: urban center (c), urban periphery (p), and rural (r). Each urban region is divided in two types of neighborhoods: slums (s) and formal urban (f). To simplify the notation, I will use $\ell \in \{cf, cs, pf, ps, r\}$ to denote the five possible residential locations. There are no trade costs and individuals can freely choose their place of residence. Location characteristics, such as productivity and amenities, are exogenous in the model. Section 1.6.3 relaxes this assumption and considers the quantitative implications of introducing endogenous agglomeration forces.

The model builds on the quantitative spatial framework of Bryan and Morten (2018), which itself relates to Ahlfeldt *et al.* (2015), Monte *et al.* (2018), and Hsieh *et al.* (2018). The model also adapts ideas from the macro-development literature on sectoral productivity gaps and human capital, particularly Gollin *et al.* (2014) and Herrendorf and Schoellman (2018).

1.4.1 Preferences and Individual Problem

Every individual i living in location ℓ has preferences over consumption $c_{i\ell}$ and housing $h_{i\ell}$ given by

$$U_{i\ell} = \left(\frac{c_{i\ell}}{1-\alpha}\right)^{1-\alpha} \left(\frac{h_{i\ell}}{\alpha}\right)^{\alpha} \mu_\ell, \tag{1.2}$$

where α is the expenditure share on housing and μ_ℓ denotes amenities in location ℓ. The value of these amenities is exogenous and governs the compensating wage differentials across locations. Individuals living in location ℓ maximize utility by choosing consumption and housing subject to their budget constraint: $c_{i\ell} + r_\ell h_{i\ell} = W_{i\ell}$, where the consumption good has been used as the numeraire; r_ℓ is the rental price of housing in location ℓ; and $W_{i\ell}$ denotes labor earnings of individual i in location ℓ. Individuals know the labor income they would have in each location and choose where to live by comparing welfare across places. As described below, this comparison depends on characteristics that make locations more or less attractive, namely housing rents, amenities, and income, which in turn depends on individual comparative advantage.

1.4.2 Human Capital and Labor Income

Labor income is a function of human capital, which depends on individual talent and education. To be specific, I build on the macro-development literature and define individual human capital (efficiency units) as

$$z_{i\ell} = \nu_{i\ell} \exp(e_i \eta_\ell), \tag{1.3}$$

where $e_i \in \{0, 1, ..., 16\}$ are education years of individual i, which are valued at rate η_ℓ in location ℓ. Parameter η_ℓ governs the elasticity between education years and human capital, and captures differences in the demand for skilled labor in each location. I assume that individual education and schooling returns are exogenous. The share of individuals in the economy with e years of education completed is denoted by ζ_e. Additionally, every individual is endowed with an idiosyncratic talent draw $\nu_{i\ell}$ for each possible residential location. The idea is that individuals have different talents and the type of jobs in each location may be different. Most individuals in rural regions work in agriculture and slums dwellers tend to concentrate in elementary occupations such as rag-picking or low-value manufacturing; similarly, industries in the urban periphery are not the same as those in the main business areas of a city. For example, banking and financial services tend to be located in the CBD of a city.

I follow Bryan and Morten (2018) and assume that talent is drawn from a multivariate Fréchet distribution

$$F(\nu_{cf}, ..., \nu_r) = \exp\left(- \left[\sum_{\ell} \nu_{\ell}^{-\frac{\tilde{\theta}}{1-\rho}} \right]^{1-\rho} \right), \tag{1.4}$$

where the shape parameter $\tilde{\theta} > 0$ governs the dispersion of talent draws. A higher value of $\tilde{\theta}$ implies lower variation in individual talent. Parameter $\rho \in [0, 1)$ governs the correlation of skills between locations.[18] As ρ get close to one, the distribution approaches the case of unidimensional talent. To characterize the rest of the model it is useful to define a shape parameter adjusted for the correlation between productivity draws: $\theta = \tilde{\theta}/(1-\rho)$. The importance of sorting in the model is stronger when there is high variation in individual talent or when there is a low correlation in talent across locations.

[18] This is a Gumbel Copula distribution. Its main advantage over other distributions is that it leads to tractable analytical solutions of the model. About the marginal Fréchet, Lagakos and Waugh (2013) argue that the talent draw in location ℓ can be thought as the task or occupation, among a large set of options, that maximizes individual labor income in that place. To the extent that the number and types of industries to which an individual has access are different across locations, individuals have a different draw in each place.

Individuals are endowed with one unit of time to supply inelastically in the labor market they face at their place of residence. There is no commuting across locations, even within the city. This is a simplifying assumption based on the observation that commuting distances are extremely short in India.[19] Then, individual labor earnings in each location are defined as

$$W_{i\ell} = \omega_\ell \, z_{i\ell}, \tag{1.5}$$

where ω_ℓ is the wage per efficiency unit in location ℓ.

1.4.3 Production Technologies

There is a final consumption good supplied by competitive producers in urban and rural regions and traded without costs, but regions differ in the technology they use. In urban regions, the final good is produced using efficiency units of labor from every urban location according to

$$Y_u = \Big[\sum_{\ell \neq r} \big(A_\ell Z_\ell \big)^{\frac{\sigma-1}{\sigma}} \Big]^{\frac{\sigma}{\sigma-1}}, \tag{1.6}$$

where σ is the elasticity of substitution across locations and Z_ℓ is the demand for efficiency units of labor from urban location ℓ. In practice, I will focus on finite values of σ greater than one, which implies that labor units from different urban locations are not perfect substitutes but also that none of them is essential. Note that the specification in (6) implies that each urban location produces an intermediate good given by $Y_\ell = A_\ell Z_\ell$. Labor productivity is potentially different in each location and

[19] In Section 1.6.1 I relax this assumption to assess the quantitative importance of allowing commuting from the urban periphery to the center. The baseline model captures access to commuting in a reduced form way through differences in local productivity.

A_ℓ is a parameter that captures spatial differences in industry mix, infrastructure, agglomeration forces, and access to non-tradable goods. The aggregation of labor to produce the final good in the city is the link across urban locations in the model, and the implicit assumption is that each location produces a differentiated good consumed by individuals in every place of the city. On the other hand, the final good in the rural region is produced using local labor according to

$$Y_r = A_r\, Z_r, \tag{1.7}$$

where A_r is a parameter of rural labor productivity.

1.4.4 Housing Production and Land

The available supply of land varies across space. First, I assume there is an infinite supply of land in rural regions and housing supply is perfectly elastic (i.e. there is no land congestion). Therefore, rural housing prices are constant: $r_r \equiv C_r$. Parameter C_r captures the marginal cost of producing a unit of housing in the rural region. Second, in each urban region, the amount of residential land is fixed, denoted by $\overline{L}_x$, and owned by landlords who only consume goods in their location.[20] Moreover, a fraction of land ϕ_x is occupied by slums in each urban region, for which landlords do not receive any payment. This means that $1 - \phi_x$ is the fraction of land available for formal development. This fraction is exogenous and captures the results of government policies and enforcement choices towards slums.

In both urban regions, formal housing is produced using land L_ℓ and intermediate inputs M_ℓ (same as consumption good) according to $H_{xf} = (L_{xf})^\psi\,(M_{xf})^{1-\psi}$, where $\psi \in (0,1)$ governs the intensity of land in housing production. Formal developers

[20] As noted by Monte *et al.* (2018), this type of assumption allows the incorporation of general equilibrium effects from changes in the value of land, without introducing an externality in the location choice of individuals from the redistribution of local land rents.

maximize profits taking housing rental prices r_ℓ and residential land prices p_x as given, that is,

$$\max_{M_{xf},L_{xf}} \quad (1-\tau)\, r_{xf}\, (L_{xf})^{\psi}(M_{xf})^{1-\psi} \;-\; M_{xf} -\; p_x L_{xf}, \tag{1.8}$$

where $\tau > 0$ is a distortion that developers take as given and captures regulations such as FAR limits, rent controls, and approval costs. The fraction of housing output that is "taxed", $\tau\, r_{xf}\, H_{xf}$, is thrown away. In equilibrium, the demand for land by formal developers has to be equal to the available supply: $(1-\phi_x)\overline{L}_x$. On the other hand, housing production in slums is given by

$$H_{xs} = C_s\, \phi_x\, \overline{L}_x, \tag{1.9}$$

where C_s is a technology parameter that governs the density of housing production in slums. The difference in housing technologies is meant to reflect the fact that production of housing in slums is more intensive in the use of land and a fixed amount of housing space can be constructed per unit of land: dwellings in slums are usually homogeneous, small one-story houses with one or two rooms at most. In comparison, the formal housing sector has the possibility to substitute intermediate inputs for land and produce different amounts of housing space per unit of land (e.g. houses or apartment buildings). Note that the supply of housing in slums is fixed in each urban region, but housing demand is determined endogenously. In equilibrium, the housing rent of slums has to clear the housing market in these locations. This rent is paid by individuals to owners who only consume goods in their location.

1.4.5 Equilibrium

I now provide a description of allocations and prices in general equilibrium. First, indirect utility of individuals with e years of education and $\nu_{i\ell}$ units of talent in location ℓ is

$$V_{i\ell|e} = \mu_x \, \nu_{i\ell} \, \omega_\ell \exp(e \, \eta_\ell) \, r_\ell^{-\alpha}. \tag{1.10}$$

Given the distribution of individual talent and expression (10) for every location, it can be shown that the share of individuals living in location j conditional on having e years of education is

$$\pi_{j|e} = \frac{\left(\omega_j \exp(e \, \eta_j) \, \mu_j \, r_j^{-\alpha}\right)^\theta}{\sum_\ell \left(\omega_\ell \exp(e \, \eta_\ell) \, \mu_\ell \, r_\ell^{-\alpha}\right)^\theta}. \tag{1.11}$$

Equation (11) says that the spatial allocation of talent depends on relative housing costs, local productivity, returns to education, and amenities. Housing prices act as congestion cost in each urban location due to the limited supply of land; moreover, formal housing prices in urban regions are relatively high due to regulation distortions. On the other hand, individuals are attracted to more productive locations and to places with a higher value of amenities. Locations with higher returns to education, where jobs use skills more intensively, attract individuals with more years of education because the function $\exp(e \, \eta_\ell)$ implies complementarity between e and η_ℓ. This means that the gain from living in a location with higher returns to schooling is larger for a person with more years of education. In equilibrium, the share of individuals who choose to live in location ℓ is given by $\pi_\ell = \sum_{e=0}^{16} \pi_{k|e} \, \zeta_e$, and market clearing implies $\sum_\ell \pi_\ell$=1.

Additionally, the supply of housing must be equal to the demand in every location, that is,

$$H_\ell \; = \; \alpha \; \frac{I_\ell}{r_\ell} \tag{1.12}$$

where $I_\ell = \sum_{e=0}^{16} \left[\int_{i \in \Omega_{\ell|e}} \omega_\ell \, z_{i\ell}(e) \, dF_i \right] \zeta_e$ is total labor income in location ℓ and $\Omega_{\ell|e}$ denotes the set of individuals with e years of education who choose to live in location ℓ. Furthermore, the price of land in each urban region has to clear the market of land, thus,

$$p_x = \psi \, \frac{(1-\tau)\, r_{xf}\, H_{xf}}{(1-\phi_x)\, \overline{L}_x}. \tag{1.13}$$

This expression implies that, holding all else constant, the price of land increases with the amount of land occupied by slums because the amount available for formal development is smaller. Also, attractive locations have a higher demand for housing, which raises the value of land in those places. To see how distortions affect the supply of formal urban housing it is useful to focus on the optimal supply of housing per unit of land: $h_{xf} = H_{xf}/L_{xf}$. In this case, the density of housing production in the formal sector is given by

$$h_{xf} = \left[(1-\tau)\,(1-\psi)\, r_{xf}\,\right]^{\frac{(1-\psi)}{\psi}}. \tag{1.14}$$

Then, a higher value of τ implies that a lower amount of formal housing can be developed with the available supply of land.

In equilibrium, labor demand in each location must equal labor supply: $Z_\ell = \sum_{e=0}^{16} \left[\int_{i\in\Omega_{\ell|e}} z_{i\ell}(e)\, dF_i\right]\zeta_e$; and workers are paid a wage per efficiency unit of labor,

$$\omega_r = A_R \qquad \text{and} \qquad \omega_\ell = \left(A_\ell\right)^{\frac{\sigma-1}{\sigma}} \left(\frac{Y_u}{Z_\ell}\right)^{1/\sigma}, \quad \ell \neq r. \tag{1.15}$$

Labor demand in a particular urban location has a negative slope and increases with local productivity A_ℓ. Note that ω_ℓ, for $\ell \neq r$, depends on labor demand and local productivity of every urban location through Y_u. The model implies that the gap in average wages between location ℓ and j, conditional on education years, is given by

$$\overline{W}_{\ell|e} \Big/ \overline{W}_{j|e} = \frac{\omega_\ell}{\omega_j}\, \frac{\exp(e\,\eta_\ell)}{\exp(e\,\eta_j)} \left(\frac{\pi_{\ell|e}}{\pi_{j|e}}\right)^{-\frac{1}{\theta}}. \tag{1.16}$$

According to expression (16), conditional wage gaps between location pairs reflect differences in local productivity captured by ω_ℓ; differences in returns to education given by η_ℓ; and the sorting effect captured by conditional population shares $\pi_{\ell|e}$. The sorting effect refers to the fact that there is a negative relationship between average wage and population share in a given location because the marginal worker has a relatively low productivity in that place. That is, as population increases in a particular location, new dwellers are added from the lower part of the distribution of talent. The strength of the sorting effect on wages depends on the variation of talent draws within locations and the correlation of draws across locations: $\theta = \tilde{\theta}/(1-\rho)$.[21] When the correlation is high, the sorting effect is weaker since individuals tend to have similar levels of talent in every location. Finally, the unconditional wage gap between locations pairs can be written as

$$\overline{W}_\ell \Big/ \overline{W}_j = \frac{\omega_\ell}{\omega_j} \; \frac{\sum_{e=0}^{16} \exp(e\,\eta_\ell)\, \pi_{\ell|e}^{1-1/\theta}\, \zeta_e}{\sum_{e=0}^{16} \exp(e\,\eta_j)\, \pi_{j|e}^{1-1/\theta}\, \zeta_e} \; \frac{\pi_{j|e}}{\pi_{\ell|e}}. \tag{1.17}$$

Thus, differences in average labor earnings depend on the level of education in each place, which is captured by the term $\zeta_e\, \pi_{\ell|e}$. Education returns in the model have a similar role in explaining wage gaps and spatial sorting as differences in skill intensities in Young (2013), given that a higher return in location ℓ attracts more educated individuals to that place; moreover, selection is key to explain wage gaps across locations in the model, though, similar to the argument in Bryan and Morten (2018), amenities and housing prices (regulations) introduce possible frictions to the spatial allocation of labor.

[21] This implies that absolute and comparative advantage are aligned in the model. See Young (2014).

1.5 Quantitative Analysis

This section presents the estimation of the model using the IHDS II and other sources of aggregate data from India. The main parts of the estimation can be summarized as follows. First, I estimate a subset of parameters using informative relationships derived from the model. Then, I implement an internal method of moments to recover unobserved location characteristics using data on wages, housing rents, and population shares. Finally, I validate the model by comparing non-targeted moments in the model and the data. I then use the estimated model to compare the effects of evicting slum dwellers from central urban areas and reallocating the land to the formal sector with the effects of reducing formal housing distortions. I focus on changes in the share of urban population, total output (which in this case is the same as output per capita), and individual welfare computed by aggregating equation (10). To provide intuition, I also quantify the effects of these policies on different population groups.

1.5.1 Calibration

First, productivity in the rural region A_r is normalized to one. Next, I restrict amenities in urban regions by assuming they have the same value in both neighborhood types within each urban region, i.e, $\mu_{xf} = \mu_{xs}$. The idea is that amenities are a key difference between urban and rural regions: access to urban services versus exposure to pollution and loss of rural networks. Differences between slums and formal urban areas are captured by local productivity, education returns, and housing technologies. Since the value of amenities governs compensating wage differentials, I normalize central urban amenities by setting μ_c equal to one. I follow Allen and Arkolakis (2014) and choose σ equal to 9. This value implies a high level of substitutability between urban locations.[22] Given that the model assumes a distorted housing market in India, I take the U.S. as an efficient benchmark to determine the value of housing expenditure share and the intensity of land in housing construction. I set α equal to 0.25 following Davis and Ortalo-Magne (2011), and ψ equal to 0.15 based on Epple *et al.* (2010). I determine the value of distortions in the formal housing sector τ using a report from KPMG (2014) on housing development in India, according to which taxes, fees and approval costs account for approximately 55% of the housing development costs. I take this as a measure of distortions in formal housing.

To obtain the supply of residential land in each urban region, I focus on the case of Mumbai. I use information from the Development Plan for Greater Mumbai 2014 - 2034 (DPGM) to obtain the amount of residential land in each subdistrict of the city. Then, I define the urban periphery as those subdistricts that are in the outskirts of the city: 20 or more kilometers away from the closest CBD.[23] This cutoff is consistent with the location of resettlements offered to evicted slums dwellers in the main cities of India.[24] Based on these definitions, I set $\overline{L}_p/\overline{L}_c$ equal to 0.43. This ratio shows that the periphery is a smaller area than the central region of the city; however, the price of land in the model will reflect the fact that most individuals want to live in the center. Additionally, based on information from the DPGM, I set ϕ_x equal to 0.33 for both urban regions, which is the share of land occupied by slums in both the center and periphery of Mumbai.

Lastly, I calculate the population share with e years of education ζ_e using data from the IHDS II. In this case I restrict the sample to focus on working age male individuals who have potentially finished their education (between 25 and 65 years

[22] Section 1.6.4 presents robustness checks with lower elasticity values.

[23] This is based on subdistricts that are part of Mumbai City and Suburban Mumbai. A map is included in Appendix A.

[24] See Dupont (2008) and Diwakar and Peter (2016).

old). Individuals are assigned to one of 17 possible categories that go from zero years of schooling to graduate education (above bachelors degree or more than 15 years of education). The five highest shares are the following: $\zeta_0 = 22.1\%$, $\zeta_{10} = 12.3\%$, $\zeta_9 = 9.5\%$, $\zeta_8 = 8.3\%$, and $\zeta_{12} = 8.1\%$.

Returns to Education

The model implies a relationship between log-wages and education years similar to a Mincer regression:

$$\log(W_{i\ell}) = \log(\omega_\ell) + \eta_\ell e_i + \log(\nu_{i\ell}) \tag{1.18}$$

Based on this expression, I use data from IHDS II to estimate location-specific returns to education η_ℓ in the same way as Section 1.2.1, but in this case I distinguish between central and peripheral urban regions. IHDS II does not provide the geographical location of households within cities; however, for rural villages, it reports the distance to the main city in the district. I use this information to define the urban periphery as those villages that are between zero and nine kilometers from the main urban center of their district. This cutoff is determined by the relationship between average wages and distance. Villages within nine kilometers of the urban center have an average wage that is much higher than the rest of the rural areas: around 33% higher. The assumption is that villages located close to main cities have similar characteristics to urban locations in the periphery of the city.

Given this definition of urban periphery and following expression (18), I estimate a Mincer regression of log-wages on education years. To do so, I control for location fixed effects to account for local differences in productivity, and a quartic in experience at the place of residence to account for individual comparative advantage. The

Table 1.4: Location-specific Mincer Returns

Location	Mincer return (η_ℓ)
Formal center	6.5%
Slum center	5.1%
Formal periphery	6.3%
Slum periphery	4.4%
Rural	4.2%

Notes: Based on male population (15-65 years).

Source: Author estimates using IHDS II 2012.

Mincer return for each location type η_ℓ is reported in Table 1.4. These results show that returns to education are significantly higher in urban regions; in particular, returns are high for individuals living in formal housing even if they are located in the periphery. On the other hand, peripheral slums have similar education returns to rural regions, but returns are higher in central slums and, thus, they attract more educated individuals to live in such locations.

These estimates suggest important differences in the use or value of skills across locations. However, as noted by Herrendorf and Schoellman (2018), differences in location-specific returns to education could be partially capturing selection. That is, urban returns to education could be higher than rural because individuals with high cognitive talent live in cities. To address this possibility, Section 1.6.4 presents a robustness check of the results taking into account this bias. That said, the fact that education has a higher value in urban occupations does not mean that moving to the city necessarily generates large productivity gains, because selection partly offsets those gains in the model, though not through returns to schooling. If most individuals with a strong comparative advantage for urban occupations already live in the city, then productivity gains from increasing urban migration would be relatively small.

Distribution of Talent Parameters

To estimate the parameters that govern the distribution of talent across locations, I use the relationship implied by the model between average wage by location and the share of population living in that location conditional on education years:

$$\log\left(\frac{\overline{W}_{\ell|e}}{\exp(\eta_\ell\, e)}\right) = \log(\omega_\ell) \;-\; \frac{1}{\theta}\log(\pi_{\ell|e}). \tag{1.19}$$

First, using data from IHDS II, I compute average wages for each location-education pair, as well as the corresponding population share. There are a total of 84 observations based on $\ell \in \{cf, cs, pf, ps, r\}$ locations and $e \in \{0, 1, ..., 16\}$ education years. Then, following expression (19) and using the Mincer returns estimated in the previous section, I estimate the elasticity between conditional average wages divided by the value of education in human capital and conditional population shares, controlling for location type. The model implies that this conditional elasticity is a measure of spatial sorting: $1/\theta$. The estimated value of θ is 4.13. Next, properties of the Fréchet distribution imply that the squared coefficient of variation of location wages conditional on education is equal to: $\Gamma\left(1-\frac{2}{\theta(1-\rho)}\right)\Big/\Gamma\left(1-\frac{1}{\theta(1-\rho)}\right)^2-1$. I calculate this moment for all the location-education pairs in the data and use the estimate of θ to find the value of ρ that fits the data. I obtain a value ρ equal to 0.29, which is the Kendall rank coefficient and represents a moderately high correlation across talent draws. The linear correlation is equal to 0.58. These estimates imply that the value of the shape parameter $\tilde{\theta} = \theta(1-\rho)$ is equal to 2.92. The estimated value of this parameter, which governs the variation in productivity draws, is in the same range of other papers that use similar parameterizations in other contexts. For example, Bryan and Morten (2018) estimate values equal to 2.69 and 3.18 across regions in the U.S. and Indonesia, respectively; and Hsieh *et al.* (2018) find values close to 2 across occupation in the U.S.[25]

Internal Calibration

There are eight remaining parameters: A_{cf}, A_{cs}, A_{pf}, A_{ps}, C_s, C_r, μ_p, and μ_r. These parameters are calibrated jointly to match eight moments: (i-iv) population shares by location; (v) the housing rent gap between urban and rural, (vi) the housing rent gap between formal urban and slums, (vii) the wage gap between urban and rural, and (viii) the wage gap between center and periphery. Population shares in the periphery are obtained from the DPGM based on subdistricts located in the periphery of Mumbai, as described in Section 1.5.1; and differences in housing rents are obtained from the reports of the NSS - Housing Conditions Round 58th in 2002. The latter takes into account imputed housing rents of non-rental households based on prevailing rents of similar dwellings in a given locality. Housing rents in the data represent both the price and quantity of housing space consumed by households. To deal with this issue, I match the gap in rental prices of housing in the model to the gap in housing rents for dwellings with a floor area of 30 to 40 sq.m. This range includes the average floor area in both rural and urban regions. The remaining moments are estimated using data from IHDS II.

The results presented in Table 1.5 show that the large wage gap between urban

[25] A possible issue with this estimation is the potential endogeneity between locations wages and population shares. To the extent that this would lead to mismeasurement of the sorting effect in the model, Section 1.6.4 presents robustness checks of the baseline results.

Table 1.5: Joint Calibration

Parameter	Value	Moment	Model	Data
μ_r	1.39	Urban-rural wage gap	1.71	1.71
μ_p	1.22	Center-periphery wage gap	1.33	1.33
C_r	0.76	Urban-rural housing rent gap	2.01	2.01
C_s	0.17	Formal-slum housing rent gap	1.60	1.60
A_{cf}	0.93	Share of population in urban	0.31	0.31
A_{cs}	0.74	Share of urban population in slums	0.36	0.36
A_{pf}	0.46	Share of formal population in periphery	0.22	0.22
A_{ps}	0.36	Share of slum population in periphery	0.26	0.26

and rural regions implies a strong preference for rural amenities; that is, individuals need to be compensated in urban regions for both the higher costs of living and the lower level of amenities. This is consistent with the findings of Lagakos *et al.* (2018a) and Munshi and Rosenzweig (2016). In particular, the latter find that caste-based insurance networks in rural India are an important barrier to urban migration. Also, the wage gap between urban center and periphery implies that individuals living in the center need to be compensated because amenities are lower. This may reflect the fact that individuals are averse to living in the busiest areas of a city because there is more pollution and noise, or it can be capturing differences in costs of non-traded goods. Not surprisingly, the results show that central urban areas are more productive

Table 1.6: Educational Attainment: Model vs. Data

Location	Model	Data
Formal center	8.9	9.6
Slum center	7.5	7.4
Formal periphery	8.7	8.6
Slum periphery	6.8	5.9
Rural	6.6	6.1

Notes: Average education years in the data are calculated based on male population (25 years or older). Education years were not targeted in the calibration of the model.

than the periphery and that is why they attract a larger share of the population. The fact that productivity is the lowest in peripheral slums reflects, among other factors, a lower access to infrastructure in those locations. While the model does not have an explicit notion of distance, a low productivity in the urban periphery captures the costs of commuting to the best jobs in the center.

Model Validation

The moments presented in Table 1.5 are all matched in the baseline economy. To assess the quantitative implications of the model for other relevant moments of the data, Table 1.6 compares average years of education across locations in the model and the data. The model does very well in matching the levels of educational attainment across locations. The gains from living in locations with high returns to schooling

Table 1.7: Average Wage Gaps: Model vs. Data

	Model	Data
Targeted		
Urban / Rural	1.71	1.71
Non-targeted		
Formal / Rural	1.83	1.89
Slums / Rural	1.50	1.44

are larger for educated individuals and, thus, formal urban locations have residents with higher educational attainment. I interpret this in the model as the fact that education is more valuable for occupations in those locations.

Additionally, as shown in Table 1.7, the model generates a ratio of average wages in slums to rural regions equal to 1.50, which is almost the same as the one observed in the data: 1.44; and the ratio of average wages in formal urban to rural regions is equal to 1.83 in the model, compared to 1.89 in the data. Lastly, Figure 1.4 presents the variance of wages by type of urban neighborhood relative to the variance in rural regions. The model does well in replicating the fact that wages variance is higher in urban regions, especially in formal locations.

To summarize, the baseline economy in the model replicates the main empirical facts presented in Section 1.2, namely urban-rural gaps in wages, education levels, and housing rents, plus it is consistent with additional moments of the data. Armed with

Figure 1.4: Baseline Economy: Urban Wages Variance

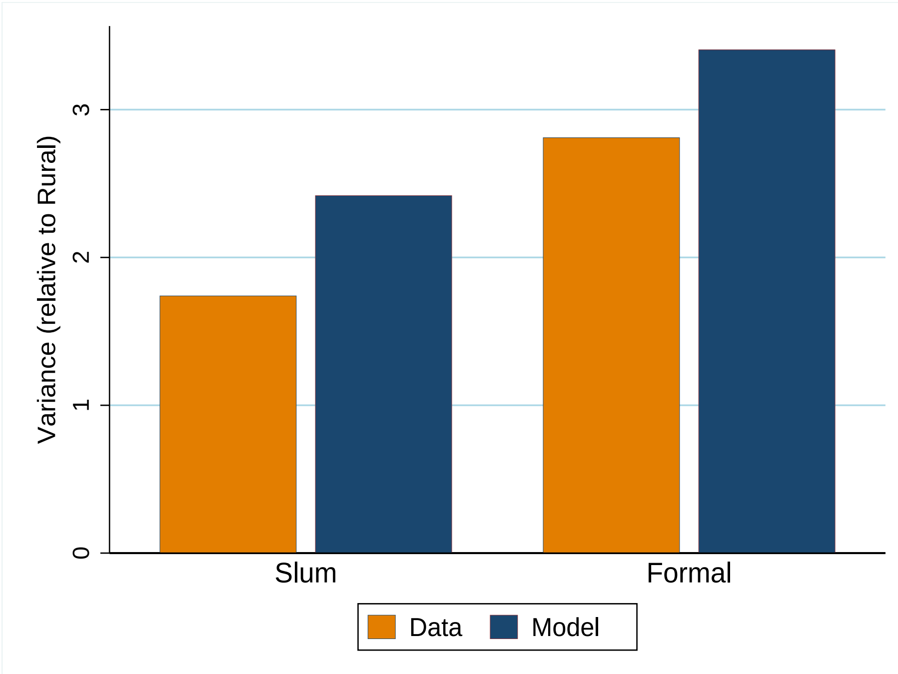

this calibrated model, I now assess the aggregate implications of anti-slum policies and housing distortions in India.

1.5.2 Counterfactual Policies

This section analyzes two types of counterfactual policies: (i) the demolition of central slums, and (ii) the reduction of distortions in formal urban housing. The forced eviction of slums dwellers is a common policy to modernize or sanitize cities by recovering land occupied by squatters and use it for alternative projects, such as commercial development or construction of high-value real estate. I compare the effects of this type of policy with an alternative supported by the model, which is reducing housing distortions in the formal housing sector. The model implies that reducing distortions could increase the density of formal housing production per unit of land, which in turn could decrease the rental price of housing in this sector. To assess the aggregate implications of these policies, I focus on changes in the share of urban population, total output, and welfare. Total output takes into account the consumption of individuals and landlords in urban locations. I also assess the effects of these policies on different groups of the population.

Demolition of Central Slums

First, I evaluate the quantitative effects of demolishing central slums. In the model, this can be introduced as an exogenous shock by reducing the share of land occupied by central slums ϕ_c and reallocating the land to the formal housing sector. To be clear, these experiments represent cases in which slums are destroyed, the land becomes available to be used in the formal sector, and the market determines who can afford the new formal housing. The potential benefit of this policy comes from the fact that it increases the supply of housing in the formal urban center, which is the most productive location in the city. To assess the impact of empirically reasonable cases, I change the value of ϕ_c to induce a reduction in the share of population living in central slums of 10%, 30%, and 50%. Table 1.8 presents the results of these experiments. If 50% of slum dwellers in the central slums are evicted, the share of urban population decreases by 11%, total output goes down by 1.6%, and individuals experience an average welfare loss of 1.2%. These aggregate outcomes seem small, but they are non-trivial considering the fact that the share of total population living in central slums is only 8.2%.

To understand these results, I simulate data and look at the effects on particular groups of the population. I focus on the case in which the share of population living in central slums is 50% lower. First, Table 1.9 presents the new location, welfare change, and average education years of evicted dwellers. These are individuals who choose to live in central slums in the baseline economy and change location when the

Table 1.8: Aggregate Effects of Reducing Central Slums Population

	Percentage change from a		
Variable	10% reduction	30% reduction	50% reduction
Urban population share	-2.0	-6.4	-11.1
Total Output	-0.2	-0.8	-1.6
Individual Welfare	-0.2	-0.7	-1.2

Notes: This table presents the results from reducing the share of population living in central slums by exogenously decreasing the share of residential land occupied by slums in the model (ϕ_c).

share of land ϕ_c is reduced. Over 75% of them move to the rural region, 17% move to formal housing in the urban center, 4% move to formal housing in the periphery, and close to 3% moves to peripheral slums. One way to interpret these results is that individuals have to relocate so far from the main areas of the city that they effectively lose access to urban markets. Alternatively, these results suggest that if central slums are destroyed, then urban migration would be even lower in India than it already is. These results are consistent with the findings of Jedwab and Vollrath (2019) regarding policies that limit urban migration into informal areas. The reason most of the evicted dwellers relocate to the rural region is that formal housing is expensive and the periphery is cheaper but not as productive, plus rental housing prices would increase in the periphery if a large number of people moves there. In contrast, the low price of housing and high value of amenities in the rural region attracts more individuals.

Table 1.9: Relocation, Welfare and Education of Evicted Slums Dwellers

New Location	Location % share	Welfare % change	Education years
Rural	75.8	-7.9	7.0
Formal Center	17.2	-8.0	9.2
Formal Periphery	4.3	-8.5	9.0
Slum Periphery	2.8	-8.4	7.3

Notes: This table presents the new location, average welfare change, and average education years of evicted slum dwellers after a change in ϕ_c induces a 50% reduction in the population share living in central slums.

Table 1.10: Welfare of Infra-marginal Individuals

Location	% Change
Formal Center	-0.3
Slum Center	-16.5
Formal Periphery	-1.1
Slum Periphery	-1.0

Notes: This table presents the average welfare change of infra-marginal individuals (those who do not relocate) after a change in ϕ_c induces a 50% reduction in the population share living in central slums.

In addition, Column 3 of Table 1.9 presents the welfare change of evicted dwellers in their new location. While the magnitude is similar for all locations, the smallest loss is for those who move to the rural region because housing is cheaper and amenities are higher. Those who move to formal housing in the center have access to more productive and skill-intensive jobs, but they are forced to pay a higher price of housing. Even if the newly available land puts a downward pressure on the rental price of formal housing in the center, distortions limit the capacity to accommodate more population and congestion partially offsets the reduction in rents. On the other hand, those who move to locations in the urban periphery have the highest welfare loss because productivity is lower, even if the value of amenities is higher in those places. Note that the welfare loss of evicted dwellers is more than six times higher than the average welfare loss in the economy.

The last column in Table 1.9 shows that the evicted dwellers with lowest level of education move to the rural region and only the most educated ones move to the formal urban sector. The latter is consistent with the argument that some slums dwellers are relatively high-skill individuals who choose to live there because the price of housing is low. That said, given that most individuals in India have a low education level, these results suggest that the relevant welfare margin in terms of spatial allocation of labor is between keeping a high concentration of population in rural areas or allowing individuals to move to urban slums, even if they occupy land illegally.

Next, Table 1.10 presents the welfare effects of this policy experiment on inframarginal individuals, those who do not move from their original location in the baseline economy. Not surprisingly, the welfare loss is very high for those who remain in central slums because housing prices increase due to the lower amount of land available for slums. This could also be interpreted as increasing expenditures on protection costs when governments are cracking down on slums. Infra-marginal individuals in the formal center have a low welfare loss because having additional land for formal development reduces the rental price of housing in this sector. Finally, the welfare loss of infra-individuals in the periphery is primarily due to the fall in urban output that is caused by the displacement of workers from the urban center.

To summarize, demolishing central slums leads to a fall in the urban population share, total output, and welfare. Allocating more land to the formal housing sector in the urban center attracts individuals who were previously on the margin of living in this location or other place; however, given that most evicted dwellers from central slums are displaced from the city, urban labor market loses individuals working in relatively productive locations and, thus, urban output goes down. Individuals do not stay in the periphery because productivity is low in those locations (e.g. lack of urban infrastructure) and housing rents would increase if a large share of individuals moves there. In this experiment, urban housing prices fall (except in central slums) because urban housing demand is now lower, which partially offsets the welfare loss due to lower urban output. Furthermore, selection plays an important role in these results. When individuals leave the city, those who remain tend to have a relatively high productivity. This limits the loss in average productivity from demolishing central slums. Finally, the magnitude of the change in welfare and total output is similar to the results of other papers based on spatial quantitative models, and to the welfare gains that Cavalcanti *et al.* (2018) obtain from counterfactual policies regarding slums in Brazil.

Reducing Distortions in Formal Housing

In this section I assess the quantitative effects of reducing distortions in the formal housing sector. I do this by decreasing τ to a level that is consistent with the U.S.. As described in Section 1.5.1, I set τ equal to 0.55 based on the costs of taxes, fees and approval procedures for housing development in India. A study from the National Association of Home Builders done by Emrath (2016) finds that regulations imposed by government at all levels account for 24% of the final price of a new home in the U.S.. I take this value as the equivalent comparison to India, which implies a reduction in τ of 31 percentage points. Even if housing regulations are costly in the U.S., this seems a reasonable comparison given that housing markets in India are highly distorted.

The results presented in Table 1.11 show that the urban population increases by 20% and the share of the urban population living in slums decreases by 19%. Furthermore, total output and individual welfare increase by 2.4% and 2.6%, respectively. These results suggest that distortions in the formal housing sector in India represent an important barrier to urban migration and keep a large share of the urban population living in slums. Eliminating burdensome regulations would not only increase the urbanization rate and the access to formal housing, but would also raise aggregate labor productivity and welfare.

The reason productivity gains are small, in comparison to the wage gaps across locations observed in the data, is because selection plays an important quantitative role in the model. That is, while individual productivity increases in cities because education has higher returns, new urban migrants have a relatively low level of urban talent. If individuals with a strong comparative advantage for urban occupations already live in the city, then productivity gains from urban migration should not be huge. This is consistent with papers that find limited welfare gains from encouraging urban migration (e.g. Lagakos *et al.* (2018a)).

To analyze the effects of this counterfactual policy on particular groups of the population, Table 1.12 presents the welfare gains of new dwellers and infra-marginal individuals. Those individuals who are now living in formal urban housing experience

Table 1.11: Aggregate Effects of Reducing Formal Housing Distortions to U.S. Level

Variable	% Change
Urban population share	20.5
Slum population (urban share)	-18.8
Total Output	2.4
Individual Welfare	2.6

Note: This table presents the results from increasing formal housing supply in urban regions by reducing τ to 0.24. The baseline value was equal to 0.55.

Table 1.12: Welfare Gains of New Dwellers and Infra-marginal Individuals

Urban housing	New Dwellers	Infra-marginal
Formal Center	4.9	10.1
Slum Center	1.0	1.9
Formal Periphery	4.9	10.1
Slum Periphery	0.9	1.9

Notes: This table presents the average welfare change (%) of new dwellers and infra-marginal individuals (those who do not relocate) after reducing τ to 0.24. The baseline value was equal to 0.55.

Figure 1.5: Interaction of Destroying Central Slums and Housing Distortions

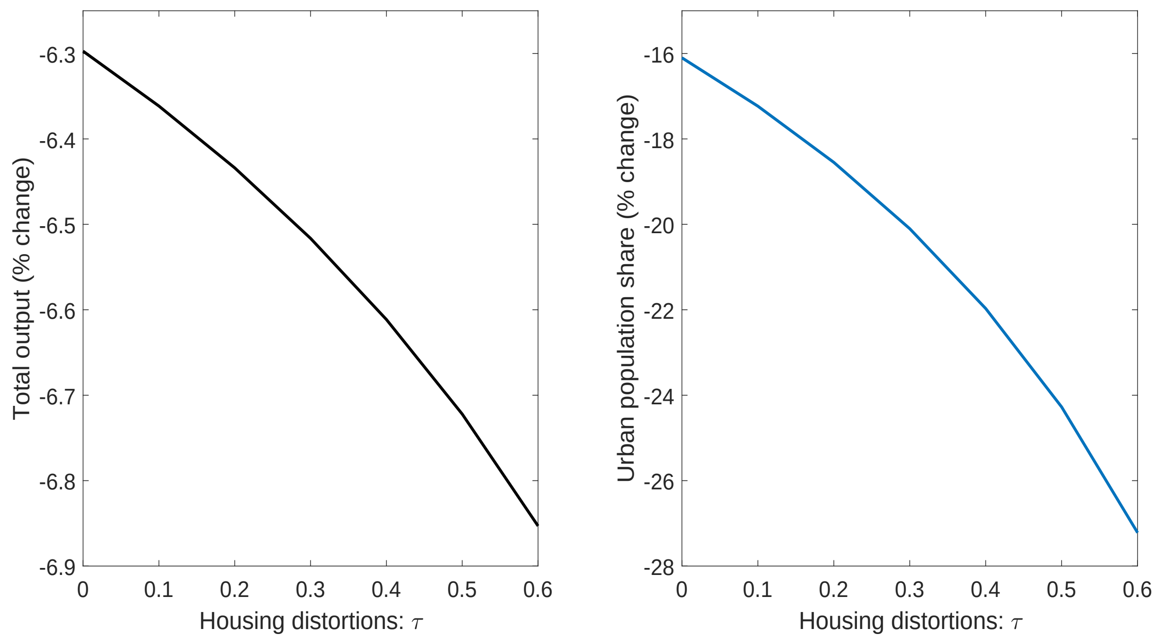

Notes: This figure presents the effects of destroying central slums in the model by setting $\phi_c = 0$.

a welfare gain of 5%, while those who move to slums gain 1%. The welfare gain of infra-marginal dwellers is around two times higher for both types of neighborhoods. These gains are due to the lower price of formal housing and the growth of urban population. The fact that a larger share of the urban population is living in formal housing implies that individual productivity increases through higher education returns and location productivity.

A key motivation to focus on the case of India was the interaction of inefficient urban policies and anti-slum policies. To analyze this, Figure 1.5 presents the effects of destroying central slums ($\phi_c = 0$) on total output and urban population for different values of housing distortions τ. The results show that housing distortions substantially amplify the effects of destroying central slums. The fall in urban population is more than 10 percentage points higher when distortions are similar to the level in

India ($\tau = 0.55$) than in the case with no distortions ($\tau = 0$). Similarly, the loss in total output is significantly lower when there are no distortions in the formal housing sector. The effects on total output are not as dramatic as the effect on the allocation of workers because the city becomes "exclusive" when central slums are destroyed, and only those individuals with a strong comparative advantage remain in urban regions. These results imply that policies aiming to sanitize cities by demolishing slums located in the most valuable locations are more likely to deter urban migration and produce negative welfare effects in the presence of restrictive building regulations.

Implications for India's Urbanization

India has a relatively low urbanization rate for its level of economic development (see Figure 1.1). The model was calibrated to match an urban population share equal to 31% and a share of the urban population in slums equal to 36%. Countries with similar level of GDP per capita have an average urban population share of 46% and an average share of the urban population in slums equal to 43%.[26]

I use the model to see how much of the gap in urban population share between India and countries with similar development can be explained by housing distortions. I compare these results to a case in which the amount of land occupied by slums is increased to match the share of the urban population living in slums in countries with similar levels of development (this is equivalent to a more lenient policy towards slums). The results presented in Table 1.13 show that eliminating distortions raises the share of urban population by 33%, which implies increasing the share from 31% to almost 41%. Thus, holding all else constant, housing distortions alone can account

[26]This is based on World Bank data from 2014. GDP per capita (PPP) in India was \$5,678 and the comparison is made with respect to countries with a GDP per capita between \$4,000 and \$7,000, which had available data on both urban and slums population.

Table 1.13: Two Ways of Increasing Urbanization in India

Variable (% change)	$\tau = 0$	$\uparrow \phi_x$
Total Output	3.6	0.3
Individual Welfare	4.4	0.7
Urban population	32.7	6.7
Slum population (urban share)	-27.5	19.4

Note: This table presents the effects (% change) of eliminating housing distortions and increasing the amount of land in slums, independently. The amount of land is increased to match a share of the urban population living in slums equal to 43%. See text for details.

for a high share of the relatively low urbanization rate in India. Moreover, output and welfare gains are substantially larger in the case where the share of urban population raises by decreasing housing distortions. In the case where slums occupy more land, the share of urban population increases by less than 7%. These results suggest that eliminating distortions in housing markets is a more effective policy to boost urbanization. Allowing slums to occupy more land increases the access to urban markets where labor productivity is higher, but it raises the price of formal housing because there is less land available for that sector.

One question suggested by the previous results is the following. Given the level of distortions and value of location fundamentals, is there an optimal level of land that should be occupied by central slums? The answer is yes. Starting from the baseline economy, increasing the share of land in central slums has small but positive effects on total output and urban population. In fact, total output and urban population display an inverted U-shaped relationship with the share of land occupied by central slums. This is due to fact that the cheap housing sector is expanding in the most productive urban region and workers with a low productivity in rural areas are moving to the city; moreover, the formal housing sector can substitute intermediate inputs for land so the population size in this sector is not affected greatly at first. However, as more land is allocated to slums, substituting intermediate inputs for land becomes increasingly costly and the formal housing sector starts collapsing, causing negative effects on total output and urban population. If all land in the urban center is allocated to slums, the value of these variables falls below their value in the baseline economy. These results imply that governments could increase the urbanization rate by making cities more inclusive in the sense of allowing the expansion of slums. But they also imply that creating a "slum city" leads to a decrease in the urban population share and total output because human capital and productivity are relatively low in those locations.

Demolition of Central Slums and Urban Sprawl

The last experiment I consider is a case where demolishing central slums is combined with an expansion of residential land in the urban periphery. This can be thought of as urban sprawl in the sense that the city is expanding spatially. To analyze this in the model, I exogenously increase the amount of residential land in the periphery so that the share of urban population remains constant after 50% of dwellers are evicted from central slums (by reducing the amount of land they occupy). To simplify the analysis of this counterfactual, I keep the share of land occupied by slums in the periphery constant. The purpose of this experiment is to avoid the fall in urban population caused by the displacement of individuals from central slums, and focus on the effects

Table 1.14: Demolition of Central Slums and Urban Sprawl

Variable (% change)	Baseline	Expanding land supply in periphery
Total Output	-1.61	-1.67
Welfare	-1.21	-0.02

Notes: This table presents the results of reducing the population share in central slums by 50% and, at the same time, increasing the total supply of residential land in the periphery to keep the share or urban population constant.

of allowing individuals to move to or stay in the city but in less productive, peripheral locations.

The comparison of results presented in Table 1.14 shows that the loss in total output is larger than in the baseline case, but the fall in average welfare is almost zero in the case with a higher amount of land in the periphery. The reason for the smaller effect on welfare is that a large increase in peripheral land is required to keep urban population constant (more than 4 times the baseline value) and the new land is more valuable for peripheral slums, the unregulated housing sector that is land intensive, so the rental price of housing in this location falls below the rural level. This is equivalent to a situation in which central slums population is replaced by more population living in peripheral slums, where land does not have a high value and housing is cheap. Thus, the welfare loss of living in peripheral slums where productivity and education returns are low is offset by the gains of paying an extremely low rental price of housing in such locations.

1.6 Extensions and Robustness

This section considers the quantitative implications of possible extensions to the baseline model, as well as robustness exercises. For each of the cases presented below, I repeat the joint calibration to match the moments in Table 1.5. To compare the same policy shock in the baseline and extensions, I focus on the case where all slums in the center are destroyed: $\phi_c = 0$; otherwise, the change in the amount of land occupied my central slums required to generate a specific change in population might vary across cases.

1.6.1 Commuting from Urban Periphery to Center

The baseline model assumes that individuals work where they live (their labor market is determined by their place of residence) based on the fact that commuting is extremely limited in India; however, in cities like Mumbia people do commute from the periphery to the center. In order to assess the quantitative importance of allowing commuting, I modify the model in the following way. I assume that individuals who live in formal housing in the urban periphery can commute and work in the same jobs as individuals living in formal housing in the center, but they face a commuting cost. Similarly, individuals living in periphery slums can commute and access the labor market of center slums subject to a commuting cost. To be clear, I assume that individuals now have two extra draws of location-specific talent $\{\nu_{pcf}, \nu_{pcs}\}$. Each of these draws is the idiosyncratic productivity of living in the periphery and working in the center and, similarly to the baseline, can be interpreted as individual talent for the types of jobs representative of that particular residence-workplace situation.

In addition, I assume that individuals lose a fraction of their income if they commute from the periphery to the center equal to κ_j, $j \in \{f, s\}$. This commuting cost varies by type of housing to reflect the fact that commuting costs represent a different share of income for poor individuals than for rich ones. Then, for example, labor income net of commuting costs of an individual living in the formal periphery and working in the formal center is equal to: $\omega_{cf}\, z_{i,pcf}\,(1-\kappa_f)$.[27] To calibrate commuting costs in this version of the model, I use data on the share of commuters by city zone in Mumbia from Baker *et al.* (2005). To be specific, I target the following moments: share of slums dwellers in the periphery that commute to the center equal to 26%, and share of formal housing dwellers in the periphery that commute to the center equal to 34%.[28]

Table 1.15 compares the aggregate effects of destroying central slums in the baseline model with the case that includes commuting. The results show that allowing for commuting leads to a lower decrease in the share of urban population, as well as lower losses of total output and individual welfare. Intuitively, the periphery is a more attractive place and more individuals remain in the city when central slums are destroyed, which means that the concentration of talent in rural areas is lower in the case with commuting. Overall, introducing commuting to the model is this particular way does not have a large impact on the results. The reason is that to match the distribution of population across locations in the city, the joint calibration implies

[27]Commuting costs are introduced as as an income tax, however, they represent the actual expenditures on commuting (e.g. bus or train), plus the income that is lost due to commuting time and the possible productivity loss caused by a stressful commuting.

[28]I assume that the periphery of Mumbai is equal to zones 4 and 6 in the cited study, which is consistent with the definition used in calibraiton of the baseline model. The study does not report commuting times by type of housing, only by income category. I use data from Table D-1 and assume that slums dwellers are the same as the poorest income category, and the formal sector represents the third income category (low to high). According to this study most commuting happens between adjacent zones and the share of commuters from the periphery to the commercial area in the south of the city is low.

Table 1.15: Aggregate Effects of Destroying Central Slums: Allowing Commuting

Variable (% change)	Baseline	With commuting
Urban population share	-25.7	-23.9
Total Output	-6.8	-5.6
Individual Welfare	-2.7	-2.5

Notes: This table presents the effects of destroying central slums in the model by setting $\phi_c = 0$. The baseline results are compared to the case where individuals can commute from the urban periphery to the center.

that peripheral locations are not attractive places to live because productive is low and commuting costs are high.

1.6.2 Decreasing Returns in Rural Production

In the baseline model rural production uses a technology with constant returns to scale. This section explores the implications of introducing decreasing returns in rural production. The possible argument to evaluate this extension is that agriculture is the main activity in rural regions and a fixed amount of land at the production unit level generates the possibility of having decreasing returns to labor. To introduce this in the model, rural output is now given by $Y_r = A_r \, Z_r^{\beta}$, where $0 < \beta < 1$. I consider a case with strong decreasing returns by setting β equal to 0.5.

Table 1.17 presents the effects of destroying central slums in the case with decreasing returns in rural production, as well as the comparison with the baseline results. According to these results, with decreasing returns in rural production the fall in urban population is smaller. While most evicted dwellers are still displaced to

Table 1.16: Aggregate Effects of Destroying Central Slums: with Decreasing Returns in Rural Production

Variable (% change)	Baseline	With rural decreasing returns
Urban population share	-25.7	-19.0
Total Output	-6.8	-7.9
Individual Welfare	-2.7	-5.1

Notes: This table presents the effects of destroying central slums in the model by setting $\phi_c = 0$. Decreasing returns in rural production are introduced by defining $Y_r = A_r Z_r^{\beta}$, with $\beta = 0.5$.

the rural region, this location becomes less attractive as people move there because congestion in rural production drives wages down. In addition, the loss in total output is larger compared to the baseline case because rural labor productivity falls as individuals move there. The welfare loss is also larger because more individuals stay in the city, but either they pay a higher housing price in the urban center, or live in the less productive peripheral region. In the baseline case, individuals displaced to the rural region paid a lower housing rent and received higher amenities, without putting downward pressure on rural wages.

These results suggest that decreasing returns in rural production might be quantitatively important; however, Lagakos *et al.* (2018a) find a value of β equal to 0.91 for the case of Bangladesh, which implies much weaker decreasing returns in rural production than the ones considered in the experiment presented.[29]

[29] The findings in Hnatkovska and Lahiri (2014) and Munshi and Rosenzweig (2016) suggest that rural wages do not increase substantially with urban migration in India. This is also not consistent with strong decreasing returns in rural production.

1.6.3 Endogenous Agglomeration Forces

The baseline model assumes exogenous location productivities and regional amenities. This section introduces endogenous agglomeration and congestion forces to urban regions in the model. In contrast to rural regions, cities feature agglomeration economies that come from firms and consumers locating close to each other; however, as people move to cities there is more traffic, crime, and pollution. To introduce these forces in the model, I define location productivities in urban locations $\ell \in \{cf, cs, pf, ps\}$ as

$$A_\ell = \overline{A}_\ell \, Z_x^{\lambda}, \tag{1.20}$$

where $\overline{A}_\ell$ represents the exogenous component of productivity and $Z_x = Z_{xf} + Z_{xs}$ is the supply of efficiency units of labor in urban region $x \in \{c, p\}$. Thus, location productivity raises with the amount and talent of residents in the region. This definition allows for human capital spillovers across types of neighborhoods in each urban region and parameter λ governs the strength of these spillovers. On the other hand, urban regional amenities are now given by

$$\mu_x = \overline{\mu}_x \, \pi_x^{\gamma}, \tag{1.21}$$

where $\overline{\mu}_x$ represents exogenous amenities and $\pi_x = \pi_{xf} + \pi_{xs}$ is the share of population living in the region. Parameter γ governs the degree of congestion (if negative). This specification implies that the share of population living in slums affects the utility of individuals living in formal housing within both urban regions.

To analyze the quantitative results of the extended model, I take the value of agglomeration and congestion parameters from the literature. Following Bryan and Morten (2018), I set λ equal to 0.05 and γ equal to -0.075. Table 1.17 compares the aggregate effects of destroying central slums with and without endogenous agglomeration and congestion. In the model with externalities, the fall in urban population is

Table 1.17: Aggregate Effects of Destroying Central Slums: with Endogenous Agglomeration Forces

Variable (% change)	Baseline	With endogenous agglomaration forces
Urban population share	-25.7	-24.1
Total Output	-6.8	-7.0
Individual Welfare	-2.7	-2.2

Notes This table presents the effects (% change) of destroying central slums in the model by setting $\phi_c = 0$.

Endogenous agglomeration forces refers to human capital spillovers and amenities congestion in urban regions.

smaller, as well as the average welfare loss. The latter is due to the fact that evicting slums dwellers reduces amenities congestion in central urban areas. However, the loss in output is now larger because human capital spillovers decreased in the city.

These results imply that introducing externalities has a small quantitative effect and exogenous location fundamentals are the main drivers of the results. That said, agglomeration and residential externalities are difficult to measure and there is not much evidence for poor countries. To address the latter, Table 1.18 presents robustness checks for the parameters governing the strength of these externalities. According to the results, demolishing central slums causes a larger fall in total output when there are strong productivity spillovers, this is because productivity and returns to education are higher in central urban regions and displacing workers from these locations implies a large loss in human capital externalities. On the other hand, the loss in output is smaller when productivity spillovers are weak and there is con-

Table 1.18: Different Strengths of Agglomeration and Congestion: Effects on Total Output (% Change)

	Amenities congestion	
Productivity spillover	$\gamma = 0.05$	$\gamma = -0.05$
$\lambda = 0.01$	-7.5	-6.6
$\lambda = 0.1$	-9.8	-8.2

Notes: This table presents the effects (% change) of destroying central slums in the model by setting $\phi_c = 0$.

gestion in amenities. In this case, evicting slums dwellers from the center decreases the level of congestion for those who stay in the region, while the agglomeration effects on productivity are small. Overall, the qualitative results of the baseline case do not change with different degrees of endogenous agglomeration forces. However, externalities could have bigger effects in a model of a city with a large number of locations interacting through commuting (e.g. Ahlfeldt *et al.* (2015) and Tsivanidis (2018)).

1.6.4 Robustness

This section evaluates the robustness of my findings to different parameter values. First, Table 1.19 presents the results of reducing the population in central slums considering lower values of the elasticity of substitution across urban locations, and using alternative values of the adjusted shape parameter in the distribution of talent. According to these results, the change in urban population is quite sensitive to the level of substitutability between labor from different urban locations. When the

Table 1.19: Robustness: Destroying Central Slums

		Elasticity Substitution		Shape parameter	
Variable (% change)	Baseline	$\sigma = 3$	$\sigma = 6$	$\theta = 3$	$\theta = 12$
Urban population	-25.7	-44.6	-29.8	-24.4	-30.2
Total Output	-6.8	-11.9	-8.0	-7.5	-5.0

Notes: This table presents the effects (% change) of destroying central slums in the model by setting $\phi_c = 0$. Baseline values are $\sigma = 9$ and $\theta = 4.13$.

elasticity of substitution is lower and, therefore, the complementarity is stronger, the fall in urban output is higher as workers from central slums are displaced. The latter also explains why the loss in total output is larger compared to the baseline.

On the other hand, the results presented in the last column of Table 1.19 show that the aggregate effects of destroying central slums are less sensitive to changes in the adjusted shape parameter. A larger value of θ implies a higher correlation across location-specific talent or a lower variation in individual productivity. Thus, with a large value of θ there is less scope for comparative advantage and the sorting effect on productivity is weaker. This explains why the loss in output is somewhat lower in comparison to the baseline; however, the fall in urban population is higher because individuals are more sensitive to location-specific characteristics such as the rental price of housing.

Finally, in the baseline case, location-specific Mincer returns were estimated assuming that they represent technological differences in the value of skills across locations. However, these differences in returns to education could also represent selection

Table 1.20: Robustness: Destroying Central Slums Accounting for Selection in Returns to Education

Variable (% change)	Baseline	Selection in education returns
Urban population	-25.7	-26.0
Total Output	-6.8	-7.2

Notes: This table presents the effects (% change) of destroying central slums in the model by setting $\phi_c = 0$. To account for selection I reduce urban returns by 50% of their difference with respect to the return in rural regions.

of talented workers. That is, returns to education could be higher in urban regions because individuals with more cognitive talent choose to live there. To assess the implications of overestimating the importance of technological differences in the use of education across locations, I reduce the value of Mincer returns in urban locations by 50% of their difference with respect to the rural region. Doing this implies that the model will underestimate the difference in educational attainment between urban and rural areas because the complementarity between individual education and schooling returns is weaker.

The results in Table 1.18 show that the qualitative implications do not change after accounting for selection in returns to education. The effects are larger in this case because, in order to match the moments of data, other location fundamentals must adjust; in particular, the value of productivity in urban locations has to be higher to be consistent with the share of population living there. Moreover, in the baseline economy there is more complementarity between education level and returns to education, so the loss in output is smaller when the least educated individuals leave the city than the case where individuals with similar levels of education leave. In other words, as differences in returns to education disappear in the model, there is less selection in education across locations.

1.7 Conclusion

Slums are a prevalent urban phenomenon in developing countries. Policy responses to their formation vary from upgrading programs to forced evictions of dwellers. India is a case in which slums have been demolished in central urban areas with the idea that displaced dwellers relocate to peripheral regions where land is less valuable. This type of anti-slum policies are also observed in other developing countries, such as Zimbabwe and Nigeria. However, slums are formed in urban centers because workers want to live close to the best jobs in a city.

This paper develops a spatial general equilibrium framework to assess the quantitative impact of anti-slum policies in a context where formal housing markets are highly distorted. The model takes into account individual selection and differences in local productivity, returns to education, regional amenities, and housing rents. The findings imply that demolishing central slums shrinks the size of urban population because only a small share of evicted dwellers can afford to stay in formal housing and the urban periphery is not attractive. This is consistent with the fact that India has a small urban population share compared to countries with similar and lower levels of development. The losses in welfare and labor productivity are significant considering the fact that only a small fraction of the total population lives in central slums. Even if valuable land becomes available for formal development when slums are destroyed, most of the evicted dwellers lose access to the labor market opportunities they had in central urban regions.

I also use the model to assess the effects of reducing formal housing distortions in

India to the U.S. level. The results imply that these distortions account for a big share of the low urbanization rate in India and keep a high fraction of the urban population living in slums. Moreover, reducing housing distortions raises total output because individuals with low talent for rural activities move to the city and gain access to locations where productivity and returns to education are high. Thus, eliminating distortions in formal housing markets, such as restrictive FARs or large approval costs, has substantial effects on the spatial allocation of talent and aggregate labor productivity.

A possible direction for future research is exploring the political economy of slums and the implications for aggregate outcomes. It seems crucial to understand under what economic conditions governments find it profitable to increase the provision of urban infrastructure and public services in slums, and under what conditions they choose forced evictions. Research could aim to understand the role of interest groups in policies regarding slums. For example, groups representing residents of formal housing seem to put pressure on governments to sanitize cities by demolishing slums. Rent seeking behaviors can have a great impact on the allocation of land in urban regions and, therefore, on the allocation of workers across space.

Chapter 2

CROP CHOICE, TRADE COSTS, AND AGRICULTURAL PRODUCTIVITY

2.1 Introduction

A large body of literature documents that studying agriculture is critical for understanding cross-country income differences.[1] The reason is twofold. First, while poor countries are much less productive in aggregate output per worker, the productivity gaps are particularly large in agriculture. Second, despite these large productivity gaps, poor countries allocate a high share of their labor force to agriculture. Combined, these two facts prompt a key question: why do poor countries devote so much labor to such an unproductive sector?

Taking a closer look at agricultural production reveals a striking pattern within this sector. Figure 2.1 compares yields and harvested land between two categories of crops: (i) grains, including the most important staples such as maize and wheat; and (ii) fruits, which are usually grown as cash crops.[2] Almost every poor country allocates a large share of their land to produce staple crops, even though yields of cash crops are significantly higher. While similar patterns can be found in richer countries, there are two reasons these facts are especially important in poor countries: (i) a high share of the population works in agriculture, which means a large share of the

[1]See Gollin *et al.* (014b) for a summary of the topic.

[2]This figure shows output value per hectare. In poor countries, this could be a good approximation to value added per hectare since intermediate inputs usage is low; however, comparing land productivity across crops, especially in rich countries, would require value added at the crop level. To the best of my knowledge, there is no public source with such data.

labor force produces low-productivity crops; and (ii) producing fruits involves labor-intensive activities such as picking and stacking, thus poor countries could exploit their production given the capital constraints and labor-intensive techniques of most farmers.

The main goal of this paper is to further refine this puzzle. I use detailed, restricted access data from Mexican farms to document two facts. First, while Mexican agriculture is much less productive than non-agriculture, the productivity gaps in value added per worker are much larger for staple crops. Second, despite these large productivity differences, farm labor in Mexico is mostly devoted to staples production. Together, these two facts suggest that we can focus on an even narrower question: why do poor countries devote so much labor to unproductive staple crops?

This paper proposes an explanation based on two key mechanisms that determine the efficient crop choice by farmers: subsistence requirements of staple crops and *interregional* trade costs. The former is based on the observation that the share of production used for family consumption is large for staples. These crops have a high caloric content and represent an important nutritional source for poor farmers who have incentives to produce their own food. The second mechanism is based on the fact that trade costs affect the relative farm price between types of crops. Trade costs are especially high for fruits, so farmers receive a smaller share of the market value and must offer relatively low prices to be competitive. This means that only those farmers who are highly productive in fruits farming produce in that sector. Moreover, enough labor and land need to be allocated to farming staple crops when food is costly to move across regions in a country, especially to densely populated urban regions.

To formalize the analysis, I build a general equilibrium model with interregional trade and self-selection of heterogeneous farmers into type of crops. The model features non-homothetic preferences, costly trade across regions, and the existence of two

Figure 2.1: Land Allocation and Yields in Poor Countries and Mexico

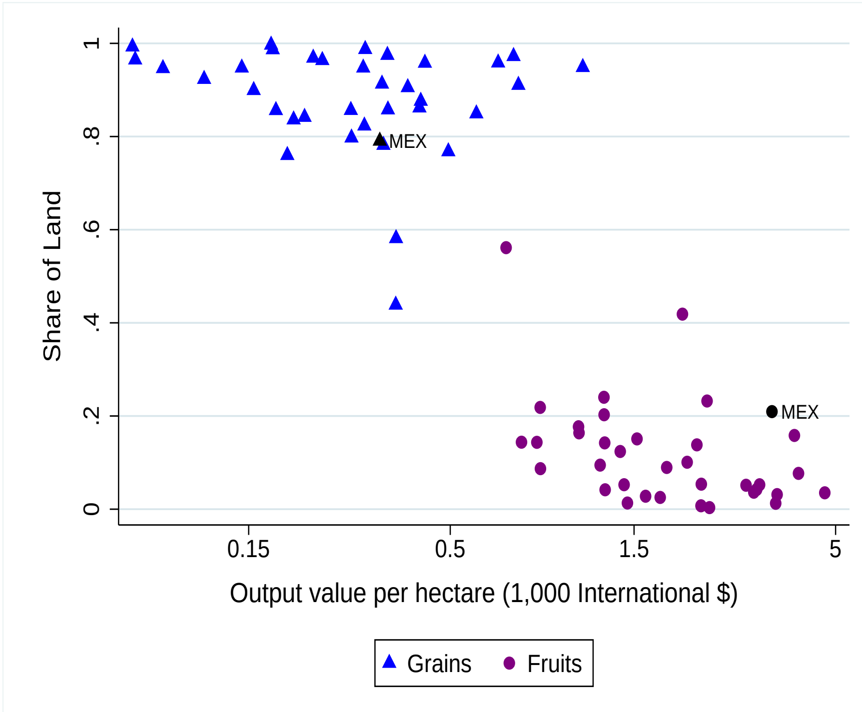

Notes: Average yields are reported for each category in a country. Yields are measured as net production value per hectare and are weighted by harvested hectares. The share of land is reported with respect to total harvested hectares in both categories. Output value (in constant 2004-2006 1,000 I$) is defined as gross output value net of agricultural inputs (seed and feed).

Includes Mexico and 35 countries classified as least developed by the United Nations.

Source: FAO, 2010.

agricultural goods: a staple crop (maize) and a cash crop (fruit). In this framework, farmers choose to produce either type of crop as an efficient response to subsistence requirements of staple food and trade costs. The existence of the latter causes a negative "income effect" in the economy, which increases the relative demand for staple food and, therefore, the relative price of maize. This is due to the fact that income elasticity of demand for staple food is less than one. Additionally, the selection of farmers into staple crops is amplified by the fact that trade costs are higher for fruits.

I calibrate the model to match features of the Mexican economy. For the topic of interest in this paper, Mexico is poor enough and the share of labor in agriculture (13.6% in 2014) is large compared to rich countries (1.6% in the U.S.). Furthermore, the yield gap between staples and cash crops is similar to poorer countries as shown in Figure 2.1, and, unlike many of the latter, it has good quality agricultural data.[3] In particular, I use detailed farm data on prices, production value, expenses, employment, and land usage at the crop level. I estimate trade costs from price gaps of homogeneous goods across regions in the country; thus, the definition of trade costs is broad and includes more than just transportation costs between distant regions, they also represent possible monopoly power of intermediaries. The quantitative results of the model imply that trade costs can account for a large share of labor allocated to maize and low labor productivity in agriculture. In a counterfactual case without trade costs, agricultural labor productivity increases by 21% and the ratio of employment in maize to fruits decreases by 19%.

This paper is related to the recent macroeconomic literature that has tried to explain agricultural productivity differences across countries. Restuccia *et al.* (2008) document that differences between rich and poor countries in Gross Domestic Product (GDP) per worker in agriculture are over two times higher than differences in aggregate GDP per worker.[4] Furthermore, they show that barriers to access intermediate inputs can account for large cross-country differences in agricultural employment and productivity.[5] Recent papers in this literature have taken into account production decisions within agriculture, for example, Adamopoulos and Restuccia (2015) develop

[3]Statistical agencies in the poorest countries have limited resources and data is usually unreliable. See Jerven (2013).

[4]Caselli (2005) finds similar aggregate and sectoral productivity differences across countries.

[5]Alternative explanations are given by Vollrath (2009), Lagakos and Waugh (2013), Young (2013), Adamopoulos and Restuccia (2014), Gollin *et al.* (014a), Herrendorf and Schoellman (2015), Herrendorf and Schoellman (2018), Donovan (2018), and Alvarez (2018).

a model that features crop choice to analyze the effects of land reforms on farm size and agricultural productivity in the Philippines. I focus on the interaction of trade costs and subsistence requirements of staple food as a determinant of resource allocation within agriculture, and how the sorting of farmers across types of crops affects agricultural labor productivity.

Additionally, this paper is closely related to recent literature that examines the effects of transportation costs on interregional trade using general equilibrium trade models. In particular, Donaldson (2018) studies the effects on trade and welfare of large transportation infrastructure projects using detailed data from India.[6] I build on the methodology of this literature to measure sector-specific trade costs across regions, and construct a unique dataset that combines farm and market data to compare prices of homogeneous goods between origins and destinations in Mexico.

Lastly, this is not the first paper that studies the relationship between trade and agricultural productivity. Tombe (2015) develops a multi-sector model with subsistence requirements to analyze the effects of trade costs on agricultural productivity differences across countries; and Sotelo (2018) studies the effects of regional trade frictions on welfare and farm productivity using data from Peru.[7] My focus is on documenting differences in labor and land productivity across crops using farm data that allows me to measure value added at the crop level, and building a quantitative model to assess the importance of trade costs for crop choices and low agricultural labor productivity. Furthermore, this paper interprets subsistence agriculture as farmers who choose to produce staple crops instead of high-value cash crops.

The rest of the paper is organized as follows. Section 2.2 presents a description

[6]See Donaldson and Hornbeck (2016) and Alder (2019) for similar approaches.

[7]Adamopoulos (2011), Herrendorf *et al.* (2012), and Gollin and Rogerson (2014) also focus on the links between trade costs, agriculture, and economic development.

of the microdata used in this paper, as well as empirical evidence on crop productivities and trade costs in agricultural markets in Mexico. Next, Section 2.3 introduces a multi-sector selection model with interregional trade. Section 2.4 provides a description of the calibration and presents the quantitative results of counterfactual experiments using the model. Finally, Section 2.5 concludes.

2.2 Empirical Evidence

To present the empirical evidence, I simplify the analysis by narrowing down the number of crops considered. Based on its production volume, harvested land, and relevance for subsistence, maize is the most important crop in Mexico. In this paper I will use it as a benchmark of staple crops and compare it to other fruits that are among the most important cash crops in the country.

I use restricted access, farm-level data from agricultural surveys in Mexico. These microdata is part of the *Encuesta Nacional Agropecuaria* (ENA) 2014.[8] The surveys were taken from a sample of 75,148 farms in 25,800 localities of the country during the agricultural cycle from fall 2013 to fall 2014. They gathered nationally representative data for 34 products that were chosen based on their contribution to GDP. The unit of observation in the survey is defined as a unit of agricultural production formed by a set of plots located in the same municipality. Since more than one crop can be grown by a unit of production throughout the agricultural year, each observation in the database represents a farm-crop pair. The target population of the surveys were all the production units that reported data for one of the products of interest in the agricultural census 2007.

[8] Access provided by the *Sistema Nacional de Información Estadística y Geográfica* (SNIEG) of the *Instituto Nacional de Estadística y Geografía* (INEGI). The views and conclusions expressed are exclusive of the author and do not reflect official positions or statistics of SNIEG, or INEGI.

The ENA 2014 surveys have detailed information by variety of crop at the farm level. For each crop that is grown in a farm, the surveys report harvested land and production volume; amounts of production used for family consumption, feed, and seed; farm-gate prices of output sold; quantities used of fertilizers; and farmers' expenses in different stages of production. The latter include expenditures on modern inputs such as chemicals, pesticides, and irrigation. In addition, hired labor and other farm expenses are reported at the farm level. The fact that most expenditures are reported at the crop level allows me to calculate valued added for each of them. This is key to make productivity comparisons across crops. The detailed information on prices, output, and expenditures of each crop produced in a farm, plus its geographical location, imply that this data is subject to confidentiality regulations.

In addition to the microdata described above, I use various sources of agricultural data provided by Mexico's government agencies. Specifically, I use data from the *Sistema Nacional de Información e Integración de Mercados* (SNIIM) to get data on wholesale prices for specific varieties of crops in every state of the country. These data provide the place of origin (state) for each variety sold in a market.

2.2.1 Crop Productivity

In this section I show that productivity in maize farming is significantly lower than productivity in fruit farming. I consider two measures of crop productivity: value added per worker and value added per hectare. I calculate value added for each farm-crop observation in the following way. First, I obtain the value of production net of the amount used for seed and animal feed; then, I subtract expenditures on fertilizers, pesticides, and irrigation. Since the amount of labor is reported at the farm level, I focus on farms producing one type of crop to measure labor productivity. See

Appendix B for more details.[9]

Panel A in Figure 2.2 shows aggregate value added per worker for different fruits relative to value added per worker in maize. For most of these fruits, labor productivity is over two times higher than in maize, and the average productivity gap is around six. In comparison, the ratio of value added per worker in non-agriculture to agriculture was 5.7 in Mexico during 2013. Thus, the agricultural productivity gap in Mexico has a similar magnitude to the productivity gap between fruits and maize. However, despite these large productivity differences, Panel B shows that labor allocated to maize is much higher than any of the fruits considered. All these crops together add up to 46% of total workers allocated to maize. These facts imply that highly unproductive crops like maize decrease the value of agricultural labor productivity in the country. Moreover, Figure 2.4 presents the gaps in aggregate value added per hectare for the same group of crops. The differences in land productivity between maize and fruits are even larger than the labor productivity gaps.

One possible concern with the aggregate results described in the previous paragraph is that such productivity gaps are driven by differences in farm size between maize and fruit producers, or by particular regions of the country that are highly productive in fruit farming. To address these issues, I estimate productivity gaps between types of crops controlling for state and farm size. The results presented in Table 2.1 show that productivity is significantly larger for fruits than for maize even if such controls are taken into account. That is, adjusting for region and the size of farms, the labor productivity gap between fruits and maize is 3.4 (the raw gap is 5.8), while the land productivity gap is 4.5 (the raw gap is 4.8). The adjusted productivity gaps are smaller, but still sizable. Appendix C provides additional evidence

[9]The results from this section are complemented in Appendix D using alternative sources of public aggregate data.

Figure 2.2: Panel A. Value Added per Worker Relative to Maize

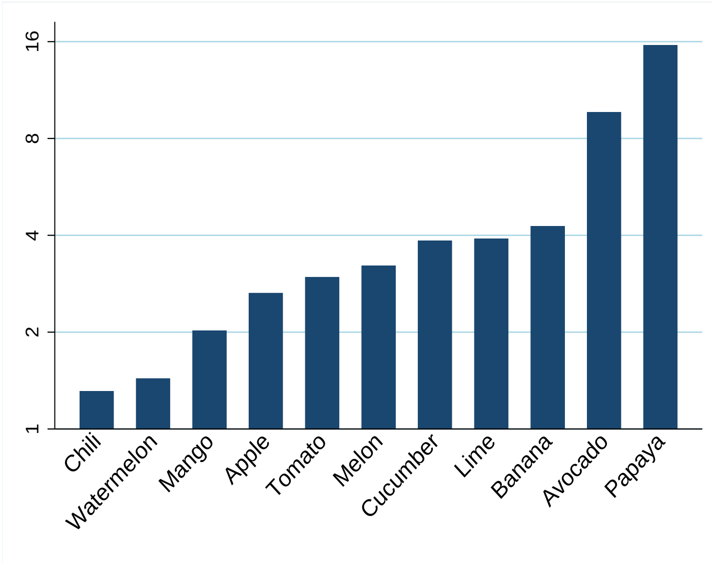

Figure 2.3: Panel B. Total Workers Relative to Maize

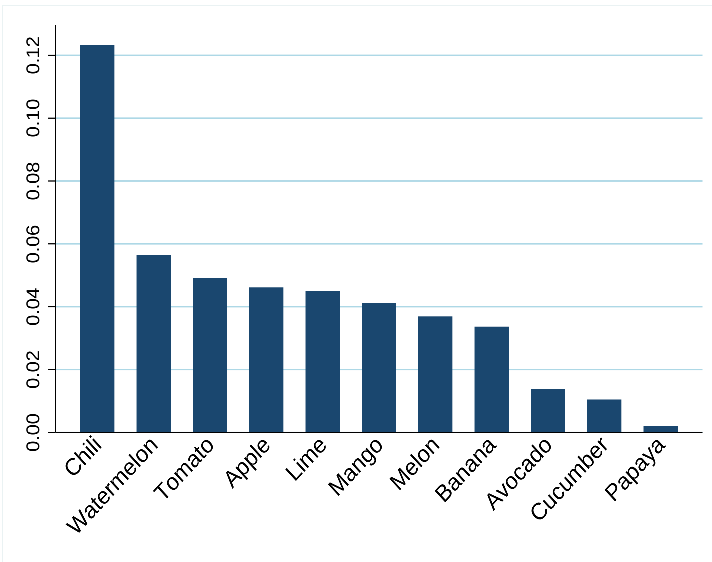

Notes: Maize is normalized to 1 in both cases. Total workers includes family members participating in farming activities.

Source: Author's estimates using data from SNIEG and INEGI: ENA 2014.

suggesting that farmers do not choose to grow maize because fruits are unproductive or infeasible to grow in their region; even if some regions of the country are more

Figure 2.4: Value Added per Hectare Relative to Maize

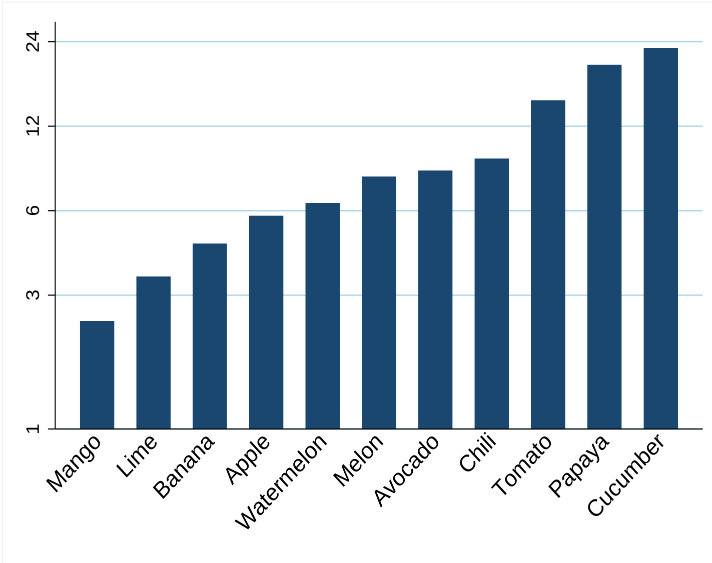

Notes: Maize is normalized to 1.

Source: Author's estimates using data from SNIEG and INEGI: ENA 2014.

suitable to grow fruits, it is not the case that maize is the only productive option for most farmers.

The empirical results presented in this section suggest the idea that a significant fraction of farmers must be relatively unproductive at producing staple crops: not every farmer has the best land to grow maize, nor the set of skills or knowledge required to manage such type of farms. The fact that most farmers decide to grow staple crops implies that there might be barriers amplifying the selection of farmers into those crops.

2.2.2 Trade Costs

This section presents evidence that trade costs are large in agricultural markets in Mexico. I measure these costs indirectly using differences in prices across regions. Thus, trade costs consist of more than just moving goods across distant regions, in-

Table 2.1: Productivity Gaps: Fruits Relative to Maize

Value Added per Worker	Raw Gap	Adjusted Gap
Gap	5.8	3.4
State fixed effects, farm size	No	Yes
Observations	26,197	26,197

Value Added per Hectare	Raw Gap	Adjusted Gap
Gap	4.8	4.5
State fixed effects, farm size	No	Yes
Observations	33,189	33,189

Notes: Results obtained from regressing log(value added per worker) and log(value added per hectare) on a dummy that takes a value of 1 for fruits and 0 for maize. The gap reported is the exponential of the estimated dummy coefficient. Controls include (log) agricultural land of the farm and state dummies. The coefficient of fruits is significant at the 1% level in every case. Regressions are weighted by worker and hectares, respectively.

Source: Author's estimates using data from INEGI-ENA 2014.

stead they reflect the quality of infrastructure for transportation and storage in each region, and how competitive markets are. My data satisfies two important characteristics to measure trade costs using spatial price gaps: (i) homogeneous products; and (ii) regions that are actually trading with each other. Donaldson (2018) uses a similar empirical strategy using regional varieties of salt in India.

First, to measure trade costs of crops, I compare farm-gate prices with wholesale market prices across states in 2014.[10] I build a unique dataset of prices for specific

[10]These wholesale markets are large supply centers of agricultural products located in the main

crop varieties by combining farm data from ENA 2014 surveys with market prices listed in SNIIM. The latter are reported monthly, so I calculate the average price in 2014 for each crop in every market. Since market prices only specify the state of origin for each product, I also aggregate farm prices to the state-level. Price gaps between each origin and destination are measure as the farm price divided by the wholesale price (which I define as the farm share). I only consider origins and destinations that are potentially trading, that is, observations in which the destination had a price greater or equal than the origin price. After this, I end up with 930 origin-destination observations for 68 fruits and grains traded across 30 states.[11]

The median farm share is 59% for grains and 36% for fruits. This means that a maize farmer receives more than half of the wholesale value, while a fruit farmer receives a little more than a third of the value. Specifically, the median farm share of maize (*white*) is 69%; in comparison, the median farm share of avocado (*hass*) and pepper (*poblano*) are 49% and 32%, respectively. One possible concern with these results is that fruits might be traded to further distances than grains. In the second column of Table 2.2, I present the raw gap in trade costs between cash crops and maize, while the third column shows the estimated gap controlling for origin-destination fixed effects. There is almost no differences between raw and adjusted gaps; thus, the differences in trade costs are not driven by differences in trading routes.

These estimates are in line with previous studies on Mexico's agricultural sector according to which farmers in fruits and vegetables markets receive between 35% and 45% of retail prices. According to these studies, the existence of few intermediaries

cities of every state.

[11] I do not consider Baja California and Baja California Sur due to their isolated geographical location.

Table 2.2: Trade Costs of Fruits Relative to White Maize

Crop (Variety)	Raw Gap	Adjusted by Routes
Avocado (Hass)	1.5	1.4
Cucumber	2.4	2.0
Pepper (Poblano)	2.1	1.9
Tomato (Saladette)	1.5	1.4
Watermelon (Rayada)	1.9	2.0

Notes: Adjusted gaps take into account origin-destination fixed effects.

Source: Author's estimates using data from INEGI-ENA 2014 and SNIIM.

controlling the distribution of cash crops was related to low prices faced by farmers. In contrast, producers in other Latin American countries receive 50% of retail prices, and in some cases of Central America between 65% and 75%.[12]

The previous paragraphs focused on trade costs of output in agricultural markets. However, trade costs of modern inputs are also relevant. To measure trade costs of fertilizers, I use a similar approach though the available data is different. First, from the ENA 2014 surveys, I obtain quantities in tons of chemicals and natural fertilizers used for production of crops. Total expenses on fertilizers are also reported for each crop. Then, I calculate implicit prices of fertilizers at the farm level dividing total expenditures by total quantity of fertilizers. A wide range of varieties of chemical fertilizers are reported by farmers, so I focus on the chemicals that account for the majority of the observations: Urea and Ammonium Sulfate. The former is mostly imported from Eastern Europe, while the latter is primarily produced in the central

[12]See OECD (2007), p. 67.

Table 2.3: Trade Costs of Fertilizers: Farm Prices Relative to Market Price in Origin

Variety	Price gap
Urea	1.62
Ammonium Sulfate	1.53

Notes: Median gap at the national level. Most of Urea is imported so the origin refers to a port (Veracruz). For Ammonium Sulfate, origin refers to the region where production plants are located (Queretaro). See text for details.

Source: Author's estimates using data from INEGI-ENA 2014, and SNIIM.

region of the country.[13]

I use data from SNIIM to get the commercial price of these fertilizers in their possible state of origin: Veracruz (a port) for Urea and Queretaro (a production plant) for Ammonium Sulfate. To have a higher number of observations, I estimate trade costs for these products by aggregating farm prices to the municipality level and dividing them by the commercial price at the state of origin. Table 2.3 presents the median farm to market ratio. The magnitude of these gaps is consistent with the fact that most farmers report that a main obstacle of production are the high costs to acquire modern inputs.

This section provided evidence on the existence of large trade costs in agricultural markets in Mexico. Intuitively, these trade costs might amplify the number of farmers producing maize relative to fruits for the following reasons. First, farmers producing fruits receive a lower share of their market value because trade costs are higher in this sector. Second, when staple food is costly to trade from farms to dense urban regions, more labor needs to be allocated to its production to guarantee that demand

[13]Source: 2006 report from ANACOFER, a national association of production and distribution of fertilizers.

is satisfied. And third, given that modern inputs are costly to acquire, more labor needs to be allocated to produce enough staple food for the population. The following section introduces a model to assess the quantitative importance of interregional trade costs for crop choices and agricultural productivity in a context with subsistence requirements of staple food.

2.3 Model

I develop a static general equilibrium model with interregional trade. The model includes production of different agricultural goods and features heterogeneous productivity across farmers, non-homothetic preferences, and trade costs. The framework builds on the selection model of Lagakos and Waugh (2013), the trade literature based on Eaton and Kortum (2002), and the interregional trade model of Herrendorf *et al.* (2012).

There is an urban region denoted by u and a rural region denoted by r. Regions are indexed by $j \in \{u, r\}$. Each region is populated by a household of size N_j. Individuals can move freely between regions so N_j is endogenous. There are three sectors in the economy: a nonagricultural good (n), and two agricultural goods: one is maize (m), a staple crop which is used for subsistence requirements, and the other is a fruit (f) or cash crop. These goods are indexed by $s \in \{n, m, f\}$. I assume that the urban region only produces non-agricultural goods, whereas the rural region only produces agricultural goods. Interregional trade is restricted by sector-specific trade costs.

The details of the decision process are presented below. Here, I describe the timing of the choices in the model. First, individuals choose to live either in the urban household or the rural household. Then, the rural household allocates its members as farmers (farm managers) or farmworkers. Finally, farmers decide to produce either maize or fruits based on their individual productivity to produce each crop.

2.3.1 Preferences and Endowments

In both regions, household preferences are defined according to the utility function

$$U(c_{jm}, c_{jf}, c_{jn}) = \epsilon_m \log\Big(c_{jm} - \overline{m}\Big) + \epsilon_f \log\Big(c_{jf}\Big) + \epsilon_n \log\Big(c_{jn}\Big) \tag{2.1}$$

where $\sum_i \epsilon_i = 1$, $\overline{m} > 0$ is the subsistence requirement of maize consumption, and c_{js} is consumption per capita of good s in region j.

Each individual is endowed with one unit of time that is inelastically supplied to the labor market. In the rural region, the household is endowed with L units of land for agricultural production and decides the fraction of its members that operate farms and the fraction of members that are workers hired by farms. Each farmer i is endowed with the same fraction of land ℓ and a pair of efficiency units of land $\{z^i_m, z^i_f\}$ to produce crops m and f, which is drawn from a distribution $G(z_m, z_f)$. However, a farmer can only produce one type of crop in her plot of land. The heterogeneity in productivity across farmers can be interpreted as differences related to both the quality of land and the managerial skills of farmers to produce each crop. I abstract from the occupational choice between agriculture and non-agriculture by assuming that land productivities are drawn after the rural household has allocated its members as farmers and workers. I do this simplification to focus on the crop choice by farmers, which is the main subject of this paper.

2.3.2 Production Technologies

The non-agricultural good is produced according to a constant returns to scale production function using labor as the only input, $Y_n = AN_n$, where A is an economy-wide productivity parameter and N_n is the amount of labor used in non-agriculture.

Given prices, the representative firm in region u maximizes profits by solving

$$\max_{N_{un}} \quad P_{un} A N_n - W_u N_n, \tag{2.2}$$

where P_{un} and W_u are the price of the non-agricultural good and the wage per unit of labor in region u, respectively.

Farmers operating in the rural region use their fraction of land to produce agricultural goods in sector $s \in \{m, f\}$, according to the production function $y_s^i = A\,(z_s^i\,\ell)^{\alpha_s}(n_s^i)^{\beta_s}(x_j^i)^{\psi_s}$, where n_s^i and x_s^i are hired labor and nonagricultural intermediate inputs, respectively, used by farmer i to produce crop s. I allow factor shares to be potentially different across agricultural goods and assume $\alpha_s + \beta_s + \psi_s = 1$. Since land is fixed for each farmer, there are decreasing returns at the farm level.

Given the choice to produce crop s, taking prices as given, a farmer maximizes profits by solving

$$\max_{\{n_s^i, x_s^i\}_{s\in\{m,f\}}} \quad P_{rs}\,y_s^i - P_{rn}\,x_s^i - W_r\,n_s^i, \tag{2.3}$$

where P_{rs} is the price of good s in region r and W_r is the labor wage in region r. Then, the payment received by each farmer is defined as $\pi_s^i = \alpha_s P_{rs}\,y_s^i$. The latter are residual earnings of a farm after input payments are made.

Interregional Trade

Goods can be traded between regions subject to iceberg costs. Region j must ship τ_s^{jk} units of good s in order for one unit to arrive in region k. Thus, $\tau_s^{jk} = 1$ implies frictionless trade and $\tau_s^{jk} \to \infty$ implies autarky. By construction, the rural region sends crops to the urban region and the latter sends non-agricultural goods to the rural region, so I omit the superscripts. Then, relative prices between regions are given by

$$\frac{P_{rn}}{P_{un}} = \tau_n, \quad \text{and} \quad \frac{P_{rs}}{P_{us}} = \frac{1}{\tau_s}, \qquad s \in \{m, f\}. \tag{2.4}$$

Trade costs generate a wedge between prices across regions. In particular, trade costs increase the price of crops in the urban region and the price of intermediate inputs in the rural region. These trade technologies imply that interregional exports and imports, E_s and M_s, respectively, must satisfy the following restrictions

$$E_s = \tau_s M_s, \qquad s \in \{n, m, f\}. \tag{2.5}$$

That is, trade costs increase the amount of goods that must be shipped to satisfy a given amount of demand in the destination region.[14]

2.3.3 Equilibrium

Farmers in the rural region choose to produce crop m or f based on their comparative advantage. In a competitive equilibrium, a farmer decides to produce maize if and only if her residual earnings of maize π^i_m are higher than her residual earnings of fruits π^i_f, that is, if and only if

$$\frac{z^i_m}{z^i_f} \geq \mathcal{K} \, \frac{\left(P_{rf}\right)^{1/\alpha_f}}{\left(P_{rm}\right)^{1/\alpha_m}} \left(W_r\right)^{\left(\frac{\beta_m}{\alpha_m} - \frac{\beta_f}{\alpha_f}\right)} \left(P_{rn}\right)^{\left(\frac{\psi_m}{\alpha_m} - \frac{\psi_f}{\alpha_f}\right)}, \tag{2.6}$$

where $\mathcal{K}$ is a constant. Holding all else fixed, a lower relative price of fruits with respect to maize leads to a higher share of farmers producing maize. The effects of labor wages and the price of non-agricultural inputs on the crop choice depends on how intensive is maize production in labor and intermediate inputs relative to fruits production. Below, I present a simplified case to illustrate how crop choices are affected by the key features of the model.

Additionally, households maximize utility by choosing consumption per capita of each good subject to income per capita I_j/N_j. Then, it can be shown that optimal

[14] Equations in 2.4 and 2.5 can be obtained from modeling the firm's maximization problem in a competitive transportation sector. See Herrendorf *et al.* (2012).

consumption allocations in both regions are given by

$$
\begin{aligned}
c_{jm} &= \frac{\epsilon_m}{P_{jm}}\left(\frac{I_j}{N_j} - P_{jm}\overline{m}\right) + \overline{m}, \\
c_{jf} &= \frac{\epsilon_f}{P_{jf}}\left(\frac{I_j}{N_j} - P_{jm}\overline{m}\right), \\
c_{jn} &= \frac{\epsilon_n}{P_{jn}}\left(\frac{I_j}{N_j} - P_{jm}\overline{m}\right).
\end{aligned} \tag{2.7}
$$

Non-homothetic preferences imply that the expenditure share of maize decreases with income, while the expenditure share of fruits and non-agricultural goods increases. These preferences are consistent with the patterns observed for budget shares of cereals and non-food products as income increases, and account for the fact that subsistence production is mostly observed for staple grains.[15] Household income in the urban region is given by labor payments, while income in the rural region is given by labor payments plus farmers earnings, that is,

$$
\begin{aligned}
I_u &= W_u N_u, \\
I_r &= W_r N_{rw} + \big(N_r - N_{rw}\big)\left(\sum_{s\in\{m,f\}} \int_{i\in\Omega s} \pi_s^i dG_i\right),
\end{aligned} \tag{2.8}
$$

where N_{rw} is the fraction of household members that are farmworkers in the rural region and Ω_s represents the set of farmers producing crop s.

To define a competitive equilibrium, I assume the non-agricultural good is the numeraire and normalize $P_{un} = 1$. Then, market clearing conditions for goods are given by

$$
N_u c_{un} + E_n = Y_n, \tag{2.9}
$$

[15]See Appendix E for evidence on subsistence production of different crops.

$$N_u c_{uf} = M_f, \tag{2.10}$$

$$N_u c_{um} = M_m, \tag{2.11}$$

$$N_r c_{rn} + X_{rm} + X_{rf} = M_n, \tag{2.12}$$

$$N_r c_{rf} + E_f = Y_f, \tag{2.13}$$

$$N_r c_{rm} + E_m = Y_m. \tag{2.14}$$

According to equations (9)-(11), total production of non-agricultural goods in the urban region is equal to local consumption by the household plus exports to the rural region, and urban consumption of crops is met by imports from the rural region. Equations (12)-(14) say that local consumption of non-agricultural goods in the rural region plus total intermediate inputs used by farmers, where $X_{rs} = (N_r - N_{rw}) \int_{i \in \Omega_s} x_s^i dG_i$, $s \in \{m, f\}$, are equal to imports from the urban region, and total production of crops, $Y_s = (N_r - N_{rw}) \int_{i \in \Omega_s} y_s^i dG_i$, $s \in \{m, f\}$, is equal to local consumption plus exports to the urban region.

Free movement of individuals across regions implies that household utilities are equalized. In addition, individuals living in the rural region are indifferent between operating as farmers and working as hired labor in farms. Therefore, labor wage in the rural region must be equal to expected earnings of farmers,

$$W_r = \sum_{s \in \{m,f\}} \int_{i \in \Omega s} \pi_s^i dG_i. \tag{2.15}$$

Market clearing of labor market in the rural region requires that total labor demand equals the total number of farmworkers, $N_{rw} = N_{rm} + N_{rf}$, where $N_{rs} = (N_R - N_{rw}) \int_{i \in \Omega_s} n_s^i dG_i$. Note that in this model total labor in agriculture is equal to hired labor plus farm operators. That is, total labor allocated to crop s is given by $\mathcal{N}_{rs} = N_{rs} + \Phi_s$, where Φ_s is number of farmers producing crop s. The latter is consistent with the way in which total labor in crop production is calculated for the empirical evidence in Section 2.2.1.[16] Finally, the fraction of land that every farmer receives satisfies $\sum_{s\in\{m,f\}} \Phi_s \, \ell = L$.

A *competitive equilibrium* with interregional trade is a set of prices of goods and inputs $\{P_{jn}, P_{jm}, P_{jf}, W_j\}$, $j \in \{u, r\}$; location choices (N_j individuals choose region j); farmers' earnings, π^i_s, $s \in \{m, f\}$; sets of households' allocations, $\{C_{jm}, C_{jf}, C_{jn}\}$, $j \in \{u, r\}$, and $\{\Phi_m, \Phi_f, N_{rw}\}$; a set of input choices in each region, $\{N_{un}, N_{rm}, N_{rf}, X_{rm}, X_{rf}\}$; and a set of interregional trade flows, $\{E_f, E_m, E_n, M_f, M_m, M_n\}$, such that: (i) given prices and farmers earnings, households maximize utility in both regions; (ii) given prices, firms and farmers maximize profits; and (iii) market clearing conditions hold.

2.3.4 Productivities Distribution

I follow the parametrization of Lagakos and Waugh (2013) and define the joint distribution of crop-specific individual productivities as

$$G_j(z_m, z_f) = \mathcal{C}[F(z_m), F(z_f)], \quad \mathcal{C}[u, v] = \frac{-1}{\rho} \log\left(1 + \frac{\left(e^{-\rho u} - 1\right)\left(e^{-\rho v} - 1\right)}{e^{-\rho} - 1}\right). \tag{2.16}$$

$\mathcal{C}[F(z_m), F(z_f)]$ is a *Frank copula* with parameter $\rho \in \{-\infty, \infty\}\backslash\{0\}$. The latter governs the correlation between productivity draws, such that a positive value of ρ implies a positive dependence between z_m and z_f. The marginal distributions are Fréchet

$$F(z_s) = \exp\left(- z_s^{-\theta_s}\right),$$

where θ_s governs the dispersion of productivity draws and the scale parameter is normalized to one. There is a negative relationship between the value of θ_s and the

[16]According to Mexican national account data from 2008, non-hired labor(e.g. owners, family members and unpaid workers) account for 63% of total labor in agriculture.

variation of land augmenting productivity in crop $s \in \{m, f\}$; that is, a lower θ_s implies a higher variation in individual productivity. The dependence across productivity draws ρ and the variation of individual productivity θ_s determine the extent of alignment between absolute advantage and comparative advantage in a particular sector, that is, the difference in productivity between the marginal farmer and average farmer in a sector.[17] The quantitative section provides more details on the role of these parameters in the model.

2.3.5 Trade Costs and Crop Choice

Farmers select into crops as an efficient response to subsistence requirements of staple food and trade costs. To see how the model works, assume that $\alpha_m = \alpha_f$, $\gamma_m = \gamma_f$, and $\psi_m = \psi_f$; and $\epsilon_m \to 0$ (maize consumption is equal to the subsistence level). Then, the cutoff that determines the crop choice of farmers is given by

$$\frac{z_m^i}{z_f^i} \geq \left(\frac{P_{rf}}{P_{rm}}\right)^{1/\alpha} = \left[\left(\frac{\overline{m}\,\epsilon_f}{C_{uf}}\right)\frac{\tau_m}{\tau_f}\left(\frac{I_u/N_u}{P_{um}\overline{m}} - 1\right)\right]^{1/\alpha}. \tag{2.17}$$

In a world without trade costs, farmers would take prices of the urban market as given and decide which crop to produce based on the relative price; however, the existence of trade costs creates a wedge across regional prices and changes the relative price between crops in the rural region. Particularly, if trade costs of fruits are higher, then the relative price of fruits with respect to maize is lower and more farmers decide to produce maize. As shown in Section 2.2.2, trade costs of fruits are significantly higher than trade costs of maize in Mexico. Thus, trade costs amplify the selection into staple crops by reducing the relative price of fruits.

Moreover, subsistence requirements raise the share of farmers producing maize by

[17] See Young (2014) for a similar discussion applied to goods and services.

increasing the relative price of this crop. To see the interaction between trade costs and subsistence requirements in the model, note that when trade costs decrease in the economy, there is a positive "income effect" that leads to an increase in the demand for non-agricultural goods and fruits that is higher than the increase in the demand for maize. The latter is due to the presence of non-homothetic preferences in the model. This implies that an economy with lower trade costs has a relatively lower demand for maize and, therefore, a smaller share of farmers producing this crop.

Finally, trade costs can be considered a barrier that affects the allocation of labor across sectors, especially across types of crops. However, this is not a source of misallocation in the model as in Hsieh and Klenow (2009); to the contrary, farming decisions are efficient given the subsistence needs for staple food and the level of trade costs in each sector. The potential welfare and productivity gains from generating a movement of farmers from maize to fruit production come from the fact that a high concentration of farmers in maize implies that many of them have a relative low productivity in that sector. In other words, farmers who have a higher productivity draw for fruit production might decide to produce maize because the relative price of this crop is high.

2.4 Quantitative Analysis

In this section I calibrate the model to match features of the Mexican economy. Then, I introduce changes to the baseline economy to evaluate the quantitative role of trade costs in allocating farmers across types of crops and generating low agricultural labor productivity. Specifically, I quantify the effects of assuming no trade costs across regions. However, while the latter case is helpful to analyze the overall importance of trade costs, it does not represent a plausible scenario for policy implications; therefore, I also consider the counterfactual case of an overall decrease in trade costs to the U.S.

level. Finally, to consider the general implications for poorer countries than Mexico, I recalibrate specific parameters of the model to match features of a typical African country and evaluate the effects of reducing trade costs.

2.4.1 Calibration

For the baseline case, I normalize the economy-wide productivity parameter by setting $A = 1$. The remaining parameters that need to be calibrated are preferences weights ϵ_s, $s \in \{m, f, n\}$; the subsistence requirement of stable food $\overline{m}$; trade costs τ_s; productivity distribution parameters θ_s and ρ; and factor income shares $\{\alpha_s, \gamma_s, \psi_s\}$. I calibrate the economy so that the urban region in the model represents the main cities where large wholesale markets of agricultural products are located in every state, and the rural region represents all locations with farming production according to ENA 2014 surveys. To compare staples and cash crops, I focus on maize and the most important exported fruits in Mexico: avocados, chili peppers, cucumber, watermelon, melon, papaya, and tomatoes.[18]

Trade costs. The estimation of trade costs is based on the idea of comparing otherwise homogeneous products across origins and destinations in Mexico. Donaldson (2018) uses this principle to estimate trade costs in India by comparing prices of regional varieties of salt. In this case, I compare prices of specific varieties of crops and chemical fertilizers. The assumption is that a crop variety is essentially the same good when it is sold by a farmer in a given state than when it is bought by a consumer in a wholesale market in another state. For example, in the case of fruits, if they did not spoil during transportation and are eatable by consumers, they are essentially the same good in the farm and the marketplace. Similarly, a specific

[18] These crops alone account for approximately 50% of the total value of agricultural exports in the last decade (not including livestock and fishery products).

type of fertilizer, such as Urea or Ammonium Sulfate, is an homogeneous chemical compound regardless of the point of sale.

Trade costs of fruits τ_f are estimated by comparing farm-gate prices and market prices of crops varieties between each point of origin and destination, as in equation 2.4 of the model. The origin price refers to the average farm-gate price in each state obtained from ENA 2014 surveys and the destination price is the wholesale price in every state where it is sold. This procedure is described in Section 2.2.2. I focus on a subset of fruits varieties that are produced and/or sold in many states. These crops include avocado (hass), tomato (bola and saladette), watermelon (cambray), cucumber, papaya (maradol), and three varieties of chili pepper (poblano, jalapeño, and serrano). I aggregate trade costs of fruits weighting each crop by their national production value in 2014. For staple crops, I estimate the mean price gap between origins and destination of maize (white). The results of this estimation are τ_f equal to 2.20 (45% farm share), and τ_m equal to 1.51 (66% farm share). These estimates imply large differences in trade costs within agriculture, which may reflect higher transportation and storage costs of fruits (e.g. refrigeration and spoilage), as well as monopoly power of intermediaries. The fact is that farmers face higher trade costs to enter fruits markets.

To estimate trade costs of non-agricultural inputs, I follow the steps described in Section 2.2.2, which imply comparing prices of fertilizers faced by farmers with market prices in the place of origin. I focus on those cases where the chemical fertilizer used is Urea, which is the most common fertilizer reported in the surveys. Given that most of this fertilizer is imported, I estimate the price gap between farms and the market price in Veracruz, one of the main ports where this product enters the country.[19]

[19]Other important port is Manzanillo in the state of Colima, but market prices of Urea in this state are practically the same as in Veracruz.

The median price gap implies that τ_n is equal to 1.62, which is consistent with the fact that most farmers surveyed report high inputs costs as their main production obstacle.[20]

Jointly calibrated parameters. I jointly calibrate the subsistence parameter $\overline{m}$ and productivity distribution parameters θ_m, θ_f and ρ to match four moments: total employment in maize relative to cash crops, 3.43; the variance of log yields measured as output value per hectare in each agricultural sector, 3.01 for fruits and 2.05 for maize; and the ratio of average output value per hectare in fruits to maize, 3.77. The value of these moments is estimated using data from ENA 14 surveys.

The reasoning behind the joint calibration is the following. First, there is a positive relationship between the size of subsistence requirements of staple food and the share of total workers (farmers plus hired labor) producing in that sector. Secondly, the variation in output value per hectare in each sector is governed by parameters θ_m and θ_f. To see this, note that in the model output value per hectare for each farmer in sector s is given by the following expression: $P_{rs}y_s^i = z_s^i\Big(P_{rs}\Big)^{\frac{1}{\alpha_s}}\left(\frac{\beta_s}{W_r}\right)^{\frac{\beta_s}{\alpha_s}}\left(\frac{\psi_s}{P_{rn}}\right)^{\frac{\psi_s}{\alpha_s}}$; therefore, $\text{var}(\log(P_{rs}\,y_s^i)) = \text{var}(\log(z_s^i))$. By matching the conditional variance of yields in the model with the variation observed in data, it is possible to discipline the parameters the govern the unconditional distribution in each agricultural sector. Finally, as described by Lagakos and Waugh (2013), the correlation parameter ρ governs the yield gap across types of crops by determining how strong is the relationship between absolute and comparative advantage. The results presented in Table 2.4 imply that variation in fruits productivity is higher than in maize ($\theta_f < \theta_m$), which may reflect the fact that there is a wider variety of goods in the fruit sector, each of which requires particular farmer skills and land qualities to grow effectively. The positive

[20] I use the median gap because the distribution of relative prices of fertilizers has a long right tail and do not want to overestimate the size of trade costs for non-agricultural inputs.

value of ρ implies a Kendall rank correlation coefficient of 0.36; thus, the calibration implies a moderate positive correlation between productivity draws.

To illustrate the results described in the last paragraph, Figure 2.5 presents the probability distribution of individual productivity draws implied by the baseline calibration. This figure shows that the equilibrium cutoff has to be quite large to generate the difference in labor allocation between maize and fruits. Moreover, only those farmers with very high productivity in fruits decide to produce in that sector and, because there is a positive correlation between productivity draws, they tend to be the farmers with a high productivity in both crops. That is, the calibration implies that the large productivity differences across types of crops is due to the fact that the farmers with the best land (and managerial farm skills) tend to select into fruits farming.

Preferences weights. Preferences weights govern expenditures shares as income tends to infinity and non-homothetic parameters become irrelevant. I follow the calibration strategy of Gollin and Rogerson (2014) and Tombe (2015), and use data from the 2005 International Comparison Program (ICP) to get budget shares for aggregates and food categories. I set these parameters to match the budget share for food, tobacco and beverages in rich countries, 0.20, and the budget share for cereals relative to fruits and vegetables in the same group of countries, 0.88. The latter implies that $\epsilon_m = 0.09$ and $\epsilon_f = 0.11$.[21]

Technology parameters. Factor income shares are calibrated using data from Mexico Input-Output tables in 2008. To estimate factor shares of maize in the model, I consider data of grains, legumes and oilseeds farming, while factor shares of fruits in the model are estimated using data of fruits, nuts and vegetables farming. Pay-

[21]To put these numbers in context, low-income countries spend 49% of their budget on food, and the share for cereals is 1.29 times greater than the share for fruits and vegetables.

Figure 2.5: Baseline Probability Distribution

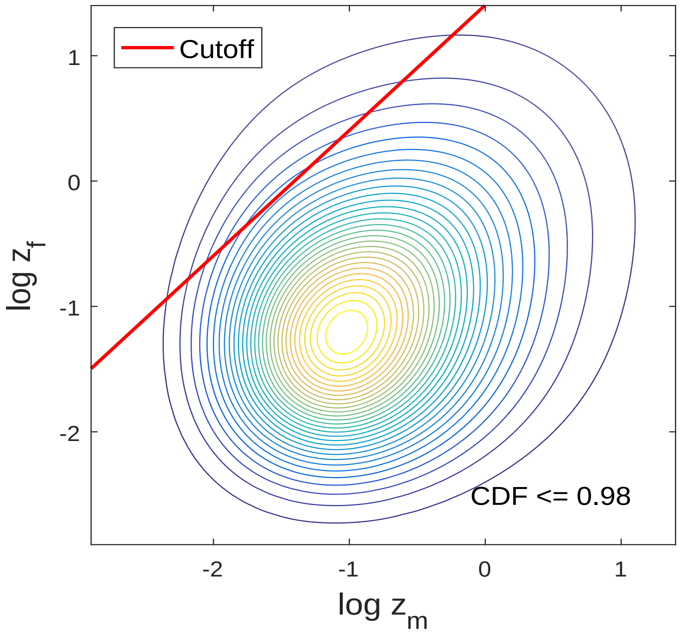

Notes: This figure is obtained by simulating the probability distribution of land augmenting productivities $G(z_m, z_f)$ implied by the internal calibration of the model.

ments to labor are calculated adjusting compensation to employees following Gollin (2002), that is, I impute the employee compensation of non-hired labor (owners, family, contract labor, and non-remunerated labor). For each sector, I calculate average employee compensation and multiply it by total workers (hired and non-hired). The share of non-agricultural intermediate inputs is computed using expenditures on inputs from non-agricultural sectors. Finally, I assume that payments to land include farm profits, so these payments are estimated as gross operating surplus minus the compensation of non-hired labor. The latter adjustment is made so gross value added in the industry remains unchanged. Results are reported in Table 2.4. According to these estimates, grains production is more intensive in labor than fruits. This is in line with labor-intensive techniques used in developing countries for subsistence

Table 2.4: Calibration Summary

Parameter	Value	Target
ϵ_m	0.09	Budget share for food, beverages and tobacco in rich countries
ϵ_f	0.11	Budget share for cereals relative to fruits in rich countries
β_m	0.35	Income share of labor in grains and oilseeds
ψ_m	0.16	Income share of non-agricultural inputs in grains and oilseeds
β_f	0.26	Income share of labor in fruits and vegetables
ψ_f	0.15	Income share of non-agricultural inputs in fruits and vegetables
τ_n	1.62	Intraregional trade costs of fertilizers
τ_m	1.51	Intraregional trade costs of maize
τ_f	2.20	Intraregional trade costs of fruits
$\overline{m}$	0.20	Labor in maize relative to fruits
θ_m	0.92	Variance of maize log yields across municipalities
θ_f	0.80	Variance of fruits log yields across municipalities
ρ	3.67	Ratio of average output value per hectare in fruits to maize

agriculture.

Lastly, I validate the model by looking at other quantitative implications. Table 2.5 compares the baseline results with non-targeted moments in the data. The model is able to replicate a similar share of agricultural employment to the one observed in the data. Moreover, the model does remarkably well in matching labor productivity differences across types of crops. The calibration targeted differences in employment and average output value per hectare, so there is no reason for the aggregate gap in value added per worker to be exactly the same in the model and the data. Additionally, the baseline economy generates higher labor and land produc-

Table 2.5: Quantitative Implications

Non-targeted Moment	Model	Data
Agriculture employment share (%)	17.50	14.00
Ratio of value added per worker in maize to fruits	0.41	0.41
Ratio of value added per hectare in maize to fruits	0.26	0.10

Notes: Aggregate productivity gaps are reported. Labor productivity is measured using total Value Added divided by total labor (hired labor plus farmers) in each sector.

tivity in fruits than in maize farming. The latter is especially important since these productivity gaps are the main empirical motivation of the paper.

A model without heterogeneous farmers could not replicate these results since differences in productivity across agricultural sectors would only reflect differences in the income shares of inputs. In such case, the productivity gaps could only be generated if maize production is significantly more intensive in both land and labor compared to fruits, or if explicit barriers or wedges are introduced to prevent the equalization of marginal products across sectors.

2.4.2 Quantitative Experiments

In this section I carry out multiple counterfactual experiments. First, I assess the impact of assuming that there is no trade costs across regions, that is, $\tau_s = 1$, for every sector. This case is useful to analyze the overall importance of trade costs. Nevertheless, since the latter is not a plausible scenario for policy implications, I consider a benchmark that is consistent with equivalent trade costs in the U.S.. Lastly, I evaluate the general implications for poor countries by recalibrating specific parameters of the model to match features of a typical African country.

To assess the quantitative importance of these experiments, I focus on the allocation of labor and land across crops, agricultural value added per worker, and total welfare gains. I also quantify the effects on the amount of modern inputs per worker used in agricultural production. For each variable, I compute changes as $\hat{x} = x'/x$, where x and x' denote the baseline and counterfactual case, respectively. Welfare gains are measured by obtaining the amount of income that would make the household of each region indifferent between the baseline case and the counterfactual economy, and calculating the average (population-weighted) of both regions. Finally, I use a Fisher price index in the rural region to compare agricultural value added in the baseline economy and the counterfactual cases.

The results from assuming zero trade costs across regions are presented in Table 2.6. Agricultural labor productivity increases by 21%, the ratio of total employment in maize to fruits decreases by 19%, and there is similar reallocation of land across crops (which is equivalent to a reallocation of farmers in the model). These results imply that the eliminating trade costs leads to an improved allocation of farmers across types of crops based on comparative advantage. Additionally, the use of intermediate inputs relative to labor increases by nearly 30%, which has a positive effect on agricultural labor productivity. It is worth mentioning that the population share in agriculture stays almost constant (increases by nearly 2 percentage points). One one hand, more people can move to the city because food is less costly to transport and more intermediate inputs are used in agriculture. On the other hand, enough labor needs to work in fruits farming to satisfy the relatively higher demand (income effect) and individuals do not need to live in the city to consume non-agricultural goods at a lower price. These two effects practically offset each other.

Next, I present more conservative results based on taking the U.S. as a low trade

Table 2.6: Counterfactual: Zero Trade Costs

Changes, relative to baseline (*ratio*)	No Trade Costs ($\tau = 1$)
Agricultural Value Added per Worker	1.21
Maize to Fruits Labor Ratio	0.81
Maize to Fruits Land Ratio	0.79
Intermediates to Labor Ratio (Agriculture)	1.29
Total Welfare	1.21

Notes: Cobb-Douglas technologies imply that changes in intermediates to labor ratio are the same for both crops.

costs benchmark. Price comparisons across regions in Mexico were based on farm-gate prices and prices in wholesale markets. To make the equivalent calculation for the U.S., I compare the farm share of total retail costs in fruits markets with the accumulated cost share of farms, transportation, and wholesale trade using data from the USDA in 2007.[22] The latter implies a farm-price share of 65%, which means that trade costs of fruits in the U.S. are around 55% lower than Mexico.[23] Using this benchmark, I reduce all trade costs by the same proportion. The idea of this experiment would be an improvement in the quality of transportation and storage facilities in Mexico, or the adoption of policies inducing competition in transportation markets, that would reduce trade costs to the U.S. level. I focus on fruits due to data availability, however, these goods are the most sensitive to transportation costs.

[22] Data from www.ers.usda.gov/data-products/food-dollar-series/food-dollar-application.aspx

[23] A farm share of 65% implies that τ_f(US)=1.54. Then, the ratio of trade costs in the U.S. to Mexico can be calculated as $(1 - \tau_f(\text{US}))/(1 - \tau_f(\text{MX}))$=0.45.

Table 2.7: Counterfactual: U.S. Benchmark

Changes, relative to baseline (*ratio*)	U.S. benchmark (55% reduction in trade costs)
Agricultural Value Added per Worker	1.14
Maize to Fruits Labor Ratio	0.83
Maize to Fruits Land Ratio	0.81
Intermediates to Labor Ratio (Agriculture)	1.19
Total Welfare	1.14

Notes: Cobb-Douglas technologies imply that changes in intermediates to labor ratio are the same for both crops.

The magnitude of the results presented in Table 2.7 are large. Agricultural labor productivity increases by 14% and the ratio of employment in maize to fruits decreases by 17%; furthermore, the lower costs of modern inputs increase the intensity with which they are used in the production of agricultural goods by 19%. Therefore, reducing trade costs to the U.S. level would raise agricultural labor productivity in Mexico by allocating more farmers in high-productivity cash crops and increasing the relative amount of modern inputs used in agricultural production. The size of the productivity gains are low compare to the large agricultural productivity gaps in Gollin *et al.* (014a); however, these results are in line with other papers looking at the effects of transportation improvements on agricultural productivity. For example, Sotelo (2018) finds an increase of 16% in average agricultural productivity from paving roads in Peru, and Donaldson (2018) finds that railroad access increased real agricultural income by 16% in colonial India.

Figure 2.6: Productivity Gains from Reducing Trade Costs

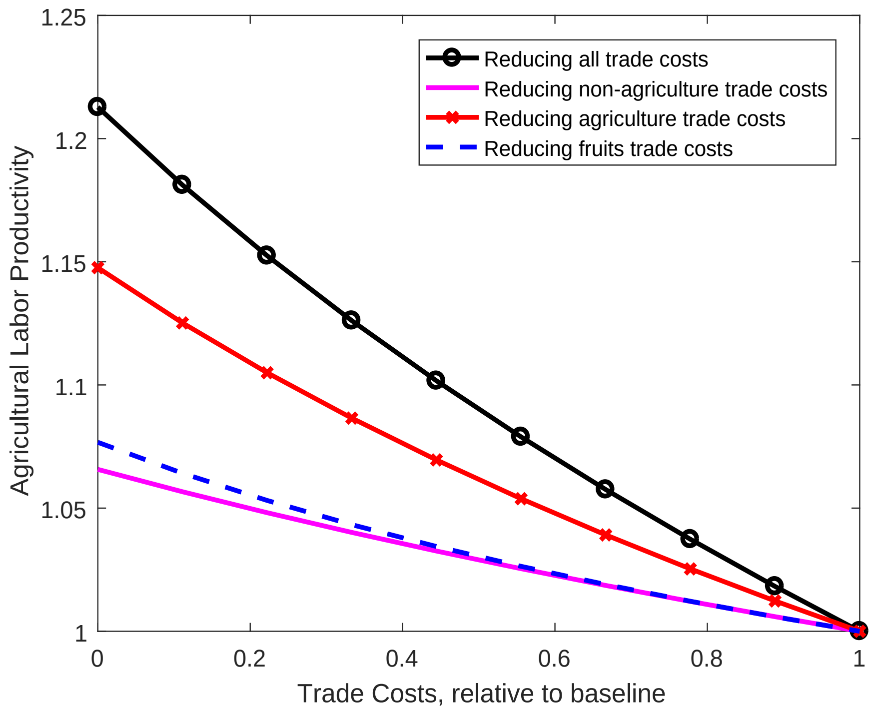

Notes: Agricultural labor productivity refers to value added per worker in agriculture relative to baseline economy. Fruits trade costs are reduced holding trade costs in other sectors constant. Agriculture trade costs refers to both maize and fruits.

To complement the previous experiments, Figure 2.6 shows the relationship between reductions in trade costs and agricultural labor productivity. In addition to decreasing trade costs in every sector, I reduce trade costs in each sector independently. These results suggest that meaningful improvements in transportation costs are needed in order to obtain significant gains in agricultural labor productivity. The latter provides support to the large amounts of resources that developing countries and international organizations allocate to improve transport infrastructure. Moreover, the results show that reducing trade costs of crops has a relatively large effect on agricultural productivity; whereas reducing trade costs of modern inputs, keeping

Table 2.8: Welfare Gains Decomposition: Zero Trade Costs ($\tau_s = 1$)

	Lower spoilage	General equilibrium
Aggregate welfare gains (%)	55.34	44.66

Notes: The welfare gains from lower spoilage are calculated by assuming that all spoilage gains are consumed by the destination region in each sector, keeping all choices fixed. Then, the difference between these and the total gains is defined as the general equilibrium gains.

trade costs of crops constant, has a relatively small effect on agricultural productivity.

The quantitative results imply that there are large welfare gains from eliminating trade costs in the economy. Given that the transportation technology is modeled as an iceberg cost, it is important to distinguish how much of the welfare gains in the model are due to the lower spoilage that results from decreasing these costs and how much is due to general equilibrium effects. To measure the former, I keep every decision of the baseline economy fixed and increase consumption of agricultural and non-agricultural goods in the urban and rural regions by the change in imports when trade costs decrease. That is, I assume that households consume the extra amount of goods that they receive without altering their baseline decisions. The general equilibrium gains are then the total welfare gains minus the welfare gains from lower spoilage in the economy. The decomposition from a counterfactual with zero trade costs is presented in Table 2.8. The general equilibrium gains are large and represent almost half of the total gains. Thus, the gains from reducing trade costs are not only due to the lower spoilage of goods, but to the fact that agents optimally react to trade improvements by reallocating resources across sectors, in particular, across crops within agriculture.

To summarize, reducing trade costs would have a large positive impact on aggregate welfare and agricultural labor productivity in Mexico. The counterfactual results imply that trade costs can account for a large share of the labor allocation across types of crops in this country. Having said that, the motivation of this paper was also based on the fact that many poor countries allocate most of their land to low-productive staple crops, even if yields in many fruits are significantly higher (see Figure 2.1).

To analyze the quantitative implications of the model for poorer countries than Mexico, I recalibrate the baseline economy to match features of the economy in Uganda. In particular, I calibrate the economy-wide productivity parameter and the trade costs to match the share of employment in agriculture and the price gap of fruits across distant regions in that country. The idea is to change the baseline economy in order to get a higher share of employment in agriculture, as in a typical poor country, and be consistent with the low quality of transportation infrastructure.

According to United Nations data from 2014, agriculture accounted for 72% of employment in Uganda. Then, I decrease A to make the modeled economy poor enough so that more people work in agriculture. The latter is a result of non-homothetic preferences in the model. Additionally, Gollin and Rogerson (2010) compute the difference between the wholesale price of Matoke (a variety of banana) at the region of origin and the wholesale price at distant points of sale. The highest ratio of destination to origin price is 4.17, with a distance between points of approximately 500 kilometers (311 miles). While I am taking the highest price ratio according to the data used by those authors, it represents the costs from transporting fruits from the southwest region of Uganda to the North region of the country and, thus, it reflects the quality of transport infrastructure across the country. Similarly to the experiment based on the U.S., I use this number to increase all trade costs in the economy

Table 2.9: Counterfactual: Case of Uganda

Variable	Change (%)
Urban population share	56.6
Slum population (urban share)	-34.5
Total Output	10.3
Welfare	2.6

proportionally.

Once the baseline economy has been calibrated to match the facts described in the previous paragraph, I reduce trade costs to the U.S. level and quantify the effects as was done for the case of Mexico. In this case, trade costs in Uganda are reduced 83% to reach the U.S. level.[24] The results presented in Table 2.9 show that the effects are much larger for a poor country like Uganda. Agricultural labor productivity increases by 191%, the ratio of employment in maize to fruits decreases by more than 70%, and the use of intermediates relative to labor increases by factor of 4.7. The latter reflects the lower cost of modern inputs and the big movement of population from the rural region to the city. These results imply that reducing trade costs in extremely poor countries can increase agricultural labor productivity by releasing individuals from this sector and allocating a larger share of agricultural labor to high-productivity crops.[25]

[24]The magnitude of the difference in trade costs between Uganda and the U.S. is consistent with the findings of Adamopoulos (2011) for differences in transportation costs between rich and poor countries.

[25]The fact that gains from trade are especially high for poor countries is a common result in the trade literature. For example, see the results in Adamopoulos (2011) and Tombe (2015) for rich and

2.5 Conclusions

This paper documents evidence that labor productivity in agriculture is much lower for staple crops than for cash crops. I use microdata from Mexican farms to show that many fruits have a higher labor productivity than maize, yet, the share of employment in the latter is significantly larger. These findings imply that the agricultural productivity puzzle is actually a staple productivity puzzle and focusing on production decisions of farmers is key to understand why agricultural labor productivity is so low in poor countries. One explanation proposed in this paper is that a high share of farmers decides to produce staple crops due to subsistence requirements of staple food and the existence of interregional trade costs in agricultural markets; in particular, the fact that trade costs are larger for cash crops amplifies the selection of farmers into staple crops.

The quantitative experiments imply that trade costs can account for low agricultural productivity and a large share of labor allocations across crops. The ratio of employment in maize to fruits in Mexico would be 17% lower if trade costs were at the U.S. level, and agricultural labor productivity would be 14% higher. These results imply that removing trade barriers in crops markets and reducing costs of modern inputs have the potential to boost economic development in rural areas where agriculture is the predominant activity.

There are alternative and complementary explanations to the one proposed in this paper. For example, switching from staple to cash crops might require large initial investments such as buying a fruit tree or acquiring modern seeds to grow attractive commercial crops; thus, barriers preventing access to capital and input markets may keep too many farmers out of the cash crops sector. Also, farmers need to have

poor countries.

accurate and updated information on crops prices in order to make the best farming decision; thus, barriers to the flow of information might be key to explain the fact that many farmers decide to grow maize, in spite of the large productivity differences. These and other possible explanations are subject of future research.

The results of this paper have important policy implications. First, reducing storage and transportation costs of crops can have large positive results on agricultural labor productivity, especially in poor countries that lack proper infrastructure for intraregional trade. Second, policies should focus on guaranteeing competitive conditions along the supply chain in agricultural markets. Reducing transaction costs and establishing competitive markets seems crucial to allow small and medium scale farmers to enter and grow in profitable agricultural markets.

REFERENCES

Adamopoulos, T., "Transportation Costs, Agricultural Productivity, and Cross-Country Income Differences", International Economic Review **52 (2)**, 489–521 (2011).

Adamopoulos, T. and D. Restuccia, "The Size Distribution of Farms and International Productivity Differences", American Economic Review **104 (6)**, 1667–1697 (2014).

Adamopoulos, T. and D. Restuccia, "Land Reform and Productivity: A Quantitative Analysis with Micro Data", **Working Paper** (2015).

Ahlfeldt, G., S. J. Redding, D. M. Sturm and N. Wolf, "The Economics of Density: Evidence from the Berlin Wall", Econometrica **83**, 2127–2189 (2015).

Alder, S., "Chinese Roads in India: The Effect of Transport Infrastructure on Economic Development", **Working Paper** (2019).

Allen, T. and C. Arkolakis, "Trade and the Topography of the Spatial Economy", The Quarterly Journal of Economics **129**, 3, 1085–1140 (2014).

Alvarez, J. A., "The Agricultural Wage Gap: Evidence from Brazilian Micro-data", **Working Paper** (2018).

Alves, G., "Slum Growth in Brazilian Cities", **CAF Working paper** (2018).

Baker, J., R. Basu, M. Cropper, S. Lall and A. Takeuchil, "Urban Poverty and Transport: The Case of Mumbai", **World Bank Policy Research Working Paper 3693** (2005).

Bertaud, A., "Mumbai FSI conundrum: The perfect storm: the four factors restricting the construction of new floor space in Mumbai", **Working Paper** (2002).

Bertaud, A. and J. K. Brueckner, "Analyzing building height restrictions: predicted impacts and welfare costs", Regional Science and Urban Economics **35**, 109–125 (2005).

Bhan, G., "Evictions, the urban poor and the right to the city in millennial Delhi", Environment and Urbanization **21**, 127–142 (2009).

Brueckner, J. K. and S. V. Lall, "Cities in Developing Countries", **5**, 1399–1455 (2015).

Brueckner, J. K. and K. S. Sridhar, "Measuring welfare gains from relaxation of land-use restrictions: The case of India's building-height limits", Regional Science and Urban Economics **42**, 1061–1067 (2012).

Bryan, G. and M. Morten, "The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia", Journal of Political Economy **Forthcoming** (2018).

Caselli, F., "Accounting for Cross-Country Income Differences", Handbook of Economic Growth **edited by Philippe Aghion and Steven Durlauf, Elsevier Science**, 679–741 (2005).

Cavalcanti, T., D. D. Mata and M. Santos, "On the Determinants of Slum Formation", Economic Journal **forthcoming** (2018).

Davis, M. A. and F. Ortalo-Magne, "Household Expenditures, Wages, Rents", Review of Economic Dynamics **14**, 2, 248–261, URL `https://ideas.repec.org/a/red/issued/09-92.html` (2011).

Diwakar, D. and V. Peter, "Resettlement of Urban Poor in Chennai, Tamil Nadu: Concerns in R&R Policy and Urban Housing Programme", Journal of Land and Rural Studies **4**, 97–110 (2016).

Donaldson, D., "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure", American Economic Review **108 (4-5)**, 899–934 (2018).

Donaldson, D. and R. Hornbeck, "Railroads and American Economic Growth: A Market Access Approach", Quarterly Journal of Economics **131 (2)**, 799–858 (2016).

Donovan, K., "Agricultural Risk, Intermediate Inputs, and Cross-Country Productivity Differences", **Working paper** (2018).

Dupont, V., "Slum Demolitions in Delhi since the 1990s : An Appraisal", Economic and Political Weekly pp. 79–87 (2008).

Eaton, J. and S. Kortum, "Technology, Geography, and Trade", Econometrica **70**, 1741–1779 (2002).

Emrath, P., "Government Regulation in the Price of a New Home. Special Study for Housing Economics", **National Association of Home Builders** (2016).

Epple, D., B. Gordon and H. Sieg, "A New Approach to Estimating the Production Function for Housing", American Economic Review **100**, 3, 905–924, URL `https://ideas.repec.org/a/aea/aecrev/v100y2010i3p905-24.html` (2010).

Gechter, M. and N. Tsivanidis, "Efficiency and Equity of Land Policy in Developing Country Cities: Evidence from the Mumbai Mills Redevelopment", **Working Paper** (2017).

Glaeser, E., "Triumph of the City: How Our Greatest Invention Makes Us Richer, Smarter, Greener, Healthier, and Happier", Penguin Press **London, UK** (2011).

Gollin, D., "Getting Income Shares Right", Journal of Political Economy **110**, 458–474 (2002).

Gollin, D., D. Lagakos and M. E. Waugh, "The Agricultural Productivity Gap", The Quarterly Journal of Economics **129**, 939–993 (2014).

Gollin, D., D. Lagakos and M. E. Waugh, "The Agricultural Productivity Gap", The Quarterly Journal of Economics **129**, 939–993 (2014a).

Gollin, D., D. Lagakos and M. E. Waugh, "Agricultural Productivity Differences across Countries", American Economic Review: Papers and Proceedings **104**, 165–170 (2014b).

Gollin, D. and R. Rogerson, "Agriculture, Roads, and Economic Development in Uganda", **NBER Working Paper No. 15863** (2010).

Gollin, D. and R. Rogerson, "Productivity, transport costs and subsistence agriculture", Journal of Development Economics **107**, 38–48 (2014).

Henderson, J. V., A. J. Venables, T. Regan and I. Samsonov, "Building functional cities", Science **352 (6288)**, 946–947 (2016).

Herrendorf, B., J. A. Schmitz and A. Teixeira, "The Role of Transportation in U.S. Economic Development: 1840-1860", International Economic Review **53 (3)**, 693–715 (2012).

Herrendorf, B. and T. Schoellman, "Why is Measured Productivity so Low in Agriculture?", Review of Economic Dynamics **18**, 1003–1022 (2015).

Herrendorf, B. and T. Schoellman, "Wages, Human Capital, and Barriers to Structural Transformation", American Economic Journal: Macroeconomics **10**, 1–23 (2018).

Hnatkovska, V. and A. Lahiri, "Structural Transformation and the Rural-Urban Divide", Working paper (2014).

Hsieh, C.-T., E. Hurst, C. I. Jones and P. J. Klenow, "The Allocation of Talent and U.S. Economic Growth", Working paper (2018).

Hsieh, C.-T. and P. J. Klenow, "Misallocation and Manufacturing TFP in China and India", Quarterly Journal of Economics **124 (4)**, 1403–1448 (2009).

Hsieh, C.-T. and E. Moretti, "Housing Constraints and Spatial Misallocation", American Economic Journal: Macroeconomics **Forthcoming** (2018).

IHDS II, "Sonalde Desai and Amaresh Dubey and Reeve Vanneman", **University of Maryland and National Council of Applied Economic Research, New Delhi, 2015. Ann Arbor, MI: Inter-university Consortium for Political and Social Research** (2012).

Jedwab, R. and D. Vollrath, "The Urban Mortality Transition and Poor Country Urbanization", American Economic Journal: Macroeconomics **11 (1)**, 223–275 (2019).

Jerven, M., "Poor Numbers: How We Are Misled by African Development Statistics and What to Do about It.", **Ithaca, NY**, Cornell University Press (2013).

KPMG, "Decoding housing for all by 2022", (2014).

Kumar, P., "Declining Number of Slums: Nature of Urban Growth", Economic and Political Weekly **XLV No 41**, 75–77 (2010).

Kuznets, S., "Modern Economic Growth: Rate, Structure, and Spread", Yale University Press **New Haven, CT** (1968).

Lagakos, D., M. Mobarak and M. E. Waugh, "The Welfare Effects of Encouraging Rural-Urban Migration", **Working Paper** (2018a).

Lagakos, D., B. Moll, T. Porzio, N. Qian and T. Schoellman, "Life Cycle Wage Growth across Countries", Journal of Political Economy **126**, 2, 797–849, URL `https://ideas.repec.org/a/ucp/jpolec/doi10.1086-696225.html` (2018b).

Lagakos, D. and M. E. Waugh, "Selection, Agriculture, and Cross-Country Productivity Differences", American Economic Review **103**, 948–980 (2013).

Marx, B., T. Stoker and T. Suri, "The Economics of Slums in the Deveolping World", Journal of Economics Perspectives **27**, 187–210 (2013).

Monge-Naranjo, A., P. C. Ferreira and L. T. de Mello Pereira, "Of Cities and Slums", **Working Paper** (2018).

Monte, F., S. Redding and E. Rossi-Hansberg, "Commuting, Migration and Local Employment Elasticities", American Economic Review **Forthcoming** (2018).

Munshi, K. and M. Rosenzweig, "Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap", American Economic Review **106**, 46–98 (2016).

OECD, "Agricultural and Fisheries Policies in Mexico", **Organization for Economic Cooperation and Development** (2007).

Restuccia, D., D. T. Yang and X. Zhu, "Agriculture and aggregate productivity: A quantitative cross-country analysis", Journal of Monetary Economics **55**, 234–250 (2008).

Rossi-Hansberg, E., R. O. III and P.-D. Sarte, "Rethinking Detroit", **Working Paper** (2017).

Sotelo, S., "Domestic Trade Frictions and Agriculture", **Working Paper** (2018).

Tombe, T., "The Missing Food Problem: Trade, Agriculture, and International Productivity Differences", American Economic Journal: Macroeconomics **7 (3)**, 226–258 (2015).

Tsivanidis, N., "The Aggregate and Distributional Effects of Urban Transit Infrastructure: Evidence from Bogots TransMilenio", **Working Paper** (2018).

United Nations, "Slum Almanac 2015/2016. Tracking Improvement in the Lives of Slum Dwellers", **UN-HABITAT** (2015).

Vollrath, D., “How Important are Dual Economy Effects for Aggregate Productivity?”, Journal of Development Economics **88**, 325334 (2009).

World Bank, “Urbanization beyond Municipal Boundaries. Nurturing Metropolitan Economies and Connecting Peri-Urban Areas in India”, **Directions in Development. Washington, DC: World Bank. doi:10.1596/978-0-8213-9840-1.** (2013).

Young, A., “Inequality, the Urban-Rural Gap, and Migration”, Quarterly Journal of Economics **128 (4)**, 17271785 (2013).

Young, A., “Structural Transformation, the Mismeasurement of Productivity Growth, and the Cost Disease of Services”, American Economic Review **104 (11)**, 3635–3667 (2014).

APPENDIX A

URBAN DIVISION IN MUMBAI

The following map shows the classification of Mumbai sub-districts into urban center or periphery used in the calibration of the model described in Section 1.5.1. The periphery are those sub-districts which are located 20 or more kilometers away from the closest CBD. The distances were obtained using Google Maps. Bandra Kurla Complex is the main CBD of the city today. Previously, it was Nariman Point, which is located in the south of the city.

Figure A.1: Mumbai Sub-districts (Wards) Division

APPENDIX B

DATA AND EMPIRICAL DETAILS

This section provides more details of the empirical evidence presented in Section 2.2. Using data from the ENA 14 surveys, the steps to calculate value added for each crop that is produced in a farm are the following. First, to obtain the value of output I multiply the volume of harvested output by the farm-gate price reported by the farmer. Many farmers do not report a price because they did not sell any output in that period of time, especially those who produce maize; for such cases, I use the average price of the crop in the municipality where the farm is located or, in cases where there is no data to compute the latter, the average price of the crop in the state. I eliminate outliers (0.5% of each tail) to compute those average prices. The next step is obtaining the value of intermediate inputs used in the production of each crop. Different categories of farm expenses are reported at the crop level; however, some categories like soil preparation and sowing may include payments to capital and labor. Therefore, I only consider expenditures on modern inputs that do not include any payments to factors of production: fertilizers; chemicals and pesticides; and irrigation. Finally, I subtract the value of production that is used for seed and feed in the farm from the value of total output.

To estimate value added per hectare, I take the amount of harvested hectares reported by the farmer for each crop. Estimating value added per worker involves additional steps. The number of workers (owner, family members and hired labor) are reported at the farm level and, thus, they might be used in the production of more than one crop within a farm. Then, I focus on farms that only produce one type of crop, otherwise there is not an obvious way to allocate labor to different crops produced in a farm. To define farms producing a single crop, I aggregate the different varieties of the crops considered in Section 2.2.2 into one category; for example, all varieties of chili pepper are considered as one type of crop. Under such considerations, from the total number of farms producing maize, fruits or both, only 5.8% of them produces maize and one of the fruits. Thus, by focusing on farms that produce one type of crop I only lose a small share of farms.[1] However, the farms omitted from the estimation of labor productivity might grow any number of different crops varieties and that is reflected in the difference of farm-crop observations between land and labor productivity gaps reported in Table 2.1.

I now describe the farm price data used in the estimation of trade costs. For crops, I use prices reported by farmers and compute the average price at the state level after eliminating outliers (0.5% of each tail). In the case of fertilizers, farmers do not report the actual price they paid, instead they report total quantities of both natural and chemical fertilizers used for crop production, as well as total expenses on fertilizers. Since I cannot split the latter between natural and chemical fertilizers, I compute the fertilizer price as total expenditure divided by total quantity. This procedure results in a distribution of fertilizers prices with fat tails. From observation and comparison with public market prices, I eliminate outliers to compute average prices at the state and municipality level (2.5% of each tail). In Section 2.2.2 of the main text, I focus on cases where the chemical fertilizer used by the farmer was Urea or Ammonium Sulfate, which account for 45% of the observations that reported fertilizers.

As described in the main text, the second most important data source used for the empirical evidence comes from the SNIIM. This is a government website that provides information of market prices in primary sectors of the economy. I build a

[1] In general, the surveyed farms produce less than two (1.7) different varieties of crops on average.

dataset with the monthly price of varieties of fruits and grains in the main wholesale markets of the country (usually located in the capital of a state). For each crop in a particular market, both the price (per kilogram most of the times) and state of origin are reported. It is worth mentioning that downloading and processing these data requires a non-trivial amount of time in order to produce a database of market prices of crop varieties in every state. Additionally, SNIIM provides data on market prices of fertilizers throughout the country. I use this to obtain the price of Urea and Ammonium Sulfate in the possible state of origin, either a port or the state where a production plant is located. These locations are obtained from a 2006 report made by ANACOFER, a national association of production and distribution of fertilizers.

Finally, I use public data from national accounts. Particularly, I use the 2008 Input-Output matrix that has data on value added, inputs expenditures, and total employment at the six-digit industry level, including 50 types of crops. I also use agricultural data aggregated to the state and national level from the *Sistema de Información Pesquera y Alimentaria* (SIAP) of the *Secretaría de Agricultura, Ganadería, Desarrollo Rural, Pesca y Alimentación* (SAGARPA). These data on production, prices, yields, and land is collected by governments offices located in many localities throughout the country and is available from 1980 to 2014.

APPENDIX C

REGIONAL PRODUCTION OF CROPS

This section documents evidence that farmers in most regions of Mexico have the possibility to produce cash crops with a higher productivity than maize. The goal is not to determine the optimal group of crops that a particular region should grow. The latter would depend in agro-ecological conditions of each subregion in the country. The aim is to show that in most localities within states there exists production of high-productivity fruits, so farmers are not always forced to grow staple crops by conditions related to climate or quality of soil in a particular area. That is, the fact that most labor and land is allocated to maize is a choice.

I look at the geographical distribution of production and yields, measured as output value per hectare, within states in Mexico using data from SIAP, which is the most comprehensive source of aggregate agricultural data. To simplify the analysis, I focus on a particular group of crops that includes the types of fruits considered in the main text: avocado, banana, chili pepper, cucumber, mango, papaya, tomato, and watermelon; and on those states with a high level of poverty and large share of agricultural employment. These states located in the central-south region are characterized by small-scale farming. Every state is divided in municipalities that cover multiple towns or cities. There are 2,458 municipalities in the country and 53% are located in the six states considered.

Figure C.1 presents the distribution of yields across municipalities in 2014. I calculate the average yield relative to maize weighted by harvested hectares. According to these maps, the majority of municipalities produce fruits with yields that are more than two times higher than those of maize. Furthermore, there are no particular subregions that seem to be limited with respect to the rest. Only Chiapas and Oaxaca have relatively big municipalities with no data, but these are surrounded by localities with presence of high productivity fruits. Thus, even if some subregions in these states seem to have higher levels of productivity, there is no evidence to conclude that other subregions are significantly constrained to produce alternative crops to maize.

In addition, to provide a general picture of the country, Figure C.8 presents a map with the fraction of municipalities in each state that produce at least one cash crop (from the group considered) with a higher yield than maize. This figure suggests that in most states of Mexico, farmers in over 50% of municipalities produce fruits that render a higher output value per hectare than maize. The states with the highest geographical concentration of production are in the north.[1] The main takeaway from these facts is that in most regions of the country farmers have the possibility to produce fruits that have a higher yield than maize. While specific subregions may be able to attain higher yields, this does not imply that staple crops like maize are the only feasible option in other areas.

[1] A possible explanation for the spatial concentration in the north is that most of this region is arid and agro-climatic conditions are less favorable; moreover, in comparison to the south, these states are characterized by large commercial farms.

Figure C.1: Distribution of Relative Fruit Yields in Poor States (Maize=1)

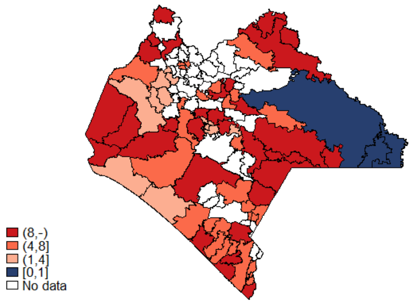

Figure C.2: Chiapas

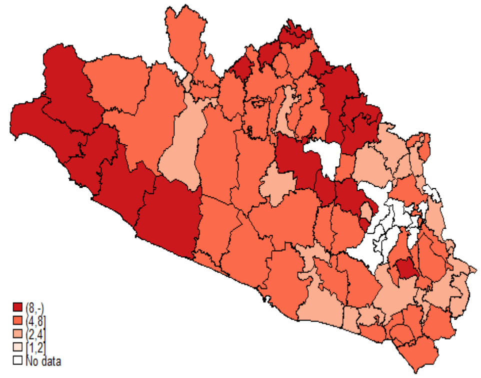

Figure C.3: Guerrero

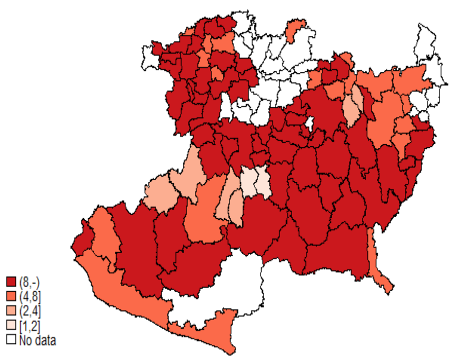

Figure C.4: Michoacan

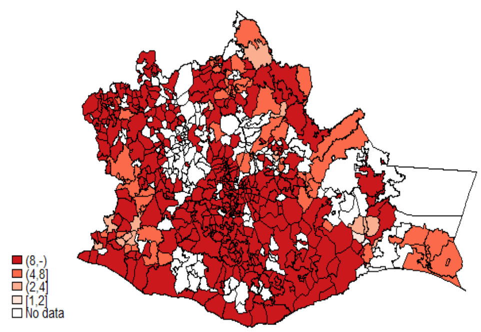

Figure C.5: Oaxaca

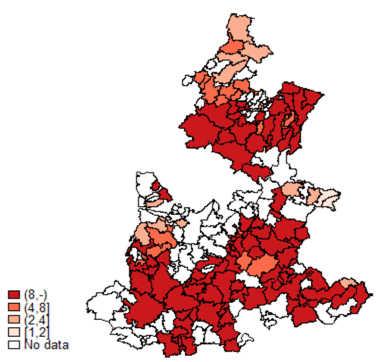

Figure C.6: Puebla

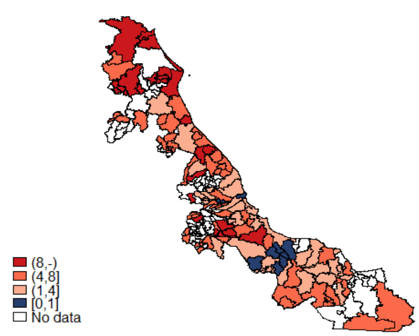

Figure C.7: Veracruz

Note: Average yields are weighted by harvested hectares.
Source: Author's estimates using SIAP data, 2014

Figure C.8: Share of Municipalities with High-productivity Fruits

Source: Author's estimates using SIAP data, 2014

APPENDIX D

AGRICULTURAL PRODUCTIVITY WITH AGGREGATE DATA

Table D.1: Productivity Gaps and Employment Shares of Crops

Crops	Value added per worker relative to non-agriculture	Employment share of total agriculture
Maize	0.15	23.1%
Other grains	0.19	15.8%
Top cash crops	0.40	8.2%
Agricultural Sector	0.16	100%

Notes: Top cash crops includes avocados, tomatoes, chili peppers, other vegetables, and other non-citrus fruits and nuts. These crops account for 64% of total exports in agriculture and 54% of Value Added in fruits and vegetables farming. Products are classified according to the NAICS.
Source: Author's calculations using Input-Output Data from Mexico, 2008.

This section complements the results from Section 2.2.1 using aggregate data of Mexican agriculture. I compute value added per worker for different industries using Input-Output data from 2008. First, Table D.1 presents labor productivity relative to non-agriculture for three categories of crops. Value added per worker in maize and other grains is less than half of value added per worker in cash crops production. Moreover, the productivity gap between agriculture and non-agriculture is almost the same as the one between maize and non-agriculture. The latter implies that the large agricultural productivity gap is actually measuring large productivity differences with respect to unproductive staple crops that have the largest employment share.

Second, Panel A in Figure D.1 shows value added per worker relative to maize for different categories of cash crops. These results confirm that there are large differences in labor productivity between fruits and maize. Moreover, Panel B shows that labor allocated to maize is much higher than any of these cash crops. Overall, these facts are consistent with the ones presented in section 2.2.1 using farm data.

Finally, to explore if the year of the ENA 2014 surveys was important for the empirical results, I compute log-yields of crops using state-panel data of agricultural production and prices in Mexico from 1980 to 2012. For each state in every year, yields are expressed in units of maize using relative prices. These yields are detrended using a linear regression with respect to time, taking 2012 as the base year. Figure D.3 shows the non-parametric densities of crop yields. According to this data, the average log-yield of each of these fruits is higher than the average yield of maize; in some cases, like with tomatoes, the difference in average log-yield is around 4. These results imply that the yield distribution of many fruits first-order stochastically dominates the yield distribution of maize. Thus, this data suggests that 2014 was not a special year and the productivity gaps between types of crops are persistent over time. If any, this motivates future research to explore other possible barriers.

Figure D.1: Panel A. Value Added per Worker Relative to Maize

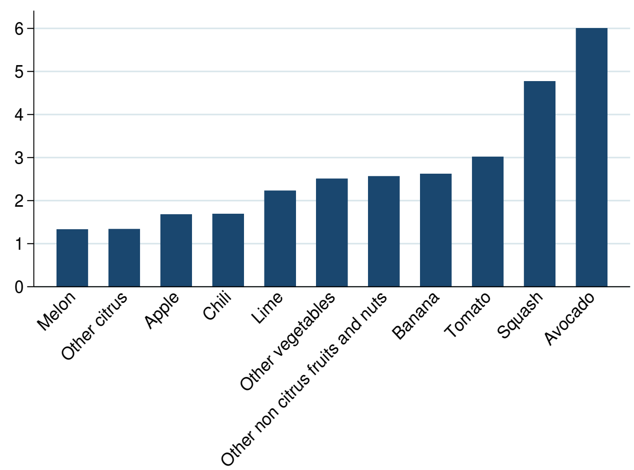

Figure D.2: Panel B. Total Workers relative to Maize

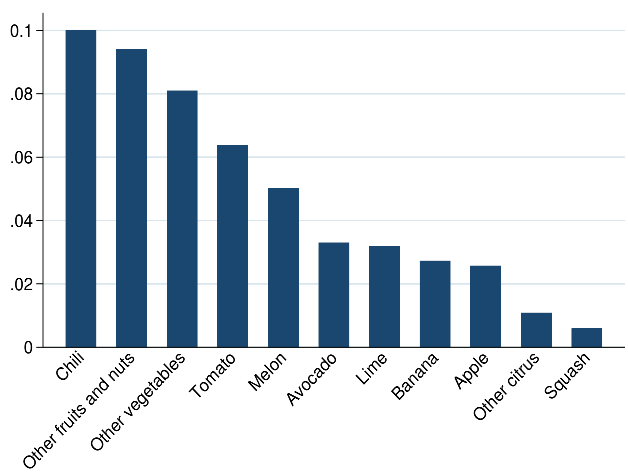

Notes: Maize is normalized to 1 in both cases.
Source: Author's calculations using National Accounts Data from INEGI: Input-Output Data. Mexico 2008.

Figure D.3: Crop Yields Distributions

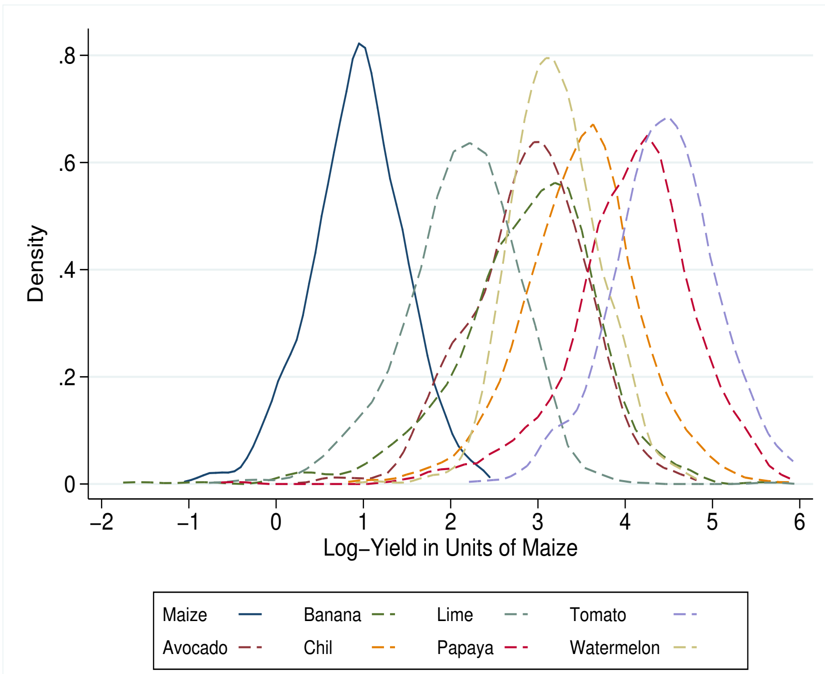

Notes: Kernel densities (Epanechnikov). Log-yields are detrended using a linear regression and taking 2012 as the base year.
Source: Author's estimates using data from SIAP. State panel from 1980 to 2012.

APPENDIX E

SUBSISTENCE PRODUCTION OF CROPS

In this section I provide evidence to support the assumption in the model regarding subsistence consumption of staple food. Staple crops like maize and rice are the main food component of a population's diet, especially in poor countries. A key distinction between staple crops and cash crops is that a relatively high share of staple crops production is used for subsistence requirements of food, while most cash crops production is sold to richer regions, within and outside of a country.

To analyze subsistence requirements by crop, Table E.1 shows the average share of farm production used for family consumption. I distinguish between farms of all sizes and farms with less than 20 hectares. The results show that the share of maize production used for subsistence is significantly higher than any other fruit. While more than one third of production of maize is use for family consumption, the range for other cash crops is between 2% and 12% of production. In fact, for most of these fruits, less than 6% of the production is used for subsistence. This pattern is the same for both groups of farms, so it is not exclusive of small-scale farming. These findings support the assumption in the model that subsistence food requirements only apply to staple crops.

Table E.1: Farm Production Used for Family Consumption by Crop

Farms	Maize	Watermelon	Mango	Avocado	Chili	Papaya	Cucumber	Tomato
All sizes	30.2%	6.0%	5.7%	5.4%	5.3%	3.6%	3.0%	2.0%
Less than 20 ha.	36.1%	9.6%	6.5%	6.0%	9.1%	4.8%	4.1%	2.8%

Notes: Average farm production by crop.
Source: Author's estimates using data from INEGI: NSA 2014, Mexico.