



Essays in Urban Economics and Development

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:40050080>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Essays in Urban Economics and Development

A dissertation presented

by

Juan Pablo Chauvin

to

The Department of Public Policy

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Public Policy

Harvard University

Cambridge, Massachusetts

May 2018

© 2018 Juan Pablo Chauvin

All rights reserved.

Dissertation Advisor:
Professor Edward L. Glaeser

Author:
Juan Pablo Chauvin

Essays in Urban Economics and Development

Abstract

This dissertation comprises three essays at the intersection of Urban and Development Economics. The first chapter explores whether well-known facts about urbanization in the United States also hold in three large developing countries: Brazil, China and India. I find that the forces that drive urban success are generally similar in rich and poor countries, with stronger agglomeration economies and correlations between education levels and city growth among the latter. Predictions based on the spatial equilibrium assumption tend to hold in the U.S., Brazil and China, but not in India. The second chapter studies how local economies in Brazil react to male-leaning vs. female-leaning labor demand shocks. I find that, while shocks that favor male employment lead to population growth and higher housing rents, the same is not true for shocks that favor female employment. I propose a spatial equilibrium framework that illustrates how, in a context with free mobility and gender segmentation in the labor market, joint mobility constraints of married couples can account for the observed patterns. The third chapter analyzes the effects of local increases in public education spending on labor market outcomes of individuals and regions. Using FUNDEF -a large federal program that redistributed education budgets across Brazilian municipalities in the late 1990s- as a source of exogenous variation, I find that education spending led to better labor market outcomes for individuals, mainly by increasing their likelihood of migrating to more productive places. The effects on regional outcomes were largely negative, and appear to be explained by sluggish local demand for educated labor.

Contents

Abstract	iii
Acknowledgments	xi
Introduction	1
1 What is Different About Urbanization in Rich and Poor Countries? Cities in Brazil, China, India and the United States	4
1.1 Introduction	4
1.2 Measuring Urban Areas in Four Countries	11
1.2.1 Defining Agglomerations	13
1.2.2 The Distribution of Populations across Area Sizes	15
1.3 City Size Distributions: Zipf's Law and Gibrat's Law	17
1.3.1 Zipf's Law across Four Countries	19
1.3.2 Gibrat's Law across Four Countries	23
1.4 Spatial Equilibrium	25
1.4.1 The Relationship between Prices and Wages	26
1.4.2 Real Wages and Amenities	33
1.4.3 Happiness across Space	36
1.4.4 Mobility in the Four Countries	39
1.5 The Determinants of Urban Success: Agglomeration and Human Capital . .	42
1.5.1 Agglomeration Economies	43
1.5.2 Human Capital Externalities	51
1.5.3 Urban Growth and Human Capital	57
1.6 Conclusion	62
2 Gender-Segmented Labor Markets and the Effects of Local Demand Shocks	65
2.1 Introduction	65
2.2 Spatial equilibrium with gender-segmented labor markets	70
2.2.1 Production and labor demand	72
2.2.2 Household utility	73
2.2.3 Labor force participation	75

2.2.4	The housing market	75
2.2.5	Key insights and predictions	76
2.3	Data, Facts and Identification Strategy	80
2.3.1	Data	80
2.3.2	Descriptive facts	81
2.3.3	Identification strategy	87
2.4	Results	92
2.4.1	Gender-specific demand shocks and household migration	93
2.4.2	Effects on population and housing rents	96
2.4.3	Employment effects	98
2.4.4	Wages and participation effects	100
2.4.5	Robustness	104
2.5	Welfare and policy implications	106
2.6	Conclusion	107
3	Education and Local Labor Market Outcomes. Evidence from a Large Federal Program in Brazil	109
3.1	Introduction	109
3.2	The FUNDEF program and its context	114
3.3	Measures, data and identification	119
3.3.1	Measures and data	120
3.3.2	Identification of individual effects	121
3.3.3	Identification of regional effects	123
3.4	Individual-level results	124
3.4.1	Effects on individual educational attainment	124
3.4.2	Effects on likelihood of migrating	128
3.4.3	Effects on individual labor outcomes	131
3.4.4	Mechanisms	135
3.5	Regional-level results	141
3.5.1	Effects on regional educational attainment levels	142
3.5.2	FUNDEF, migration, and the educational attainment levels of regions	145
3.5.3	FUNDEF and regional labor market outcomes	147
3.6	Conclusion	149
	References	151
	Appendix A Appendix to Chapter 1	166
A.1	Supplementary Figures for Chapter 1	166
A.2	Supplementary Tables for Chapter 1	171

Appendix B	Appendix to Chapter 2	173
B.1	Supplementary Figures for Chapter 2	173
B.2	Supplementary Tables for Chapter 2	178
B.3	Model solutions appendix	187
B.3.1	Equilibrium under autarky	187
B.3.2	Equilibrium in the open region	189
Appendix C	Appendix to Chapter 3	192
C.1	Supplementary Figures for Chapter 3	192
C.2	Supplementary Tables for Chapter 3	197
Appendix D	Data Appendix	204

List of Tables

1.1	Share of people living in urban areas of different size	15
1.2	Income distributions, 1990-2010	18
1.3	Urban population Kolmogorov–Smirnov two-sample tests	22
1.4	Gibrat’s Law: Urban population growth and initial urban population	24
1.5	Regressions of local prices on wages, 2010	28
1.6	Climate amenities regressions, 2010	34
1.7	Percentage of the population living in a different locality five years ago	40
1.8	Income and agglomeration, 2010	45
1.9	Real income and agglomeration, 2010	49
1.10	Human capital externalities, 2010	54
1.11	Human capital and growth, 1980-2010	59
2.1	Effects of gender-specific demand shocks on migrant population	94
2.2	Effects of gender-specific demand shocks on population	96
2.3	Effects of gender-specific demand shocks on rents	97
2.4	Effects of gender-specific demand shocks on employment	99
2.5	Effects on the employment gap	100
2.6	Effects of gender-specific demand shocks on wages	101
2.7	Effects of gender-specific demand shocks on non-participant population	102
2.8	Effects on wage gaps	104
3.1	Enrollment in Fundamental Education in Brazil in 1997, by system	116
3.2	Effects of FUNDEF on probability of reaching of reaching a specific educational attainment	128
3.3	Individual effects of FUNDEF on migration	131
3.4	Effects of FUNDEF on individual labor market outcomes	133
3.5	Effects of FUNDEF on local education attainment	144
3.6	Effects of FUNDEF on regional population	146
3.7	Effects of FUNDEF on local labor market outcomes	148
A.1	Happiness Regressions	171

A.2	Local prices and agglomeration, 2010	172
B.1	Summary statistics, 1990s	178
B.2	Summary statistics, 2000s	179
B.3	Correlations, 1990s	180
B.4	Correlations, 2000s	181
B.5	Population and work in Brazil between 1991 and 2000	182
B.6	Geographic Mobility working-age population in Brazil 1991-2000	183
B.7	Migrant population and work in Brazil, 2000	184
B.8	Gender-specific Bartik shocks and start-year characteristics	185
B.9	Pre-trends tests	186
C.1	Individual summary statistics, 2010	197
C.2	Regional summary statistics, 2010	199
C.3	Regional summary statistics, 2000	200
C.4	Correlations among regional-level variables, 2010	201
C.5	Correlations among regional-level variables, 2000	202
C.6	Effects of FUNDEF on migrant population by educational attainment group	203
C.7	Effects of FUNDEF on local education attainment controlling for 1990s' migration composition	203
D.1	Databases Used	204
D.2	Country-level variables (Chapter 1)	206
D.3	Individual-level variables (Chapter 1)	207
D.4	Individual-level variables (Chapters 2 and 3)	212
D.5	Area-level variables (Chapter 1)	214
D.6	Area-level variables (Chapters 2 and 3)	220

List of Figures

1.1	Share of total population living in urban areas, 1960-2014	12
1.2	Zipf's Law. Urban populations and urban population ranks, 2010	21
1.3	Income and rents, 2010	29
1.4	Happiness and income levels	37
1.5	University graduates share and population growth 1980-2010	60
2.1	Distribution of labor income across local labor markets in Brazil and the U.S.A.	82
2.2	Employment growth by industry and gender, Brazil 1991-2010	84
2.3	Changes in Employment and the Gender Wage Gap	86
2.4	Female Bartik shocks and base year female non-participation	91
3.1	Percentage of Brazilian population age 23 or older in each educational attainment category	117
3.2	Changes in labor market outcomes 2000-2010 by educational attainment category	118
3.3	Effects of FUNDEF on probability of reaching a specific educational attainment in 2010 by cohort	125
3.4	Effects of FUNDEF on probability of being a migrant in 2010 by cohort and gender	129
3.5	Effects of FUNDEF on wages and labor force participation	132
3.6	Effects of FUNDEF on wages and labor force participation by migrant status	136
3.7	Effects of education attainment on labor market outcomes (2SLS estimates) .	139
3.8	Effects of FUNDEF on growth of the share of educated people among adults	142
A.1	Income and housing prices, 2010	166
A.2	Urban population and income residuals, 2010	167
A.3	Population density and income residuals, 2010	168
A.4	Happiness and population size	169
A.5	University graduates share and wage residuals, 2010	170
B.1	Distributions of Gender-specific Bartik shocks	173
B.2	Labor force participation across microregions in Brazilian macro-regions, 2000	174

B.3	Geographic distribution of gender-specific Bartik shocks, Brazil 1991-2000 . .	175
B.4	Effects of gender-specific shocks on migrant population by age	176
B.5	Effects of gender-specific shocks on the employment gap, predictive margins	177
B.6	Effects of gender-specific shocks on the wage gap, predictive margins	177
C.1	Changes in labor market outcomes 2000-2010 by educational attainment category and gender	192
C.2	Distribution of FUNDEF shock across localities	193
C.3	Effects of FUNDEF on individual's probability of reaching a specific educa- tional attainment in 2010 by cohort and by gender	193
C.4	Effects of FUNDEF on wages using alternative measures	194
C.5	Effects of FUNDEF on wages and labor force participation by gender	195
C.6	Effects of FUNDEF on wages by migrant status using alternative measures .	196
C.7	Percentage of Brazilian population age 23 or older in each educational attain- ment category by migrant status, 2010	196

Acknowledgments

I am deeply grateful to my academic advisors, family, friends and colleagues. The completion of this work would not have been possible without their support and the inspiration they provided.

First, I thank the members of my committee for their guidance and encouragement. Ed Glaeser, my committee chair, has been an exceptionally generous mentor, teacher and friend; and his many lessons -from those on economic modeling to those on work-family balance- will always stay with me. Melissa Dell has been a source of inspiration from day one, and her advice has shaped my approach to empirical work and the way I communicate information and ideas. Ricardo Hausmann has constantly pushed me to think outside the box, challenged me to make my work relevant in the real world, and inspired me to approach it with passion, honesty, and a sense of humor.

Just as important has been the continuous support of my family. My wife Jasmina and I have navigated together, over the past few years, the adventures of our doctoral programs and of parenthood. She has been my partner, my most trusted collaborator, and my best friend. I am grateful to her and to our children, Sofia and Adrian, for their love and for inspiring me to try to be better every day. My parents, Rigo and Noshy, have been the most influential force in my life, and I will never stop thanking them for their relentless love and support. I also want to thank the rest of my family in Ecuador and California for bringing so much joy to my life, and to my family in Cambridge, Susana, Clark, Jamil and, Annita, for making me feel at home all these years.

I am also grateful to those who contributed to this work with insightful conversations, comments and suggestions, including Alberto Abadie, Sam Asher, Carlos Azzoni, Ricardo Paes de Barros, Daniel Da Mata, Dario Diodado, Sergio Firpo, Claudia Goldin, Rema Hanna, Stephan Heblich, Adriana Kugler, Ljubica Nedelkoska, Naercio Meneses Filho, Frank Neffke, Rohini Pande, André Portela, Daniel Shoag, William Strange, Hildegunn Stokke, Matthew Turner, Gabriel Ulyssea, an anonymous referee, and seminar participants at the Economic

Development Workshop, Labor Economics Workshop, and CID-Growth Lab Seminar at Harvard, Urban Economics Association 12th annual meeting, Harvard-Dartmouth Cities and Development mini-conference, Universidade de São Paulo, IPEA, the IADB Research Department, the Economics Departments at Syracuse University and Florida International University, the Real Estate Department at CUNY, the McCourt School at Georgetown, the Ross School at the University of Michigan. I would also like to thank Fernanda Estevan, Brian Kovak and Christopher Poliquin for sharing data.

Finally, I would like to gratefully acknowledge the support from the Harvard Center for International Development, the Weiss Family Fund for research in Development Economics, and the Harvard-Brazil Cities initiative from the David Rockefeller Center for Latin American Studies.

To my parents, who showed me different paths, let me choose, and always found a way to walk along with me.

Introduction

By the year 2100 the planet will have, according to forecasts, twice as many people living in cities as it does today. The lion's share of this transformation will take place in the developing world, where the urban population could go from 2.6 billion to close to 8 billion in a century (Fuller and Romer 2014.) What will this new chapter in the history of urbanization look like? Will it be a "re-make" of what we saw in the 20th century in now-rich countries? Will it come accompanied by stark improvements in living conditions as it did before? How will this depend on the actions of local and national policymakers over the next couple of decades?

These are broad and complex questions, that have major policy implications. They are ultimately about the economic opportunities that will be available -or not- for billions of people. And about who among them will have access to those opportunities. Yet we are still far from having enough satisfactory answers. Urban Economics has traditionally focused on the U.S. and other high-income countries (Glaeser and Henderson 2017), and we are only beginning to explore what is similar and what is different about urbanization in high- and low-income countries. My dissertation is a contribution towards closing this knowledge gap.

In the first chapter, co-authored with Edward Glaeser, Yurean Ma and Kristina Tobio, I assess whether the well-known facts about urbanization in the United States also true for the developing world. The essay compares American metropolitan areas with analogous geographic units in Brazil, China and India. Both Gibrat's Law and Zipf's Law seem to hold as well in Brazil as in the U.S., but China and India look quite different. In Brazil and

China, the implications of the spatial equilibrium hypothesis, the central organizing idea of urban economics, are not rejected. The India data, however, repeatedly rejects tests inspired by the spatial equilibrium assumption. One hypothesis is that spatial equilibrium only emerges with economic development, as markets replace social relationships and as human capital spreads more widely. In all four countries there is strong evidence of agglomeration economies and human capital externalities. The correlation between density and earnings is stronger in both China and India than in the U.S., strongest in China. In India the gap between urban and rural wages is very large, but the correlation between city size and earnings is more modest. The cross-sectional relationship between area-level skills and both earnings and area-level growth are also stronger in the developing world than in the U.S. The forces that drive urban success seem similar in the rich and poor world, even if limited migration and difficult housing markets make it harder for a spatial equilibrium to develop.

In the second chapter I explore how segmentation by gender in the labor markets shape the way local economies react to local labor demand shocks. Gender segmentation in the labor market is widespread. However, most existing studies of the effects of labor demand shocks on local economies assume away gender. In this paper, I show that local labor demand shocks can lead to different outcomes depending on whether they favor male or female employment. I develop a spatial equilibrium model that features gender segmented labor markets and joint mobility frictions, which predicts that couples are more likely to migrate in response to male opportunities. As a result, positive shocks to local labor demand for men lead to population growth, increases in female labor supply, and housing demand growth. Meanwhile, equivalent shocks to labor demand for women lead to smaller inflows of migrant workers, and labor force participation is a relatively more important margin of adjustment in this case. I find strong empirical support for the model's predictions in the context of Brazil during 1991-2010. Comparing the effects of gender-specific labor demand shocks, I show that male shocks produce a higher migratory response and make localities more populated and expensive. These results imply that place-making policies that create jobs for females are more likely to benefit residents while those that create male jobs are

more likely to benefit immigrants and landlords.

The third chapter studies, also in the Brazilian context, the effect of increasing local spending in public education on individual and regional labor market outcomes. In principle, education could lead to higher local productivity, but potential benefits to local economies could be muted if the educated workers leave in search of better opportunities, or if shifts in the supply of skills outpace demand growth. I use a large program that redistributed public education finance across Brazilian municipalities (FUNDEF) as a source of exogenous variation to empirically study the effects of expansions in public education expenditure on attainment and labor market outcomes at the individual and the local economy levels. The program was successful at improving educational attainment levels for individuals and regions, specially at the primary school level. For individuals, education led to higher wages -mainly by enabling workers to migrate to more productive places- but my estimates of returns to schooling turn negative when I control for region-of-work characteristics. For regions, the program worsened wages and other labor market outcomes but not employment, suggesting that the increased supply of educated workers outpaced demand growth.

These contributions are part of a still-small but growing body of studies focusing on the economy of developing-country cities. The questions that remain open are vast, and the stakes of getting the answers right are high. If forecasts are accurate, close to 600 million people will live in Indian cities by 2030 (McKinsey Global Institute 2010), China will to go from 50% to 70% urban over the same period, adding 350 million residents to cities (World Bank 2014), and two-thirds of the urban infrastructure in Africa will be built between now and 2050 (Collier and Venables 2016.) Getting urbanization right could substantively improve the lives of millions of people in poor countries in just a few decades. Getting it wrong could seriously constraint the development opportunities in those places for generations.

Chapter 1

What is Different About Urbanization in Rich and Poor Countries? Cities in Brazil, China, India and the United States ¹

1.1 Introduction

The majority of the world's urban population will soon live in places that are far poorer than the U.S. and Europe. This creates a knowledge mismatch, for urban economists have predominantly focused on the cities of the wealthy west. The relevance of the long literatures on wealthy world urbanization depends on the similarity between poor world urbanization and rich world urbanization. This paper asks whether the major stylized facts about cities in the U.S. also hold for Brazil, China and India.

Economists frequently assume that our models work everywhere, although different levels of income and education may create marginal differences. Yet the enormous social and political differences between the U.S. and countries like Brazil, India and China may belie that assumption. For example, the central organizing model of urban economics is the spatial equilibrium hypothesis, which in its standard version assumes free mobility across metropolitan areas. Does that assumption make sense in a country like China,

¹Co-authored with Edward L. Glaeser, Yurean Ma and Kristina Tobio

which historically imposed legal barriers to mobility such as the Hukuo system (Au and Henderson, 2006)?

We focus on three major areas of research: core facts about city size, characterized by Zipf's and Gibrat's Law; the Rosen (1979) and Roback (1982) spatial equilibrium; and the determinants of urban success, including agglomeration economies and human capital effects on wages and on city growth. The transferability of Zipf's and Gibrat's Law is of primarily academic interest. The transferability of the spatial equilibrium framework determines our ability to rely on that framework's many implications, such as the implication that the benefits of new infrastructure for local renters will be muted by higher prices. Economists might want to be far more circumspect about championing human capital and agglomeration if there is little evidence that human capital externalities and agglomeration economies exist in the developing world.

Section 1.2 of this paper describes the data, which can be particularly problematic in the developing world. For the U.S., we will work with Census-defined metropolitan areas using standardized geographic boundaries based on the latest definitions. We tried to duplicate this structure for the other three countries, relying whenever possible on standard Census-like products, but even the definition of metropolitan areas could be difficult. In the case of India, for example, we use districts, but include only the urban population. Our time frame runs from 1980 to 2010.

In Section 1.3 we present the basic facts about the distributions of populations across city sizes. While Zipf's Law is often considered to be a universal truth, like Soo (2014) we do not find it so. Standard statistical tests reject the hypothesis that China, India and the U.S. are characterized by the same power law distribution. Brazil, and most notably China and India, have fewer extremely large sizes than would be predicted by Zipf's Law. Gibrat's Law, which claims that growth rates are independent of initial population levels, holds roughly for the U.S. and Brazil. It does not hold for India and China. In both of these countries, urban population levels show substantial mean reversion from 1980 to 2010. Following the logic of Gabaix (1999), the failure of Gibrat's Law in these countries may explain why Zipf's

Law also fails to hold, perhaps because India and China are still finding their way towards an urban steady state.

Section 1.4 turns to spatial equilibrium, which has long been the organizing principle of urban economics. We do not focus on the intra-urban implications of the spatial equilibrium hypothesis, developed by Alonso (1964), but rather than inter-urban implications developed by Rosen (1979) and Roback (1982). Perhaps the most basic implication of that model is that the advantages of a place, such as particularly good weather, should be offset by countervailing disadvantages, such as long commutes. Higher wages should be offset by either lower amenities or higher housing costs. When there are forces that limit mobility, including place-specific tastes and human capital, moving costs or even legal barriers, then the predictions of the spatial equilibrium model will soften. Results that appear to reject a frictionless spatial equilibrium model may not reject a spatial equilibrium model with severe frictions.

In the U.S., a one-log-point increase in area incomes (estimated as the residual from a regression of earnings on human capital and demographics) is associated with a 1.6 log-points increase in annual rents and a 2.9 log-points increase in housing values. The rents-income relationship is actually too small, relative to the predictions of the Rosen-Roback model, unless higher income areas have low amenities or higher levels of unobserved human capital. The values-income relationship is closer to the predictions of the Rosen-Roback framework.

The comparable elasticities of rents to area earnings for Brazil and China are 1.4 and 1.8 respectively. In China, we also estimate a 1.1 elasticity of housing values to area earnings. As in the U.S., the earnings-rent relationships in these countries are quite strong, but smaller in magnitude than theory would suggest. By contrast, the relationship between earnings and rents in India is practically non-existent. This finding can imply either that Indian rental data is problematic, Indian rental markets are dysfunctional, or that the frictionless spatial equilibrium hypothesis does not hold in India. We suspect that the truth involves some combination of all three explanations.

A second implication of spatial equilibrium is that real wages should be lower in areas with better natural amenities. Within the U.S., real wages rise, primarily because housing costs fall, in areas with less temperate climate. In Brazil, real wages are higher in more temperate areas, primarily because nominal wages are much lower in the hottest areas of the country. We suspect this reflects a combination of omitted human capital differences and imperfect mobility. There is no relationship between climate and real wages in either India or China, perhaps because these countries are not rich enough for ordinary workers to sacrifice earnings for nicer weather.

We also look at income and self-reported happiness across space in the U.S., India and China (data is not available for Brazil). Income and happiness are only weakly related across U.S. cities, which suggests that higher incomes in U.S. metropolitan areas are not generating outsized improvements in personal welfare. Across Chinese and Indian metropolitan areas, the income-happiness relationship appears stronger, even if it is imprecisely measured. A stronger relationship could suggest that differences in unobserved human capital are larger across cities in developing countries than in the U.S., or again, that the spatial equilibrium hypothesis has weaker predictive power in these countries.

The fundamental idea behind spatial equilibrium is that migrants move to equalize welfare levels across space, which seemed distinctly plausible in the highly mobile U.S. Five-year mobility rates in China and Brazil are lower than historic U.S. mobility rates, but the drop in U.S. mobility since 2000 and the rise in Chinese mobility means that the three countries look broadly similar today. India, however, appears to be far less mobile, which may explain why the Indian data does not seem well explained by the frictionless spatial equilibrium model.

The case for a frictionless spatial equilibrium is stronger in the U.S. than in the three developing countries analyzed. Brazil and China do have reasonably high migration rates and a strong correlation between income and housing rents. India has low migration rates and essentially no correlation between income and rents. There is no compensation for less temperate climates in any of the developing countries. We conclude from this subsection

that the spatial equilibrium framework can be used, if it is used warily, in Brazil and China. We see little reason for confidence in the standard framework when applied to India.

Section 1.5 turns to the determinants of local success, such as agglomeration economies and human capital spillovers. As is well known (e.g. Glaeser and Gottlieb, 2009a), there are two standard problems with agglomeration regressions: unobserved personal heterogeneity and unobserved place-based heterogeneity. We address these issues in the limited ways that are standard in the literature (see Combes and Gobillon, 2015, for a discussion), controlling for observable human capital and instrumenting for current population levels with population levels from 1980 and the start of the 20th century.

In the U.S., we estimate an agglomeration coefficient of .054 when the logarithm of male earnings is regressed on metropolitan area population. The coefficient on the logarithm of density is slightly smaller (.046). Our Brazilian estimates are similar to those in the U.S. The elasticity of wages with respect to area population is .052 in Brazil, and the elasticity of wages with respect to area density is .026. In the U.S., we estimate a “real wage” (defined as wages controlling for area rents) elasticity of approximately .02. In Brazil, the elasticity is 0.01, which is not statistically significant.

By contrast, the estimated agglomeration effects are noticeably higher in both China and India, especially with regard to area density. The density elasticity in China is .19. The population elasticity is half the size, which is still higher than the estimated U.S. elasticity, but the Chinese coefficient is not statistically significant. The Indian density elasticity is .076, which is similar to its population elasticity. In India there is also a substantial real-wage premium associated with denser areas and larger urban populations, which again suggests either the unobserved human capital differences are enormous or that India is not characterized by a spatial equilibrium. In China, the real-wage elasticity to density (.052) is comparable to the one in India, but the population elasticity is negative and statistically insignificant.

We then estimate human capital externalities by following Rauch (1993) and Moretti (2004) and regressing the logarithm of earnings on area-level education (measured as the

share of adults with tertiary degrees), individual education and other demographic variables. We acknowledge the significant problem that unobserved human capital may be correlated with measured area-level human capital, but we have no way of solving that problem. Our estimated coefficient for the U.S. is 1.0, suggesting that a ten percent increase in the share of adults with college degrees is associated with an approximately 10 percent increase in earnings.

The comparable coefficients for Brazil, China and India are 4.7, 5.3 and 1.9 respectively. These results suggest that an area-level increase in education in those countries is associated with a far higher increase in the logarithm of earnings than in the U.S. The share of the population with a college degree varies far more across U.S. cities than across developing world cities, so that the impact of a standard deviation increase in area level education is more similar across the four countries.

Finally, we end with the correlation between human capital and the growth of urban populations and income levels. In the U.S., a one percentage point increase in the share of adults with college degrees in 1980 is associated with a 2.2 percentage point increase in population growth between 1980 and 2010 and a .9 percentage point increase in income growth. These results have been taken as evidence for skills-enhancing local productivity growth (Glaeser *et al.*, 1995) or the increasing importance of human capital externalities (Glaeser and Saiz, 2004).

The comparable effects in Brazil are far stronger. A one percentage point increase in the share of adults with a college degree in Brazil is associated with an almost six percentage point increase in population growth from 1980 to 2010 and a twelve percentage point increase in income growth. Again, the impact of skill seems larger in the developing world, although the differences narrow if we consider the impact of a one standard deviation increase in the skills measure.

The impact of skills on population growth in China is even larger. A one percentage point higher increase share of adults with a college degree is associated with a 22 percentage point increase in population growth from 1980 to 2010. The measured impact on income growth

is negative and statistically insignificant. We do not have results on income change in India, but education is a weaker predictor of population growth than in the other developing-world countries. A one percentage point increase in the college educated share in 1980 is associated with only a .34 percentage point increase in population growth over the next thirty years, and the estimate is not statistically significant.

Agglomeration economies and human capital externalities appear robust in the developing world. The correlation between skills and urban growth is extremely strong in China and Brazil. Consequently, two major policy lessons from U.S. data –skills matter for urban success and agglomeration increases productivity– seem to be quite relevant for the developing world.

We conclude that, while there are many similarities between the four countries, there are also important differences. Typically, the Brazilian results are the most similar to those found in the U.S. The Indian results are most different. The results on human capital and agglomeration suggests that these forces are, if anything, more important in the developing world than in the U.S. But anyone who assumes that India is in a frictionless spatial equilibrium is making a leap of faith.

The data cannot reject the possibility that there isn't some kind of spatial equilibrium, with someone on the margin between different locations in India. No data could. But the simplified spatial equilibrium model that has guided much of the empirical work over the past 40 years in the U.S. does poorly when applied to Indian data. One interpretation of these results is that the standard spatial equilibrium framework is of little empirical relevance in poor, traditional economies, where human-capital heterogeneity is enormous and people remain rooted to the communities of their birth. Another interpretation is the property markets are particularly dysfunctional in poorer countries.

As to why the standard spatial equilibrium model fits the richer countries in our sample better than the poorer countries, it is possible that a frictionless spatial equilibrium emerges with development as human capital becomes more widespread and as people turn to markets instead of traditional social arrangements in their home villages. European research

also typically finds results compatible with standard spatial equilibrium assumptions, despite low migration rates (e.g. Buettner and Ebertz, 2009; Hiller and Lerbs, 2015). The transition to a frictionless spatial equilibrium seems like a fertile topic for future research.

Developing world cities contain the lion's share of the world's urban population and it is important that economists understand them better. In 2014, there were more urbanites in Africa and Latin America than in Europe and North America. We also believe that studying these cities will teach us about which urban facts are fundamental and apply generally, and which depend on the level of development. Our tentative conclusion from this paper is that the power of cities and skills to enhance productivity appears common, but the welfare-smoothing function of property markets seems much less ubiquitous.

1.2 Measuring Urban Areas in Four Countries

American urban research often examines variation across metropolitan areas. This research is possible because the United States has a dispersed urban system with a large number of metropolitan areas that have a rich variety of sizes, education levels and incomes. To examine the differences and similarities between the developed and developing world, we chose three countries that are also large, populous and endowed with a dispersed urban hierarchy: Brazil, China and India. These three countries are notable not only for their size, but also for the fact that they are not dominated by a single urban giant, such as Buenos Aires, Jakarta or Mexico City.

While these three countries are frequently grouped together as large emerging markets, they have substantially different income levels. Per capita GDP in India is approximately one-third of per capita income in Brazil, and China lies between these two extremes. Figure 1.1 shows that the paths of urbanization (defined as the percentage of the population living in what each national statistics office calls "urban areas") also differed across these countries. In 1965, Brazil was already one-half urban, while India and China were overwhelmingly rural.

Brazil's high level of urbanization was part of the classic 1960s puzzle of high Latin

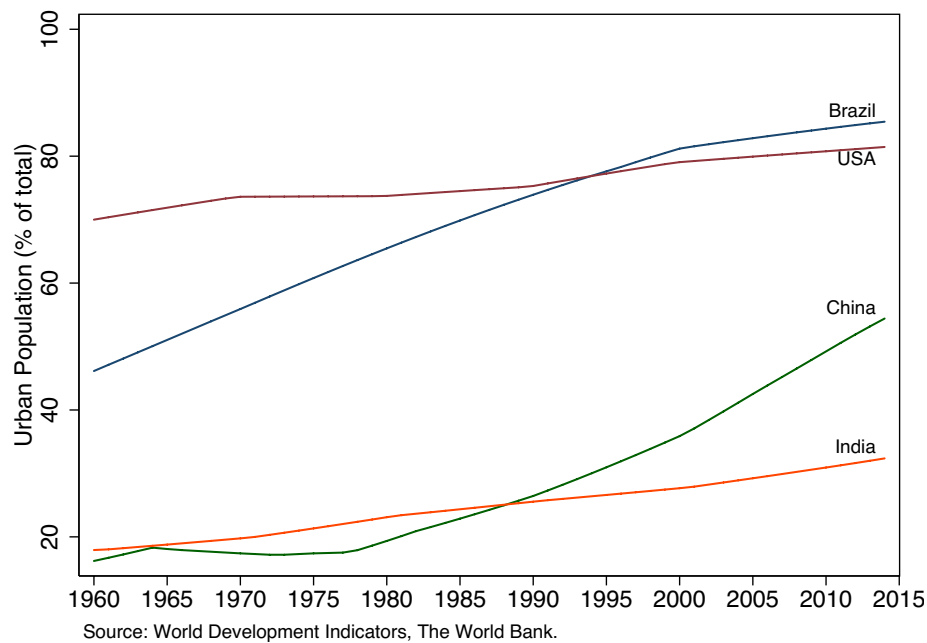


Figure 1.1: *Share of total population living in urban areas, 1960-2014*

American urbanization. Social scientists noted that “Latin America, on the whole, is more urbanized than it is industrialized or developed in other respects” (Durand and Pelaez, 1965), and that “urbanization is occurring without any industrialization” (Arriaga, 1968). While American per capita GDP was \$7,500 (in 2012 dollars) in the 1920s, when the U.S. became 50 percent urban, Brazilian per capita GDP only reached that level in 2011, when it was 80 percent urban. Indeed, today Brazil is more urbanized than the United States despite being far less wealthy.

By contrast, India’s urbanization has shown a slow but steady growth from 18 percent in 1960 to 31 percent in 2010. India is still predominantly poor and predominantly rural. Yet India’s vast size means that it has extensive mega-cities, despite having a low urbanization rate.

Before 1800, China had the globe’s greatest track record of city building, yet despite that history China’s urbanization rate remained below 20 percent when Mao died in 1976. After that point, and the economic opening that came with Deng Xiaoping’s Southern Strategy, China’s urbanization rate exploded. Chinese income and urbanization levels are now far

higher than those in India. China has even more vast cities, most of whom westerners – even western urbanists– cannot name. According to the OECD (2015), in 2010 there were 643 million Chinese living in 127 metropolitan areas with more than 1.5 million people. By contrast, there are only 11 such metropolitan areas all together in the United Kingdom, France, Belgium, the Netherlands, Spain, Portugal and Switzerland (OECD, 2012).

1.2.1 Defining Agglomerations

In order to produce results comparable to U.S. urban research, we need to define comparable geographic units. Even more challenging, we will need to define geographic units that can be identified in large data sets with individual-level information. Typically, U.S. research uses metropolitan areas, which are multi-county agglomerations, defined by the U.S. Census. Since the Census definitions change, we follow the convention of using the definitions as of 2010 and using Consolidated Metropolitan Statistical Areas, which are relatively large groupings. The U.S. Census considers all Americans who live within metropolitan areas to be urban, since they are part of a large urban labor market, even if their own home is surrounded by considerable greenery.

The OECD (2012) and other organizations have already defined functional urban units for a large number of countries, and for many purposes it would be better to simply use their definitions. Yet our purpose is to replicate the U.S. urban literature, which uses individual level-data and, as such, we must also use Brazilian, Chinese and Indian censuses and surveys containing large numbers of individuals. These data sources don't use OECD urban area definitions, and typically contain geographic identifiers based on political boundaries. We will define metropolitan areas using those boundaries, typically excluding non-urban respondents both from our tests and from our definition of area-level variables, such as aggregate population, density and skill levels.

In the case of Brazil we use microregions, which are agglomerations of contiguous and economically integrated municipalities that have similar economic features, defined by the Brazilian Institute for Geography and Statistics (IBGE, 2002). These areas capture better the

notion of local labor markets than municipalities, which are more similar to U.S. counties in that they can differ dramatically in size and economic characteristics. Using legally defined metropolitan regions was not a plausible alternative either, because the Brazilian constitution of 1988 delegates to the states the right to establish them, and the criteria used to form these regions varies significantly across states.

For China we use administrative “cities”, including provincial-level and prefecture-level areas. The name can be misleading, since these geographical units are typically regions that comprise both urban and rural territories. While there is not a single spatial administrative structure for cities, a typical “large” city (provincial or prefecture level) includes both an urban core and large rural areas with scattered towns. The urban core and its surroundings are in turn divided into districts, and the rural areas into “counties” (Chan, 2007).

In the case of India we use districts, the second-level administrative division of the country after states and union territories. This choice enables us to merge the available microdata to area-level aggregates available from the Indian Census and other sources. However, Indian districts are, for the most, geographically extensive areas that contain large numbers of rural dwellers.

In order to make the units of analysis from Brazil, China and India more comparable to those from the U.S., we restrict the samples to urban individuals and urban-only area aggregates throughout most of the study. Moreover, we try to homogenize the sample compositions by restricting them to the higher end of the urban population distributions (areas with 100,000 urban dwellers or more).

Like U.S. definitions of metropolitan areas, the boundaries of Brazilian microregions, Chinese cities and Indian districts change over time. In these large developing countries, the process is mostly driven by the breakup of existing administrative units to create smaller ones. Thus, we need to construct time-consistent geographies for the cases where we perform inter-temporal comparisons.

In the case of Brazil, we employ data from municipality-level border changes for the period 1980-2010 to aggregate microregions when necessary and construct time-consistent

Table 1.1: *Share of people living in urban areas of different size*

	Areas of 100K – 250K (percent)	Areas of 250K – 500K (percent)	Areas of 500K – 1M (percent)	Areas of 1M – 1.5M (percent)	Areas of 1.5M+ (percent)	Population in areas 100K+ (millions)
2000						
USA (MSAs)	5%	8%	8%	6%	38%	184
Brazil (Microregions)	17%	12%	9%	5%	30%	123
China (Cities)	0.3%	1.2%	6%	8%	21%	458
India (Districts)	2%	5%	6%	4%	10%	279
2010 (2011 for India)						
USA (MSAs)	5%	7%	9%	5%	41%	207
Brazil (Microregions)	16%	12%	11%	6%	32%	148
China (Cities)	0.2%	0.8%	4%	6%	39%	669
India (Districts)	2%	4%	6%	4%	14%	373

Note: Population bins are based on the size of the urban population of each area. All figures are expressed as a percent of the total population in the country. **Sources:** See data appendix.

borders following Kovak (2013). In the cases of China and India, we use GIS historical data to geo-match current borders to 1980 borders, and use the 1980 area definitions aggregating smaller areas when necessary.

We recognize that our definitions are debatable, but we believe that these are reasonable choices with the goal of creating a standard definition across quite different countries.

1.2.2 The Distribution of Populations across Area Sizes

Table 1.1 shows the distribution of population across different population sizes in 2000 and 2010. The first five columns show the share of the total national population living in metropolitan groupings of different sizes. The last column shows the total size of the population living in all of these groups put together. All four countries have more than 100 million people living in these areas. In 2010, collectively these urban populations included 1.4 billion people, about one fifth of the world's population.

The U.S. population distribution is heavily skewed towards the larger metropolitan areas, with 38 percent of the population in such areas in 2000 and 41 percent in 2010. Collectively the other four population groupings contain only 26 percent of the U.S. population in 2010.

Brazil also has a large share of its population (32 percent in 2010) in the largest metropolitan areas, but it also has a large share in the smaller areas. Twenty-eight percent of the Brazilian urban population lives in microregions with fewer than 500,000 urban inhabitants. Some of these smaller areas might not even be classified as metropolitan areas within the U.S. We highlight this to emphasize that the data issues make these comparisons challenging, especially when we are dealing with the less populated areas.

By contrast, only one percent of Chinese in 2010 lived in cities of less than 500,000 and 39 percent of Chinese live in metropolitan areas with more than 1.5 million people. While definitional issues might explain some of the absence of smaller Chinese agglomerations, there is no doubt that a large number of Chinese live in extremely large metropolitan areas. Perhaps the most striking fact is that between 2000 and 2010, the share of Chinese living in such areas increased by 18 percent, which reflects both migration and the rapidly expanding populations of many Chinese mega-cities.

Even in India, the share of the population living in the largest urban areas increased significantly between 2000 and 2010 – from ten percent to fourteen percent. India may be the least urbanized country in the group, but it has 373 million urbanites living in cities with more than 100 thousand people, according to our classification. This represents the second largest urban population in the world. The typical urbanite in 2010 is far more likely to reside in Beijing or Shanghai than in London or New York.

Before taking a closer look at these city size distributions, we briefly discuss income heterogeneity in the four countries, both across and within cities. Table 1.2 shows the national income distributions in the four countries and the gulf between urban and rural incomes. Despite the enormous attention given to inequality in the U.S., America is not particularly unequal among these countries. Brazil is the standout in inequality, with both the highest share of its income going to the top ten percent (42 percent) and the lowest share going to the bottom ten percent (one percent); but inequality has dropped noticeably in recent years, particularly between 2000 and 2010.

In India, China and the U.S, between 28.8 and 30 percent of national income goes to the

top ten percent of the income distribution. In China and the U.S., 4.7 percent of income goes to the bottom fifth of the income distribution. The poorest quintile of Indians does much better, as a share of national income, earning 8.6 percent of national income. In all three cases, inequality widened over the 1990-2010 period.

We are particularly interested in the gulf between urban and rural citizens, which is displayed in the bottom panel of the table. In the U.S., urban incomes are 30 percent higher than rural incomes, which is a significant gap, albeit one that is offset by higher urban costs of living. In China, urban incomes are 44 percent higher than rural incomes, which is significant but not extreme.

By contrast, urban Indians in our sample earn 122 percent more than rural Indians. Urban Brazilians earn 176 percent more than the rural Brazilians. These gulfs are enormous, suggestive of huge productivity differences between urban and rural areas. Presumably, a significant fraction of these gulfs reflect unobserved and observed human capital characteristics (Young, 2013), and perhaps also non-pecuniary compensation in rural areas. Given the enormous differences between rural and urban Brazil and India, we will include only urbanites in the tests that follow. Nonetheless, we will still have to grapple with unusually large earnings differences across space that do not seem to be fully offset by differences in housing costs.

1.3 City Size Distributions: Zipf's Law and Gibrat's Law

Before discussing facts related to economic theories, we follow Rosen and Resnick (1980) and Soo (2005), and turn to two stylized facts about city size distributions: Zipf's Law and Gibrat's Law. We choose to begin here despite the large international literature on these laws (e.g. Rose, 2006; Giesen and Sudekum, 2011; Soo, 2014), because we are using slightly different city size definitions and because it is important to duplicate past results using our attempt at producing consistent data. Zipf's Law was originally posed as the rank size rule: the population of the N th largest city is $1/N$ times the population of the largest city. In large

Table 1.2: Income distributions, 1990-2010

	USA			Brazil		
	1990	2000	2010	1991	2000	2010
National income distribution						
Income share held by lowest 20%	5.4%	5.4%	4.7%	2.3%	2.4%	3.3%
Income share held by second 20%	11.2%	10.7%	10.4%	5.5%	5.9%	7.5%
Income share held by third 20%	16.7%	15.7%	15.8%	9.7%	10.3%	12.3%
Income share held by fourth 20%	23.7%	22.4%	23.1%	17.9%	18.0%	19.4%
Income share held by highest 20%	43.1%	45.9%	46.0%	64.6%	63.4%	57.6%
Income share held by lowest 10%	1.8%	1.8%	1.4%	0.8%	0.7%	1.0%
Income share held by highest 10%	26.7%	29.9%	29.6%	48.1%	47.3%	41.9%
Income levels (in 2014 \$, PPP)						
	Income per capita			Income per capita		
	1990	2000	2010	1991	2000	2010
Urban Areas	37,195	44,071	45,124	3,899	5,830	7,543
Rural areas	26,816	31,342	34,835	1,141	1,846	2,731
Difference	10,379	12,729	10,289	2,758	3,984	4,812
	China			India		
	1990	1999	2010	1993	2004	2009
National income distribution						
Income share held by lowest 20%	8.0%	6.4%	4.7%	9.1%	8.6%	8.6%
Income share held by second 20%	12.2%	10.3%	9.7%	12.8%	12.2%	12.1%
Income share held by third 20%	16.5%	15.0%	15.3%	16.5%	15.8%	15.7%
Income share held by fourth 20%	22.6%	22.2%	23.2%	21.5%	21.0%	20.8%
Income share held by highest 20%	40.7%	46.1%	47.1%	40.1%	42.4%	42.8%
Income share held by lowest 10%	3.5%	2.7%	1.7%	4.0%	3.8%	3.7%
Income share held by highest 10%	25.3%	29.7%	30.0%	26.0%	28.2%	28.8%
Income levels (in 2014 \$, PPP)						
	Income per capita			Earnings per capita		
	1990	1999	2010	2005	2011	
Urban Areas	1,451	3,064	6,179	3,123	5,027	
Rural areas	660	1,099	4,265	1,129	2,265	
Difference	792	1,965	1,914	1,994	2,762	

Sources: See data appendix.

samples, this claim is equivalent to the city size distribution being characterized by a power law distribution with a coefficient of minus one.

Gibrat's Law is dynamic. It states that the growth rate of population is unrelated to the initial population. Researchers typically test Gibrat's Law by regressing the change in the logarithm of city population on the initial level of city population, and testing whether the coefficient is statistically distinct from zero. Champernowne (1953) and Gabaix (1999) linked the two facts and showed that if Gibrat's Law holds for city growth rates, then the equilibrium distribution of city sizes will display Zipf's Law. This result is mathematical, not economic, and it requires no assumptions about the motives of migrants or the productivity of firms.

Our purpose is not to revisit the many controversies around Zipf's Law (e.g. Holmes and Lee, 2010) or the methodological issues related to measurement. We will use the simplest established techniques and compare across countries. We will also use our measures of metropolitan area population, which are also debatable. Throughout this paper, we aim to reproduce simple, transparent facts about space, not to advance methods or debate nuanced issues within the established literature.

1.3.1 Zipf's Law across Four Countries

Figure 1.2 shows the cumulative distribution of city sizes for our four countries. In all cases, we consider only those areas with more than 100,000 urban inhabitants. The plots show the relationship between the logarithm of the city rank and the logarithm of the area population.

We follow Gabaix and Ibragimov (2011) who present theory and simulations showing that the relationship between the logarithm of area population and the log of rank minus one-half is a far better estimate of the coefficient on the power law distribution for area sizes than the relationship between the log of area population and the log of the rank in the city population. We also plot the fitted line that is implied by this procedure.

The results for the U.S. show a coefficient of -.91, which is lower in magnitude than the -1.05 estimate discussed by Gabaix and Ibragimov (2011). We differ from them primarily

because we consider a broader range of metropolitan areas. The figure shows how the size-rank relationship steepens at higher ranks. If we restricted our sample to the 135 largest metropolitan areas, as they do, our estimate would be larger in magnitude and closer to their estimate.

The next figure shows the results for Brazil. The fit of the relationship is extremely tight. Brazil's city populations do seem well characterized by a power law distribution, with less non-linearity than in the U.S. However, the -1.18 estimated coefficient is much higher than in the U.S. and higher than predicted by Zipf's Law. This high coefficient means that population rises too slowly as rank falls, or that Brazil's biggest cities are smaller than Zipf's Law would predict. Soo (2014) finds an estimate of -.94 for Brazil across his entire sample, but the coefficient rises as he restricts the sample to larger cities. Rose (2006) found a coefficient of -1.23 for Brazil which is quite close to our estimate.

The third figure shows results for China, following Anderson and Ge (2005). The estimated coefficient of -.91 seems reassuringly close to the U.S., but the figure suggests that such comfort is mistaken. The -.91 coefficient masks strong non-linearity in the rank-size relationship, and the r-squared is quite low (.79) relative to the U.S. (.94) or Brazil (.99). The steep curve among the larger Chinese cities suggests that when it comes to big areas, China is more like Brazil than like the U.S. China also has far fewer extremely large cities than Zipf's Law would suggest. The -.91 estimate is larger in magnitude than Soo (2014), but smaller than Schaffar and Dimou (2012) and Rose (2006).

The final panel shows results for India, which has -1.03 coefficient, suggesting that Zipf's Law appears to work for Indian cities even more strongly than it does in the U.S. This estimate is close to Soo (2005) but lower than Rose (2006). Our r-squared (.92) is somewhat lower than the U.S. and there is also some concavity in the relationship between rank and size, suggesting that there are again fewer large cities than Zipf's Law would suggest.

In Table 1.3, we test whether the city size distributions are the same across the four countries. The top panel in the table shows results for all city sizes. The bottom panel shows results only for city sizes above 500,000, where we are more confident that metropolitan area definitions are not driving the results.

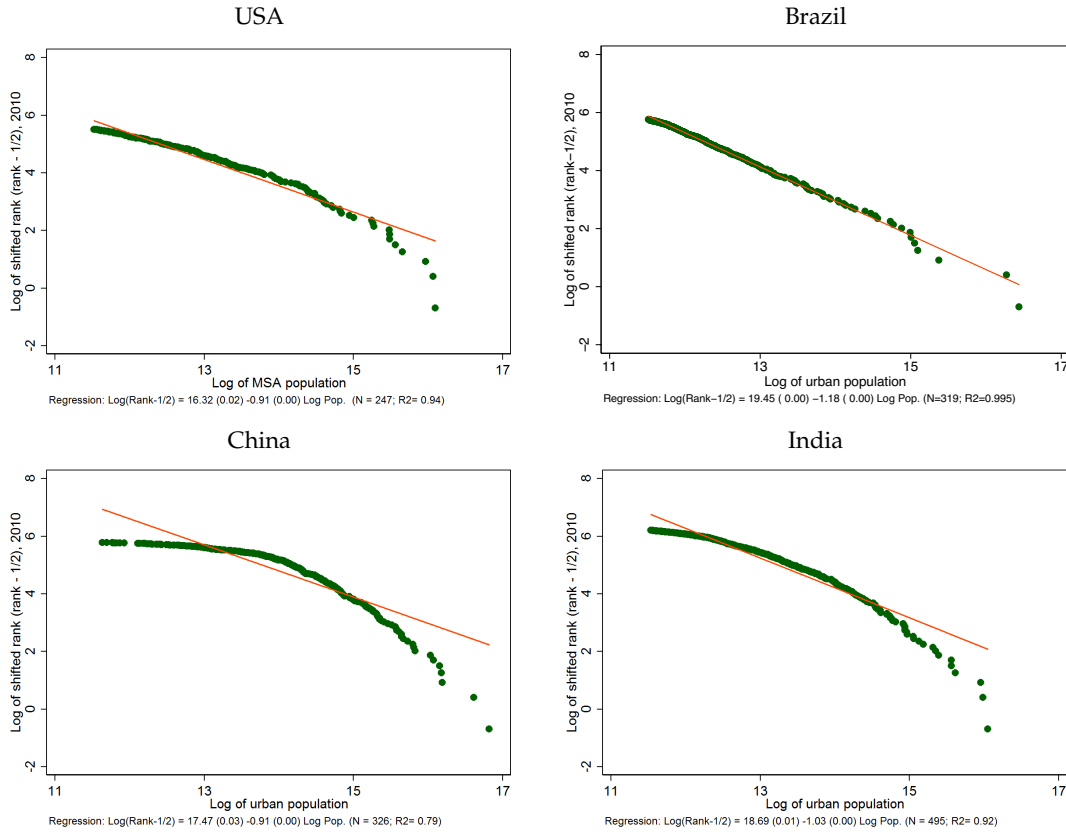


Figure 1.2: Zipf's Law. Urban populations and urban population ranks, 2010

Note: Regression specifications and standard errors based on Gabaix and Ibragimov (2011). Samples restricted to areas with urban population of 100,000 or larger. **Sources:** See data appendix.

The table reports D statistics from two-sample Kolmogorov-Smirnov tests that compare each pair of city size distributions. For the whole sample of cities, for every pair of countries we can reject that the distributions are identical. China's city size distribution is particularly distinct statistically, with all D statistics above .5. India and the U.S. appear to have the most similar city size distributions, with a D statistics less than .2.

When we turn only to the larger metropolitan areas, the differences become muted. The U.S., Brazilian and Indian city size distributions are no longer statistically distinct. China's city size distribution is, however, statistically different from the other three. The primary difference is again that China has fewer ultra-large cities than the U.S. city size distribution would predict if it was applied to the number and total population of Chinese cities.

Table 1.3: *Urban population Kolmogorov–Smirnov two-sample tests*

	Brazil (Microregions)	China (Cities)	India (Districts)
Full Sample			
USA (MSAs)	0.396 (0.000)	0.534 (0.000)	0.194 (0.000)
Brazil (Microregions)		0.779 (0.000)	0.346 (0.000)
China (Cities)			0.564 (0.000)
Cities with urban population of 500,000 or more			
USA (MSAs)	0.148 (0.432)	0.229 (0.001)	0.123 (0.286)
Brazil (Microregions)		0.342 (0.000)	0.085 (0.911)
China (Cities)			0.301 (0.000)

Note: Figures are D test-statistic scores, p-values in parentheses. The observations in the full sample are: US = 258, Brazil = 548, China = 345 and India = 632. The observations in the restricted sample are: US = 93, Brazil = 55, China = 296 and India = 204. **Sources:** See data appendix.

There are many possible explanations for these differences. China's population has exploded so rapidly that it may be far from steady state. China's governments are far more active in planning city populations than any of the other countries. The growth of ultra large Chinese cities may also be blocked by disamenities of size that can become extreme for urban populations over 20 million. Finally, both China and India may be better seen as continents rather than standard countries and this may also explain some of the difference.

The differences between China and the other countries do raise the possibility that in the long run China's urban populations will be much more skewed towards ultra large areas like Beijing and Shanghai. The attempts of many local governments to boost growth in middle size (Tier 3 and Tier 4) cities seem to have led to fiscal difficulties. Over time, more vertical construction and congestion pricing may ease the disamenities of crowding and congestion. China's city size distribution may eventually look far more like Zipf's Law, and to examine that possibility we now turn to the dynamics of city growth and Gibrat's Law.

1.3.2 Gibrat's Law across Four Countries

Table 1.4 shows our results on the mean reversion of city populations. In all cases, we report coefficients where the change in the logarithm of area urban population is regressed on the logarithm of initial area urban population. Gibrat's Law implies that the coefficient should be statistically indistinguishable from zero. While Gibrat's Law does hold for the U.S. in recent decades, it does not appear to hold well in Germany (Bosker *et al.*, 2008) or for the U.S. historically (Glaeser *et al.*, 2014b).

The first column shows results for the United States for 1980-2010. We first show the coefficient for the entire time period and then results for each of the three decades separately. Over the entire 1980-2010 period, there is no correlation between initial population and subsequent population growth. The r-squared in the regression is 0 to two decimal places. The estimated coefficient is .009. The standard error around that coefficient means that we cannot rule out the possibility that the coefficient is .03, but even that coefficient is quite small for a thirty-year period.

Gibrat's Law holds less perfectly within the U.S. during each independent decade. In both the 1990s and the 2000s, the estimated coefficient is close to .01 and statistically significant. Yet given the strong correlation that exists between metropolitan area growth and other variables, such as January temperature and education, these coefficients are quite compatible with Gibrat's Law.

Gibrat's Law has failed during many periods of U.S. history. Glaeser *et al.* (2014b) examine population growth among eastern counties of the U.S. from 1860 until today. For example, during the 1970s, there was sharp mean reversion in population levels in those counties. During the 1960s, population growth was much faster in more populous counties. Gibrat's Law has not been a permanent feature of U.S. urban dynamics, and perhaps it should not be expected to hold in countries experiencing far more rapid urban change.

The second column shows the results for Brazil, which are generally statistically indistinguishable from the U.S. Over the entire time period, like Soo (2014), we cannot reject the hypothesis that Gibrat's law holds in Brazil. During the 1980s, there was slight mean

Table 1.4: *Gibrat's Law: Urban population growth and initial urban population*

	USA (MSAs)	Brazil (Microregions)	China (Cities)	India (Districts)
1980 - 2010	0.009 (0.020) N=217 R2=0.001	-0.038 (0.023) N = 144 R2 = 0.015	-0.447*** (0.053) N=187 R2=0.280	-0.052** (0.023) N=237 R2=0.021
1980 - 1990	0.008 (0.008) N=217 R2=0.004	-0.026** (0.013) N = 144 R2 = 0.020	-0.310*** (0.054) N=187 R2=0.151	0.063* (0.034) N=237 R2=0.015
1990 - 2000	0.014** (0.007) N=217 R2=0.019	0.001 (0.010) N = 144 R2 = 0.000	-0.308*** (0.036) N=187 R2=0.280	0.005 (0.020) N=237 R2=0.00
2000 - 2010	0.012** (0.006) N=217 R2=0.018	0.006 (0.006) N = 144 R2 = 0.006	0.019 (0.021) N=187 R2=0.005	-0.013 (0.015) N=237 R2=0.004

Note: All figures reported correspond to area-level regressions of the log change in urban population on the log of initial urban populations in the specified period. Regression restricted to areas with urban population of 100,000 or more in 1980. Robust standard errors in parentheses.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

reversion, but during the 1990s and 2000s, Gibrat's Law seems to describe the data well. These results also echo Resende (2004).

China's results are shown in the third column. There is strong mean reversion over the entire time period and during individual decades, except for the 2000s. As China liberalized and migration increased, smaller and middle-sized cities grew faster than the most populous. These patterns don't look at all like Gibrat's Law, which is perhaps why Zipf's Law also seems to fail for China.

The fourth column shows the coefficients for India. Over the entire time period, the coefficient is significantly negative. If a city's population was 1 log point higher in 1980, then it grew on average by .052 log points less over the next 30 years. This negative coefficient does not imply that India has once great cities that are declining, but rather that growth

was particularly robust in smaller agglomerations.

When we split the Indian growth by decades, we see that the 1980s were marked by positive serial correlation, where higher populations led to faster growth, while this trend disappeared in the 1990s and the 2000s. One possible explanation for this shift is that prior to the economic liberalization in the early 1990s, regulation tended to keep the urban hierarchy in places.

Brazil and the U.S. both appear to adhere broadly to both Zipf's and Gibrat's Laws. China and India do not. Perhaps the most natural reason why Brazil and the U.S. are similar is that they are both moderately sized places, which have long been largely urban. China and India are both much larger, and many of their cities are much newer. If they have not reached a dynamic steady state, then perhaps Gibrat's and Zipf's Laws may eventually appear in their urban systems.

1.4 Spatial Equilibrium

We now turn to empirical tests that are motivated by economics. For over fifty years, the spatial equilibrium hypothesis has been the organizing principle of urban economics. It was first applied to land prices and land usages within metropolitan areas by Alonso (1964) and Muth (1969) and then it was applied to income and price differences across metropolitan areas by Rosen (1979) and Roback (1982). The core idea of spatial equilibrium is that locations don't offer a free lunch. If a place has high wages and decent amenities, then real estate costs should be high. If a place has nice amenities, then real wages (i.e. wages controlling for local prices) should be lower.

We look at four different empirical patterns that are related to the spatial equilibrium hypothesis. We begin by testing whether the costs of living rise with wages across metropolitan areas. We then test whether real wages are lower in places that have more attractive natural climates. Third, we examine whether self-reported life satisfaction is higher in places with higher income and we end this section by looking at overall migration patterns. Migration is not itself a prediction of the spatial equilibrium model, but it is one channel through which

spatial equilibrium is produced. When migration is low, we might be less confident about the predictions of the standard spatial equilibrium model.

The first three tests all try to assess whether people in one area are receiving a higher welfare level than people in another area. But if these tests fail, then there are always two quite plausible explanations. First, a spatial equilibrium might not exist because of legal or preference-based barriers to mobility. Second, the people in the more successful area might have fundamentally different levels of unobserved human capital than the people in the less successful area. If the two areas have very different people, then we would not expect them to deliver the same welfare levels. While we can control for observable human capital measures, such as years of schooling, we can never reject the possibility that unobserved human capital is driving our results.

1.4.1 The Relationship between Prices and Wages

The starting point for spatial equilibrium is the assumption that utility is equalized over space for any homogenous set of workers who are living in multiple cities. Individual heterogeneity can come in the form of place-independent heterogeneity, such as different levels of human capital or tastes for particular amenities, or place-dependent heterogeneity, such as taste for living in a particular locale. Both types of heterogeneity can be modeled. For example, Glaeser (2008) discusses models of heterogeneous worker human capital. Diamond (2016) works with heterogeneous tastes for cities, as well as heterogeneous human capital. For expositional purposes, we will stick with the most standard and simple assumption of worker homogeneity.

In this case, we can define an indirect utility function over wages, prices and amenities $V((1-t)W, P, A)$, where $(1-t)W$ reflects after-tax wages, P reflects prices and A reflects amenities. This indirect utility function is typically either operationalized as a log-separable or a linear-separable function. The log-separable function is justified by a Cobb-Douglas utility function defined over a general consumption good and housing. This can produce an indirect utility function of $A(1-t)WP^{-\alpha}$, where A represents an index of amenity values

which is assumed to multiply welfare and α represents the share of housing in the utility function and household spending.

The linear-separable structure is justified by assuming that every person consumes exactly one unit of housing and, consequently, people's after-housing income is $W - P$. In the linear separable formulation, it is convenient to assume that the amenity index is just added to net earnings so that total welfare is just $(1 - t)W - P + A$. In the log-separable formulation, nation-wide proportional taxes are irrelevant to the relationship between wages and prices. In the linear-separable formulation, nation-wide proportional taxes will matter, unless housing costs are deductible. In the U.S., housing prices are partially deductible because of the home mortgage interest deduction.²

The log-separable formulation suggests the relationship:

$$\text{Log}(\text{Price}) = \frac{1}{\alpha} (\text{Log}(\text{Wage}) + \text{Log}(\text{Amenities})) \quad (1.1)$$

The linear-separable formulation suggests:

$$\text{Price} = \text{After Tax Wage} + \text{Amenities} \quad (1.2)$$

If log price is regressed on log wages, then the first formulation implies that the coefficient will be $\frac{1}{\alpha} \left(1 + \frac{\text{Cov}(\text{Log}(\text{Wage}), \text{Log}(\text{Amenities}))}{\text{Var}(\text{Log}(\text{Wage}))} \right)$. As the historic share of spending that goes towards housing is approximately one-third, this suggests a benchmark coefficient of three. The second formulation suggests that if price is regressed on wage, the coefficient will be $1 - t + \frac{\text{Cov}(\text{Wage}, \text{Amenities})}{\text{Var}(\text{Wage})}$. The two models yield tight predictions only if we know the correlation of the amenity index and wages, which we unfortunately do not. Our approach is not to attempt to definitively disprove the spatial equilibrium predictions, but rather to test whether reality is roughly compatible with the predictions of the model in our four countries.

An added complication is that measured wages and measured housing rents and prices will necessarily vary because of differences in human capital and the physical characteristics

²Albouy (2015) provides a comprehensive discussion of the connection between deductibility and spatial equilibrium.

Table 1.5: *Regressions of local prices on wages, 2010*

	USA (MSAs)	Brazil (Microregions)	China (Cities)	India (Districts)	USA (MSAs)	China (Cities)
	Log of rents				Log of prices	
Average log wage	1.225*** (0.106) N = 29M R2=0.208	1.011*** (0.044) N = 819 K R2=0.560	0.853*** (0.157) N=6.5 K R2 =0.187	-0.044 (0.052) N=1,484 R2=0.304	1.922*** (0.172) N = 56M R2 =0.396	1.122 *** (0.073) N = 24.5K R2 = 0.521
Average log wage residual	1.612*** (0.159) N = 29M R2 = 0.202	1.367*** (0.076) N = 819 K R2 = 0.552	1.810*** (0.167) N=6.5 K R2 =0.311	-0.019 (0.060) N=1,484 R2=0.304	2.887*** (0.256) N = 56M R2 = 0.403	1.097 *** (0.122) N = 24.8K R2 = 0.515
Dwelling characteristics	Yes	Yes	Yes	Yes	Yes	Yes

Note: Regressions at the urban household level, restricted to areas with urban population of 100,000 or more. All regressions include a constant. Standard errors clustered at the area level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

of the house. Our approach to this issue is to estimate a wage residual from a regression in which the logarithm of wages is regressed on individual human capital characteristics, including years of schooling and age. To promote comparability, we will only include males in this wage regression. Excluding women from the wage regression seems particularly necessary because female labor force participation is much lower in India than in the other countries, but we recognize that there might be further differences between countries if women were included. We then include this residual in a regression in which housing rent or price is regressed on this area-level wage and other housing characteristics, especially the physical characteristics of the home.

We begin with the United States. Table 1.5 shows the coefficient when the logarithm of housing rents or of housing prices (at the household level) is regressed on two measures of area-level income. The first row shows results when we define income as the average of the logarithm of income in the area. The second row instead uses the average of the residual from a regression of the logarithm of wages on human capital characteristics. For the logarithm of housing rents as dependent variable, the first coefficient is 1.23 and the second coefficient is 1.61. For the logarithm of housing prices as dependent variable the

coefficients are 1.92 and 2.89, respectively.

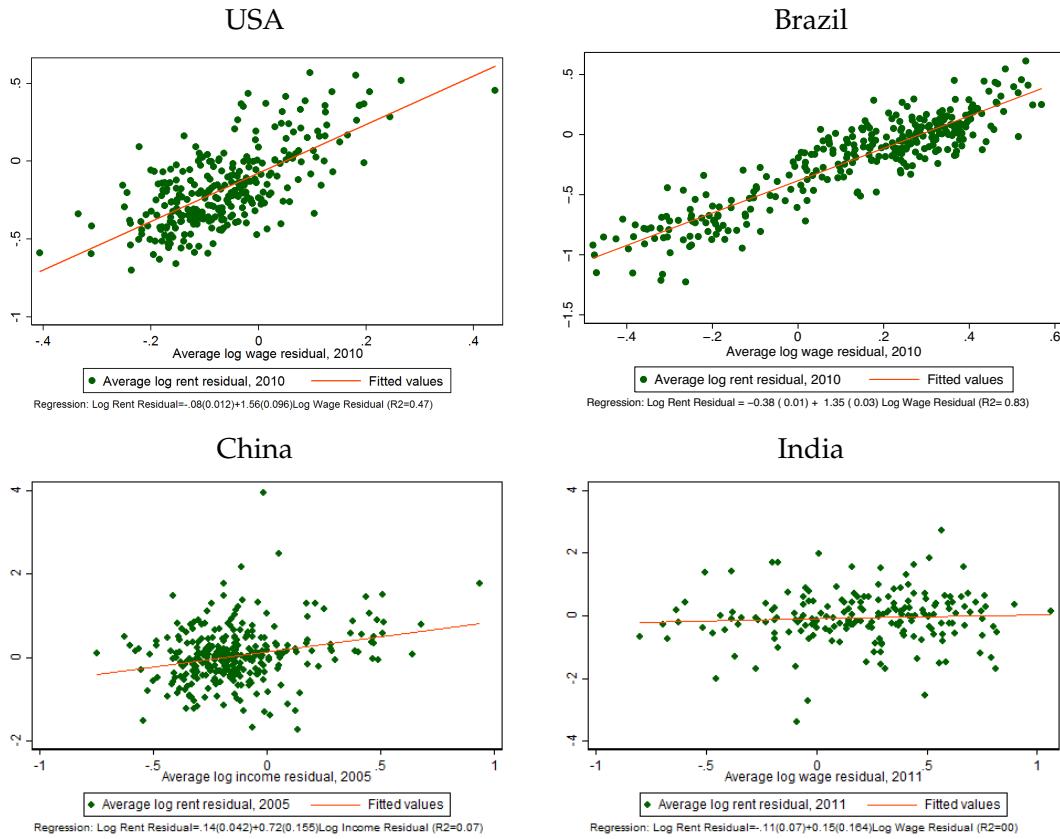


Figure 1.3: *Income and rents, 2010*

Note: Samples restricted to areas with urban population of 100,000 or more. **Sources:** See data appendix.

Figure 1.3 shows the core relationship visually at the area level using housing rents, and Appendix Figure A.1 does it for housing prices in the countries where we have this data. For the U.S. the plot shows the metropolitan area log wage residual and the metropolitan area log rent residual. At the metropolitan area level, the r-squared is .47, but the coefficient seems too small. Given that Americans spend, $1/3$ of their incomes on housing, the predicted coefficient should be three, unless urban amenities move with housing costs. When we rerun the regression in levels, we estimate a coefficient of .13, which is certainly much lower than the value of one minus the tax rate, which is predicted by theory.

There are several possible explanations for finding a coefficient below that suggested by the Rosen-Roback model. Most obviously, amenities may be negatively associated with

wages in the U.S., and there is some evidence to support that view. The share of workers with commute times over 20 minutes is significantly higher in metropolitan areas with higher incomes. January temperatures are lower in areas with higher incomes.

A second hypothesis is that the independent variable is mismeasured badly, which will naturally lead to attenuation bias. Many renters receive public assistance or are in public housing. Consequently, their rents may be artificially low. Building quality levels may differ systematically across areas.

A third view is that since the majority of Americans are owners, and since rental apartments tend to be lower quality, we are not capturing the true cost of living in a particular place. Column 5 duplicates these results with self-reported housing values from the Census, assuming that ownership costs (including finance, depreciation and maintenance) are approximately ten percent of housing values. We find that the logarithmic specification yields a coefficient much closer to one than to three when regressed on the average log wage, but much closer to three (2.89) when regressed on the average wage residual. The levels coefficient is also small, although substantially larger than the rent coefficient. Housing values are also an imperfect measure of housing costs because they are partially shaped by expectations of future housing appreciation, and that expected appreciation lowers the effective price of housing.

The second column of Table 1.5 and the second graph in Figure 1.3 show the basic results for Brazil with housing rents as dependent variable. We were not able to obtain individual-level housing value data for this country. The estimated coefficients range from 1.01 to 1.37. The microregion level r-squared is comparable to the U.S. metropolitan area sample. These results corroborate Azzoni and Servo (2002) who also find higher costs of living in higher-wage Brazilian regions.

The Brazilian figure should be larger than the U.S. figure because Brazilian spending on housing is a smaller share of total income, approximately 15 percent according the Credit Suisse Emerging Consumer Survey. If that is correct, then the predicted log coefficient could be as high as seven. The same explanations for the low estimate exist in Brazil as well as the

U.S.: a negative amenity correlation with high incomes and mismeasurement of both the dependent and independent variable.

Overall, though, the Rosen-Roback inspired wage-to-rent relationship looks pretty similar in the U.S. and Brazil. In both cases, area-level rents are tightly correlated with area-level incomes. In both cases, the coefficients are close to one, which is a far smaller relationship than is predicted by economic theory. The similarities between the Brazilian and American results leave us optimistic that the tradition of Rosen-Roback inspired hedonic regressions that has been so successful with U.S. data can proceed in Brazil. However, we must stress that the similarity is also somewhat strange given that neither of the rent-income coefficients is close to the parameter predicted by the model unless there is a negative correlation between income and amenities. We hope that future researchers who work on the similarity between the countries will help resolve why results in both countries seem to miss the model by roughly the same amount.

The third column of Table 1.5 and the third graph in Figure 1.3 provide the results for China using housing rents as dependent variable, and the sixth column of Table 1.5 and second graph in Appendix Figure A.1 do the same using housing values. Most of the table's estimated coefficients are again close to one, which suggests that Chinese do pay more when incomes are higher, as in Long *et al.* (2009). Nonetheless, the coefficient of one seems low, since the Chinese spend an even smaller share of their incomes (ten percent) on housing than the Brazilians. The graph shows that the r-squared of the relationship between income and rents (.07) is much smaller than in the U.S. and Brazil. The r-squared of the relationship between income and housing values in Figure A.1 (0.34) is larger, but still not as large as in the U.S. The goodness of fit in the table and the figure can be quite different, because the table reflects individual-level data while the figure looks at the area-level relationship, which is not weighted by the number of people in each area.

Chinese rents are correlated with incomes across areas, but the link is much weaker than either the U.S. or Brazil. One explanation for this is that amenity correlation with income is even more negative and that is certainly possible. Another possibility is the barriers to

mobility in China, especially the famous Hukou system make it difficult for migration to equalize welfare levels.

Yet a third possibility is that public interventions in the housing market are particularly important in China, and these act to distort market prices. Moreover, only 10 percent of Chinese and 13.4 percent of Indians rent. A standard explanation for these low rates is that rental markets are dysfunctional and distorted by rent-control-like rules. It is notable that speculators in Chinese real estate will often buy apartments and leave them empty rather than taking the risk of renting them out. Consequently, people are unable to rent and must buy low quality housing instead. The rental market that does exist may only reflect a very unusual part of the housing market. The fact that we obtain larger coefficients when using housing values as the dependent variable (although still significantly smaller than in the U.S.) is consistent with this view.

Finally, with those concerns in mind, we turn to the results from India. In this case too, we only have data for housing rents. The last graph in Figure 1.3 shows that there is essentially no correlation between urban income and rents across Indian districts. This non-result is repeated in the fourth column in Table 1.5. Possibly, this non-correlation reflects profound amenity differences across Indian districts, but it seems just as likely to reflect terrible measurement of housing rents. For example, in many cases, landlords are not allowed to raise rents and cannot eject tenants. Indeed, it is hard to see any pattern in our Indian rent data, which leads us to suspect that without better data on the cost of living, any hedonic estimates pursued in this context are risky. Certainly, we cannot conclude that this provides any support for the usefulness of the standard spatial equilibrium model in India.

Across the four countries, the patterns in Brazil and the U.S. were broadly similar. Both countries have a tight correlation between income and rents. In both countries, the estimated relationship was smaller than would be predicted by the core Rosen-Roback model. The link between income and rents was weaker in China, but still significant. The link disappears entirely in India. For some reason, the standard spatial equilibrium model appears to be least effective in the poorest nation, either because mobility frictions are large, because an

equilibrium does not exist, or because measurement is most problematic. We now turn to the equilibrium pricing of amenities.

1.4.2 Real Wages and Amenities

Equations 1.1 and 1.2 also provide implications about the connection between amenities and real wages, where real wages are defined either as $\text{Log}(Wage) - \alpha \text{Log}(Price)$, in the logarithmic model or as $Wage - Price$, in the linear model. When it comes to amenities, the models do not yield any implication about the magnitude of the effect. That will depend upon consumer valuation of amenities. The model does, however, imply that areas with positive amenities will have lower real wages.

We will focus on climate-related amenities, which have the advantage of being exogenous to the local economy, and generally well measured. Climate appears to be valued as an amenity in wealthy nations including the U.S. (Glaeser and Gottlieb, 2009a), France (Cavailhes *et al.*, 2009) and Germany (Rehdanz and Maddison, 2009). In our samples, we will typically have a measure of Winter and Summer temperatures (corresponding to January and July in the U.S.) and an average precipitation measure. For January and July temperatures, we will transform the variables by taking the absolute value of the difference between the average temperature and the equivalent in Celsius of 70 degrees Fahrenheit (21.11 degrees Celsius). Consequently, the value can increase if the area is either particularly hot or particularly cold. The choice of 70 degrees represents a middle ground within the 65 to 75 degree range that is often discussed an ideal for human comfort. We recognize that this choice is relatively arbitrary within this range. Our results are not particularly sensitive to minor perturbations in the assumed bliss point for temperature. For average annual rainfall, we will use standardized values.

Table 1.6 shows our results. The first panel shows the findings for wages, rents and real wages in the U.S. regressed on our three distinct measures of climate amenities. The first column shows that there is no relationship between nominal wages and these variables, controlling for our core human capital attributes (age, race and years of schooling). The

Table 1.6: *Climate amenities regressions, 2010*

	USA (MSAs)				Brazil (Microregions)		
	Log wage	Log real wage	Log rent	Log price	Log wage	Log real wage	Log rent
Absolute difference from ideal temperature in the summer	0.001 (0.003)	0.006*** (0.002)	-0.027*** (0.008)	-0.066*** (0.015)	-0.084*** (0.007)	-0.045*** (0.003)	-0.110*** (0.016)
Absolute difference from ideal temperature in the winter	0.002 (0.002)	0.005*** (0.001)	-0.018*** (0.003)	-0.032*** (0.007)	-0.015** (0.006)	-0.005 (0.004)	-0.012 (0.010)
Average annual rainfall (std. deviations from the mean)	-0.006 (0.008)	0.005 (0.007)	-0.054** (0.026)	-0.129*** (0.033)	0.063*** (0.015)	0.010 (0.009)	0.179*** (0.028)
Education groups controls	Y	Y	N	N	Y	Y	N
Age groups controls	Y	Y	N	N	Y	Y	N
Dwelling characteristics controls	N	N	Y	Y	N	N	Y
Observations (thousands)	28,237	8,497	24,125	44,765	2,157	2,157	2,157
Adjusted R-squared	0.249	0.158	0.117	0.372	0.341	0.315	0.477
	China (Cities)				India (Districts)		
	Log wage	Log real wage	Log rent	Log price	Log wage	Log real wage	Log rent
Absolute difference from ideal temperature in the summer	-0.005 (0.018)	-0.006 (0.015)	-0.001 (0.021)	0.000 (0.037)	0.000 (0.004)	-0.004 (0.006)	0.001 (0.001)
Absolute difference from ideal temperature in the winter	0.003 (0.009)	-0.004 (0.009)	0.019** (0.009)	0.035* (0.018)	-0.001 (0.003)	0.003 (0.004)	0.000 (0.001)
Average annual rainfall (std. deviations from the mean)	0.109 (0.067)	0.021 (0.055)	0.256*** (0.069)	0.164 (0.142)	0.063** (0.025)	0.049* (0.036)	- 0.005 (0.013)
Education groups controls	Y	Y	N	N	Y	Y	N
Age groups controls	Y	Y	N	N	Y	Y	N
Dwelling characteristics controls	N	N	Y	Y	N	N	Y
Observations (thousands)	5.8	4.2	3.4	6.1	8.4	1.8	2.9
Adjusted R-squared	0.145	0.118	0.079	0.070	0.235	0.228	0.762

Note: Regressions at the individual level, restricted to urban prime-age males or urban households in areas with urban population of 100,000 or more. All regressions include a constant. Standard errors clustered at the area level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

second column shows that a ten degree Celsius difference from ideal temperatures in either the winter or the summer is associated with an approximately five percent increase in real wages. Americans do seem to be paid higher real wages when they live in less temperate climates. The third and fourth columns show that these differences are largely driven by lower rents or lower housing values in less temperate areas. In Europe as well, positive amenities are more likely to drive up housing costs than drive down wages (Buettner and Ebertz, 2009). Although the effects of rainfall on real wages are not statistically significant, rainier places also tend to have lower rents and housing values.

These climate relationships with rents and real wages are a prediction of the Rosen-Roback model that is confirmed within the U.S. Do real wages rise with bad climates in the developing world? The second panel in Table 1.6 shows the results for Brazil. Brazil has lower nominal wages in places that have less ideal temperatures. Wages are higher in areas with more rain as well.

These differences are driven primarily by the huge gaps in the level of development between northern and southern Brazil. There are large human capital differences between north and south which are assuredly only imperfectly captured by our coarse control variables and which are correlated with the weather (Azzoni *et al.*, 2000). There are also differences in the level of capital stock and infrastructure. Other work (Mueller, 2005) finds that Brazilians do seem to value climate differences that are not correlated with region.

Finally, the third regression shows results with rents, which are indeed also lower in places with more extreme weather. While these results are compatible with the Rosen-Roback framework, the coefficients are not large enough to reverse the also negative relationship with wages and hence we have the anomalous result that people in Brazil earn more in real terms when the climate is worse.

The patterns for China and India are almost identical. In almost all cases, there is almost no correlation between climate and any of our variables. China's economic divide runs east-to-west rather than north-to-south, like Brazil, which may explain why there isn't a large correlation between climate and nominal wages. Overall, perhaps the most natural

explanation is that the Chinese and the Indians are not wealthy enough to be willing to pay a significant premium to live in places with more temperate climates. Liu and Shen (2014) also find a far weaker relationship between climate and population growth in China than in the U.S. or Europe.

1.4.3 Happiness across Space

Economists like Jeremy Bentham and John Stuart Mill argued that human beings both should and typically do try to maximize pleasure and minimize pain. If they were right, then the modern economists' concept of utility should be synonymous with self-reported happiness or life satisfaction. Yet many if not most economists now reject the Benthamite hedonist approach that equates happiness and welfare. Utility functions, in their modern use, are simply rankings based on preferences. People may sensibly make decisions that appear to lower their level of reported life satisfaction. Parents of young children, for example, typically report lower levels of life satisfaction, perhaps because of enormous time costs in caring for infants (McLanahan and Adams, 1987). This does not mean that those parents have made a mistake. Having children may be rational and increase utility even if it does not increase happiness.

Nonetheless, we follow Glaeser *et al.* (2014a) in believing that happiness can be useful for examining the spatial equilibrium hypothesis even if happiness is not equivalent to utility. Heterogeneity in happiness across space should provide a test of the spatial equilibrium model. Strong differences in happiness can be taken as evidence against spatial equilibrium, and there are modest differences in happiness across space both within the U.S. and across Europe (Pittau *et al.*, 2010). Given the difficulties in attributing heterogeneity in self-reported happiness to small samples or local cultures, we focus on the narrower question of whether happiness rises with income across areas.

If spatial equilibrium holds, we expect to find little or no relationship between happiness and area income. Indeed, happiness may be a proxy for certain urban amenities, and then the spatial equilibrium model might predict that happiness should be lower, not higher, in

richer areas. If there is a positive relationship between income and happiness across areas, then this suggests either that spatial equilibrium doesn't hold, or that different people live in different cities, or that happiness is not capturing welfare.

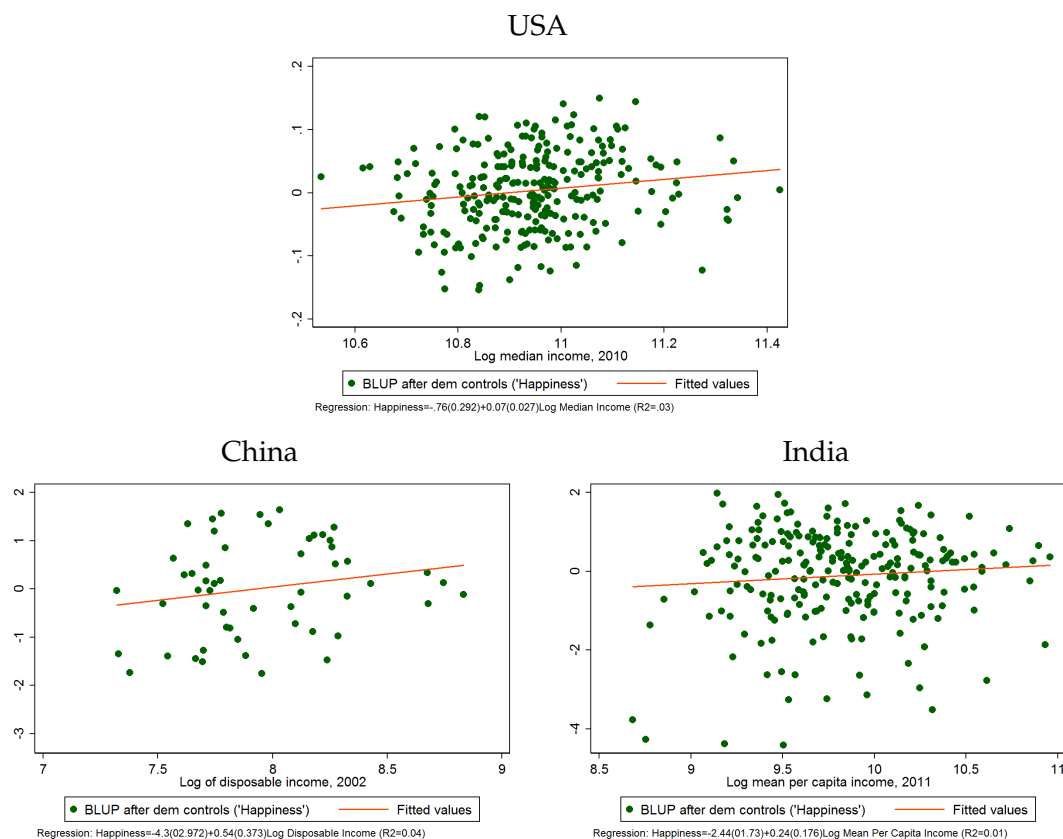


Figure 1.4: *Happiness and income levels*

Note: Samples restricted to areas with urban population of 100,000 or more. **Sources:** See data appendix.

Figure 1.4 shows the relationship between area income and area happiness for the U.S., India and China. The corresponding regressions are also reported in Appendix Table A.1. For the U.S., the relationship is positive but small. If the income of an area doubles, then self-reported life satisfaction increases by seven tenths of a standard deviation. Certainly, given that richer places also have people with higher levels of human capital, this is not enough to challenge the spatial equilibrium assumption in the U.S.

We do not have comparable data for Brazil, but an IPEA (2012) report finds that happiness is actually lower in wealthy southern Brazil and highest in the country's poor and rural

northeast. This finding seems to support the view that there is not a spatial arbitrage opportunity available in moving to Brazil's wealthier area. Other work (Corbi and Menezes-Filho, 2006) confirms that across individuals, Brazilian patterns resemble those in other countries, and that happiness rises with income at the individual level.

The estimated coefficient for Chinese cities is also on the margin of statistical significance, but the point estimate is much larger. As income doubles, self-reported life satisfaction increases by more than five tenths of a standard deviation. There is a great deal of noise in the Chinese data but the coefficient is almost eight times the size of the U.S. coefficient.

India displays a point estimate that is three times larger than the U.S., but the coefficient is imprecisely measured so that we cannot statistically rule-out a coefficient smaller than the one in the U.S. It does seem that richer cities are happier in the developing countries, more so than in the U.S., but the evidence is far from conclusive.

There are at least two interpretations of these results that are quite compatible with the spatial equilibrium framework. One interpretation, again, is omitted human capital, and in that case, richer cities have people who are richer because they are more skilled and we might expect them to have higher happiness levels. A large literature finds a positive link between individual-level education and happiness in many countries (e.g. Gerdtham and Johannesson, 2001; Chen, 2012; Cuñado and Pérez de Gracia, 2012). Another interpretation is that happiness is reflecting urban amenities, which are higher in richer places. In that second case, however, the failure to find much higher prices in high income areas in India becomes even more of a puzzle.

The third explanation is that the spatial equilibrium assumption is just not particularly useful when thinking about China, and especially India. Perhaps the most natural reason why the frictionless spatial equilibrium assumption could fail is that mobility is limited, either because of strong tastes for remaining in places or because of barriers to mobility like China's Hukou system. We now turn to facts about spatial mobility in the four countries.

1.4.4 Mobility in the Four Countries

In principle, the standard spatial equilibrium does not require much mobility. Even if no one moves, housing costs can adjust to keep welfare equal across space. However, immobility can be a sign that the assumptions needed for the Rosen-Roback model do not hold. For example, if individuals were forbidden from moving across state lines, then there would be no reason for utility levels to be equalized. Without labor mobility, regional models revert to having the implications of national models, which certainly do not predict constant utility across space.

In reality, mobility is possible but imperfect, and when immobility is caused by heterogeneity of tastes, then the implications of the Rosen-Roback model weaken. Imagine if villagers have the option to move to a city but they have tastes for living in the village of their youth and those tastes are heterogeneous. The real wage gap between the village and the city will equal the taste for the village held by the marginal migrant. Strong tastes will mean that the real wage gap can be quite large, and immobility is a sign that tastes for one's home are strong. Moreover, there are good reasons to believe that the limitations on mobility differ across countries (Bell and Muhidin, 2009), which may explain the differences in the empirical value of the spatial equilibrium concept.

The U.S. has traditionally been a highly mobile nation, which presumably suggests that many Americans have only weak tastes for their home locales. As Table 1.7 shows, according to the 2000 Census, 21 percent of Americans lived in a different county, state, or country five years ago. About one-half of these moves were just across county lines, which could mean within a single metropolitan area. Still, about one-in-ten American in 2000 had made a major move over the preceding 5 years. The figures were comparable in 1990 and for many previous years.

In 2010, only 13.8 percent of Americans had moved counties, states or countries during the previous five years. Only 7.1 percent had changed states or countries. While these figures are still relatively high by global standards, they do represent a dramatic drop, which is presumably best understood as a reflection of the Great Recession. Underwater

Table 1.7: *Percentage of the population living in a different locality five years ago*

	USA			Brazil		
	1990	2000	2010	1991	2000	2010
Migrants in the last 5 years (% of population)	21.3%	21.0%	13.8%	9.5%	9.1%	7.1%
From same state/prov., different county / dist.	9.7%	9.7%	6.7%	6.0%	5.4%	4.5%
From different state/province	9.4%	8.4%	5.6%	3.5%	3.6%	2.4%
From abroad	2.2%	2.9%	1.5%	0.04%	0.1%	0.14%
	China			India		
	2000	2010		2001	2011	
Migrants in the last 5 years (% of population)	6.3%	12.8%	1.9%	2.6%	2.0%	
From same state/prov., different county / dist.	2.9%	6.4%	1.3%	1.5%	1.2%	
From different state/province	3.4%	6.4%	0.6%	1.0%	0.8%	
From abroad	N/A	N/A	0.02%	0.1%	0.03%	

Sources: See data appendix.

homeowners may have been unable to sell their homes to move during the downturn. Younger people often chose to stay at home during the recession to save costs.

Comparable mobility figures for our other three countries are reported in Table 1.7. Again, the standard is to use a retrospective question of current residents, asking them where they lived five years ago. Censuses typically provide us with this information. We have attempted to use major and minor geographic units in each country that are comparable to states and counties within the United States.

Brazilians are mobile (Fiess and Verner, 2003) but they are less mobile than Americans. Brazil's mobility rate has also declined over time. In 2000, 9.1. percent of the population had made a major or minor move over the previous five years. In 2010, 7.1 percent had made a major or minor move. Major moves are particularly rare. Only 2.4 percent of the population had changed regions, and about one-tenth of one percent of the population were international immigrants. The high fraction of foreign-born remains a relatively special aspect of American society.

In China, our data begins in 2000 and there has been a large jump in mobility between 2000 and 2010. In 2000, 6.3 percent of the population had made a major or minor move

over the previous five years. In 2010, 12.8 percent of the population had moved. Shen (2013) also documents this increase in mobility. Somewhat remarkably, China is now a more geographically mobile country than the U.S., when we consider only major moves. Chinese mobility is particularly remarkable because the Hukou system limits the benefits from moving. If American mobility supports a spatial equilibrium, then surely Chinese mobility does as well.

By contrast, mobility is extremely low in India. Only two percent of the sample had moved during the preceding five years in 2011, and that figure replicates results for 2001 and 1993. Less than one percent of the population had made a major move. Munshi and Rosenzweig (2009) also document a low Indian mobility rate and suggest that immobility may reflect informal risk-sharing relationships in villages. There is very little in-migration from outside the country. These low migration rates seem puzzling given the enormous growth of Indian cities.

It is quite possible that the Indian data actually understates the true amount of functional migration that occurs. This data misses the temporary migrants discussed by Morten (2013). There could be other measurement issues, like listing the migrant's primary place of residence as their home village, even though they are working elsewhere. It is also possible that the surveyors may have undercounted many of the residents living in urban slums. Jedwab and Vollrath (2016) document that urban growth in India and Africa is also driven by high levels of fertility, which suggests how India can combine low mobility rates and substantial urban growth.

The migration rates in Brazil and China are lower than migration rates have historically been for the United States, but they are not dramatically lower than migration rates in the U.S. today. Consequently, there would seem to be enough migration in those countries to make spatial equilibrium models sensible tools for thinking about these countries. However, since China's migration was much lower historically, it may take some time for the process to fully equilibrate. By contrast, migration within India is quite low, and this may help explain some of the Indian facts that seem at odds with the predictions of the frictionless

spatial equilibrium model.

Data cannot prove that a spatial equilibrium exists, but in the U.S., a wide range of facts are quite compatible with the existence of an approximate spatial equilibrium. Our reading of the data suggests that assuming a spatial equilibrium is also reasonable in Brazil, where mobility rates are high and housing costs track incomes. Even China, despite its regulatory limitations on mobility, seems to be moving towards a spatial equilibrium. India is the big outlier, which systematically fails every test of a frictionless spatial equilibrium. Perhaps, these failures represent human capital heterogeneity across space, or strong Indian preferences to remain in birth locations.

The failure of the standard spatial equilibrium model in India seems like an important topic for future research, especially if the Indian results are duplicated for other extremely poor places. Tests that correctly support the existence of a frictionless spatial equilibrium require both that a spatial equilibrium exists, which relies on factor mobility, and that omitted heterogeneity is not so large that it drives the results. We suspect that the spatial equilibrium hypothesis may end up having relatively little empirical power in many places as poor as India because there are profound limits on spatial mobility, deep social ties to place, and profound heterogeneity in human capital and places.

1.5 The Determinants of Urban Success: Agglomeration and Human Capital

The spatial equilibrium hypothesis represents one major part of work on cities within the wealthy world. A second large body of research focuses on the determinants of success across space. One literature focuses on agglomeration economies and why incomes appear to rise with city size. A second literature focuses on skills and urban success, whether at a point in time (Rauch, 1993; Moretti, 2004) or dynamically (Glaeser *et al.*, 1995). We divide this section into three subsections, dealing with agglomeration economies, human capital externalities and the connection between skills and growth in city population and incomes.

1.5.1 Agglomeration Economies

The core idea of agglomeration economies is that productivity increases with the geographic proximity of economic activity. The strength of agglomeration economies helps to determine the actual and optimal size of cities. If agglomeration economies are strong in Brazil, China and India, then this helps to explain why these areas are urbanizing so rapidly. Strong agglomeration economies would also seem to work against policies that act to limit the size of city growth.³

The central fact that justifies economists' belief in agglomeration economies is that wages are higher in larger, denser cities. This fact is buttressed by the connection between productivity and city size and the high commercial rents that are paid within city centers. There are two primary empirical problems with interpreting high urban wages as agglomeration economies. First, some places may be intrinsically more productive, causing wages to rise and density to form. Second, more able people may sort into denser cities (D' Costa and Overman, 2014). Cities may also attract more productive firms (Combes *et al.*, 2012a).

While there has been copious work on these issues for decades (Glaeser and Maré, 2001; Combes *et al.*, 2010), they are still not fully resolved within the U.S. We are not going to resolve them for Brazil, China and India either. Our one approach to address causality is to instrument for current population and density levels with historic values of population and density, which is an approach used by Ciccone and Hall (1996); Rice *et al.* (2006); Combes *et al.* (2010) and other researchers. Regional land area is another common instrument used in European research (Ciccone, 2002; Foster and Stehrer, 2009), but it seems ill-suited for the developing world.

The use of lagged population as an instrument does nothing for the problem of sorting into cities. Even if sorting is only a contemporary event, this sorting may still be shaped

³Theory could justify such policies, even in the presence of agglomeration economies, if there are strong negative externalities associated with contagious disease and congestion. In practice, they can be driven by non-welfare-maximizing motives, or have negative unintended consequences. For example, in Brazil, Feler and Henderson (2011) document strategic under-provision of water to small houses in likely low-income migrant destinations during the period 1980-2000. These "exclusionary policies" can help explain the slower growth and the presence of informal and under-served neighborhoods in economically successful localities.

by historic city sizes. It may slightly reduce the problem of unobserved local productivity shocks if we believe that the shocks that are relevant to current productivity are relatively recent and unrelated to historic population. We see this instrumental variables strategy as a robustness check rather than as any proper solution to the two great problems inherent with estimating agglomeration economies.

Our goal remains a comparison of the four countries using standard methods, not advancing new methods for solving old problems. As such, we will stick with standard—if flawed—approaches that are easy to replicate across countries. We will estimate the coefficient on density and overall population size in different regressions. We do not see these two coefficients as estimating different things, but rather different ways of getting at the same concept of agglomeration.

Each of these measures can be flawed, and it seems sensible for us to show results for both. For example, consider two agglomerations that are intrinsically identical. In one case, the metropolitan area is drawn to include a lot of extraneous farmland. In the other case, the borders are drawn tightly around the agglomeration. The density level will be misleading here, but the population level will not. Conversely, consider a case in which there are three distinct and identical agglomerations, two of which are grouped into a single metropolitan area. Population will, in this case, be misleading since we are counting two as one, but density will accurately capture the effective agglomeration size.

The first column of Table 1.8 shows the results for the U.S. All specifications include human capital controls. The agglomeration elasticity coefficient of log wage on either log density or log population is about .05. This suggests that if either density or population size doubles, wages rise by about five percent. When both effects are included in the same regression, the coefficient on each is about .03, but we think that these two measures of agglomeration are too similar to put much faith in the ability of a regression to estimate the effects separately.

In the second panel of Table 1.8, we show results using population or density in 1980 as an instrument for population or density today. U.S. metropolitan-area population levels move slowly, so unsurprisingly the coefficients are quite similar to the results in the first

Table 1.8: *Income and agglomeration, 2010*

	USA (MSAs)	Brazil (Microregions)	China (Cities)	India (Districts)
	Log wage	Log wage	Log wage	Log wage
OLS regressions				
Log of urban population	0.0538*** (0.00720) R2=0.255	0.052*** (0.013) R2=0.321	0.0875 (0.0708) R2=0.014	0.0770*** (0.0264) R2=0.251
Log of density	0.0457*** (0.00865) R2=0.235	0.026** (0.010) R2 = 0.318	0.192*** (0.0321) R2=0.237	0.0760*** (0.0195) R2=0.257
Observations	28.5M	2,172 K	147K	9,778
IV1 regressions				
Log of urban population	0.0559*** (0.00753) R2=0.256	0.051*** (0.014) R2 = 0.321	0.0320 (0.102) R2=0.173	0.160 (0.0998) R2=0.237
Log of density	0.0431*** (0.00888) R2=0.253	0.026** (0.011) R2 = 0.318	0.169*** (0.0367) R2=0.240	0.0828*** (0.0218) R2=0.253
Observations	28.5M	2,172 K	143K	7,627
IV2 regressions				
Log of urban population	0.0764*** (0.0130) R2=0.255	0.015 (0.021) R2 = 0.315	0.320* (0.156) R2=0.117	0.233** (0.0963) R2=0.224
Log of density	0.0493*** (0.0173) R2=0.253	0.015 (0.012) R2 = 0.315	0.323*** (0.0847) R2=0.242	0.0749*** (0.0229) R2=0.256
Observations	28.5M	1,998 K	112K	5,245
Educational attainment controls	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes

Note: Regressions at the individual level, restricted to urban prime-age males in areas with urban population of 100,000 or more. All regressions include a constant. Standard errors clustered at the area level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

panel. In the third panel, we use population or density in 1900 as instruments. If anything, the coefficients become slightly stronger but they do not alter significantly. The core U.S. coefficients are quite robust and don't change much if we rely on historic population levels for our variation in population or density.

The second column gives results on Brazil. A meta-analysis by Melo *et al.* (2009) has shown sizable differences in measured agglomeration economies across countries, so we don't expect this elasticity to match that found in the U.S. Yet, the population elasticity is .05, which is essentially equivalent to the U.S. The density coefficient is .026, which is slightly lower than in the U.S. These coefficients are somewhat smaller than those estimated by Fontes *et al.* (2009), probably because they also include non-urban residents. These coefficient are also comparable to results found in Great Britain (Rice *et al.*, 2006), Netherlands (Groot *et al.*, 2014), and Europe more broadly (Ciccone, 2002; Foster and Stehrer, 2009). The gap between urban and rural earnings in Brazil is enormous, and much larger than in the U.S., but the relationship between earnings and density across urbanites is somewhat weaker in Brazil than in the U.S.

The second panel shows that using population and density in 1980 as instruments has almost no impact on the estimated coefficients. Using 1920 population and density as instruments cause the estimates coefficients to drop and become insignificant as shown in the third panel. One interpretation of that fall might be that some of the current correlation between agglomeration and income in Brazil reflects omitted 20th century local productivity shocks that pushed both variables in the same direction.

The Chinese data, shown in the third column, is somewhat unusual. The coefficient on area population is larger than the coefficients for either the U.S. or Brazil, but statistically indistinct from zero. This may reflect much more noise in both variables. The coefficient on density is extremely large, close to .2, and statistically quite robust. Combes *et al.* (2013) also find an agglomeration coefficient in China that is roughly three times as large as standard coefficients found in the west. Results for Japan are also higher than in the west (Tabuchi and Yoshida, 2000). In the second panel, using 1980 population as an instrument, we find that

the population coefficient is small and insignificant. Using 1950 population as an instrument, the population coefficient grows dramatically and becomes marginally significant. We suspect part of the issue is that many of the Chinese “cities” are quite large and may include workers who are not really in the same metropolitan agglomeration. This would lead to measurement error in the dependent variable, which should bias the estimated coefficient towards zero. Alternatively, this weak relationship may reflect an underlying non-linear relationship between population and productivity as found by Xu (2009).

By contrast, the density results are large and robust in all three specifications. The high coefficients suggest China is experiencing dramatic agglomeration economies, but that they are better measured by density than total population. The density coefficient is about four times higher than in the U.S., which suggests that productivity is dramatically higher in places where population is concentrated. In the case of China, the instrumental variable estimates help dispel the fear that this correlation is the reflection of post-1980 political shocks to particular areas, like the special economic zones.

Many other studies have also found agglomeration economies in China (World Bank, 2014). Pan and Zhang (2002) use firm production data and show that as city size doubles, firm productivity increases by 3.6 percent. Lin *et al.* (2011) find significant agglomeration economies in the textile industry. Ke and Yu (2014) find that productivity growth is tightly tied to industrial agglomeration. Hering and Poncet (2010) find that market access significantly determines wage differences across Chinese metropolitan areas, which is one explanation for these agglomeration economies. Interestingly, Ke (2010) finds that it is the size of the industrial sector, not employment density that determines productivity.

The Indian results in the fourth column show agglomeration effects that are somewhat larger than in the U.S., which echoes the findings of Lall *et al.* (2004). The coefficients on both current population and density are approximately .075, which is about 50 percent larger than in the U.S. In the second panel, we use population in 1980 as an instrument, and the coefficients increase in magnitude. The coefficient on population becomes statistically insignificant, however. In the third panel, we use population and density in 1951 as

instruments. The coefficient on population rises further still and becomes statistically significant again. The coefficient on density remains quite close to the ordinary least squares estimate.

One of the standard tests for examining whether the estimated agglomeration economies represent productivity or sorting is to look at real wages. If workers in cities have higher levels of human capital then they should earn more in real terms, not merely in nominal terms. If workers in cities are intrinsically identical to workers elsewhere, then they should not be earning higher real wages. Naturally, one difficulty with interpreting these real wage coefficients is that higher real wages in cities might also reflect compensation for adverse urban amenities.

Table 1.9 shows the correlations between real wages and density in these four countries. Again, we use historical values of density and population as instruments for the current population and density levels. We also define real wage using the rent data as the logarithm of wage minus .33 times the logarithm of rents. Appendix figures A.2 and A.3 show the relationship between area average log earnings residuals and area urban population and area density, respectively.

In the U.S., we find the real wages coefficient of both variables is .02. The coefficients remain about the same using the 1980 value of population as an instrument, but the coefficients rise significantly when we use the 1900 values as instruments. These results differ slightly from Glaeser and Maré (2001), who found no relationship between real wages and area population across American metropolitan areas, and Glaeser and Gottlieb (2006) who found that a real wage premium existed in 1970 but not in 2000.

There are two natural reasons why these results differ. First, real wages can be measured with significantly more precision in the U.S. using better data, such as the American Chamber of Commerce Real Estate Association price indices. We did not do that here to ensure comparability with other countries. Second, in those papers the regressions are weighted at the metropolitan area level, while here they are weighted at the person level. Nevertheless, these results still suggest that the majority of the agglomeration effect in the

Table 1.9: *Real income and agglomeration, 2010*

	USA (MSAs)	Brazil (Microregions)	China (Cities)	India (Districts)
	Log real wage	Log real wage	Log real wage	Log real wage
OLS regressions				
Log of urban population	0.0190** (0.00916) R2= 0.067	0.011 (0.010) R2=0.310	-0.0313 (0.0307) R=0.174	0.0688** (0.0298) R2=0.240
Log of density	0.0219 (0.0134) R2=0.068	0.002 (0.007) R2=0.309	0.0516** (0.0166) R2=0.179	0.0691*** (0.0213) R2=0.244
Observations	28.5M	2,172 K	147K	2,102
IV1 regressions				
Log of urban population	0.0209** (0.0102) R2=0.068	0.009 (0.010) R2 = 0.310	-0.0664 (0.0485) R2=0.174	0.116 (0.0927) R2=0.243
Log of density	0.0230* (0.0134) R2=0.068	0.001 (0.007) R2 = 0.309	0.0345* (0.0175) R2=0.179	0.0647** (0.0255) R2=0.241
Observations	28.5M	2,172 K	143K	1,649
IV2 regressions				
Log of urban population	0.0466** (0.0190) R2=0.065	-0.017 (0.016) R2 = 0.305	0.0648 (0.0743) R2=0.161	0.208** (0.0840) R2=0.244
Log of density	0.0419** (0.0163) R2=0.067	-0.008 (0.008) R2 = 0.307	0.0665 (0.0625) R2=0.179	0.0512* (0.0263) R2=0.241
Observations	28.5M	1,998 K	112K	1,141
Educational attainment controls	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes

Note: Regressions at the individual level, restricted to urban prime-age males in areas with urban population of 100,000 or more. All regressions include a constant. Standard errors clustered at the area level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

U.S. does not reflect sorting, since these coefficients are so much smaller than the coefficient on nominal wages.

For Brazil, there is no evidence suggesting that there is a real-wage premium in bigger cities. All of the coefficients are quite close to zero. This fact pushes against the view that the Brazilian urban wage premium reflects omitted human capital characteristics.

In the third column, we show results for China. In previous work, ? found that real incomes initially rise with city size in China and then decline in the largest cities. This is consistent with the fact that we find in our linear specification a negative but statistically insignificant coefficient. The density coefficient remains positive, but is much smaller than the density coefficient for nominal wages. Unless amenities are much higher in denser areas in China, this suggests that most of the nominal-wage premium received in denser Chinese areas does not reflect sorting of higher-ability people into those areas.

Finally, the Indian regressions in the last column show that the nominal-wage premium in India that is associated with both density and population is essentially the same as the real-wage premium. The results are comparable when we use the recent lag of population as an instrument and implausibly large when we use the long lag of population as an instrument. In a sense, these results are unsurprising given that we have already seen that rents and incomes are essentially unrelated across Indian cities and real wages, in our regressions, are just nominal wages corrected for rents.

There are two plausible interpretations for this strong relationship between real wages and agglomeration in India. First, it is possible that Indian urbanites do, in fact, have much higher human capital levels than rural Indians. Migration rates are low, and education quality may be quite different between urban and rural areas. Second, it is also possible that our rent measures are so bad that these regressions have basically no information value.

We supplement these real-wage regressions with Appendix Figure A.4, which shows the connection between happiness and area population for the U.S., China and India, and Appendix Table A.1, which report the corresponding regressions. In the U.S., there is essentially no correlation between area population and happiness, which corroborates the

finding that there is little real-wage premium in larger cities. The correlation for China is also statistically indistinct from zero. This finding is corroborated by a small literature on happiness and urbanization within China. Knight and Gunatilaka (2010) find that wealthier urbanites in China are actually less happy than rural dwellers, perhaps suggesting that migrants are forgoing current well-being for future economic prosperity. Cheng *et al.* (2014) find that second generation rural-urban migrants are less satisfied than the first generation of migrants. Within India, however, there is a large correlation between self-reported happiness and area population, which is consistent with the positive real-wage premium associated with area size in that country.

The first four columns of appendix Table A.2 repeats these regressions where rent, rather than income or real income, is the dependent variable. Columns five and six of the same table report these regressions with housing value as the dependent variable for the U.S. and China. Rents rise significantly with population and density in Brazil, China and the U.S. Housing values raise significantly with population and density in the U.S., but only with density in China. Rents are unrelated to population and density in India. Again, this highlights the difficulties of using rental data in the developing world.

Taking this evidence at face value, it appears that agglomeration economies are stronger in India and China than in Brazil or the U.S. There is some possibility that the robust agglomeration effects observed in India are driven by sorting. Still, there is every reason to believe that the literature that explains and examines agglomeration economies in the developed world will continue to be relevant in the developing world. We hope that future work will look at whether developing world settings also replicate other facts that appear to be true in the west, such as faster wage growth in cities (Glaeser and Maré, 2001; ?) and rising agglomeration effects over time (Brülhart and Mathys, 2008).

1.5.2 Human Capital Externalities

One prominent theory of agglomeration economies is that knowledge and ideas spread across space. The theory also predicts human capital externalities: people who work and

live in better educated areas will themselves become more productive because they will accumulate more human capital. Rauch (1993) and Moretti (2004) are two key contributors to this literature, which consistently estimates a significant relationship between area-level education and earnings, holding individual-level education constant. Human capital externalities have also been documented in Germany (Heuermann, 2011), Italy (Dalmazzo and De Blasio, 2007), Spain (Serrano, 2003) and Europe more generally (Rodríguez-Pose and Tselios, 2012).

Naturally, the same two problems that bedevil agglomeration regressions also trouble human capital externality regressions. It is possible that some omitted productivity variable both disproportionately attracts skilled people to an area and increases wages. It is even more likely that places with higher levels of education also have people with higher levels of unobserved human capital. The issue of sorting on unobservable human capital is even more severe than in the case of agglomeration economies, because the key independent variable is the level of observable human capital, not the level of density. It seems quite likely that unobservable and observable human capital move together.

Acemoglu and Angrist (2001) attempt to address the omitted variables problems by using the variation in compulsory education rules by state to estimate human capital spillovers. They find no spillovers with this approach, but their source of variation is not ideal for estimating standard human capital spillover models. For example, Glaeser (1999) provides a model of learning from neighbors in which raising the top of the human capital distribution will generate spillovers, but raising the bottom of the distribution (which is essentially the effect of raising the minimum school-leaving age) will not. Moretti (2004) tries different instrumental variable approaches and finds consistent support for the existence of such spillovers. Falck *et al.* (2011) use the location of Baroque Opera Houses as an instrument for human capital in Germany, and find evidence for both static and dynamic effects of local skills.

Again, we have no magic bullet for addressing the sorting and omitted variables problems. Primarily for robustness, we will show two instrumental variable regressions as well

as the ordinary least squares regressions. We instrument using a recent but lagged value of the local education variable, typically from 1980 in the first panel of the table. We also follow Moretti (2004) and use the demographic structure of the city as an instrument for education for the US, Brazil and India in the second panel. Specifically, we predict age-group shares of the area population based on the population shares of each group in the area in 1980. We then attribute to each age group the education level that is typical for that age group in 2010. Essentially, our instrument is the predicted education for the area if it had kept its age structure from 1980 and if everybody had his or her age group's average education. For China, we use the number of educational institutions in 1948 as the instrument in the second panel. In all cases, our education measure will be the share of the population with a post-secondary degree, which has the advantage of being readily available and relatively comparable across areas.

In Table 1.10 we present two columns for each country. In the first column we show the impact of area-level education with individual human capital controls, and the in the second we also control for area-level density. In most cases, controlling for area-level density makes little difference, since income is much more tightly linked to area-level education than area-level density. The density coefficients, however, typically fall when we also control for area-level education. In the U.S., India and China, density remains significant and positive when we control for education. The density coefficient actually reverses sign in Brazil. Appendix Figure A.5 shows the relationship between area average log earnings residuals and area-level education.

The first two columns of Table 1.10 show our results for the United States. The first row in the first two columns shows that as the share of adults in the area increases by 10 percentage points, wages increase by about .10 log points or about ten percent, holding an individual's own (measured) human capital constant. These effects are somewhat larger than those reported by Moretti (2004), where a ten percent increase in the share of adults with college degrees is associated with an eight percent increase in earnings, holding individual human capital constant. Perhaps the most natural explanation for this difference is that the measured human capital externalities have been rising over time (Glaeser and Saiz, 2004).

Table 1.10: Human capital externalities, 2010

	USA (MSAs)		Brazil (Microregions)		China (Cities)		India (Districts)	
	Log wage	Log wage	Log wage	Log wage	Log wage	Log wage	Log wage	Log wage
OLS regressions								
Share of Adult population with BA	1.272*** (0.155)	1.001*** (0.200)	3.616*** (0.269)	4.719*** (0.440)	6.743*** (1.088)	5.262*** (0.862)	3.215*** (0.851)	1.938** (0.841)
Log of density		0.0241*** (0.00746)		-0.029*** (0.008)		0.112*** (0.0199)		0.0542*** (0.0169)
R-squared	0.26	0.255	0.342	0.346	0.120	0.139	0.256	0.255
Observations (thousands)	34M	27M	2,172 K	2,172 K	147K	147K	12K	12K
IV1 regressions								
Share of Adult population with BA	1.237*** (0.202)	1.126*** (0.231)	2.985*** (0.332)	3.784*** (0.486)	6.572*** (0.925)		2.911*** (0.988)	2.124** (1.074)
Log of density		0.0216*** (0.00769)		-0.018** (0.009)				0.0425** (0.0178)
R-squared	0.254	0.255	0.341	0.344	0.120		0.240	0.243
Observations	27M	27M	2,172K	2,172 K	147K		11 K	11K
IV2 regressions								
Share of Adult population with BA	1.594*** (0.380)	0.956** (0.396)	4.166*** (1.059)	6.705*** (1.756)	7.189*** (1.437)		8.126** (3.458)	7.989 (5.521)
Log of density		0.00654 (0.0155)		-0.052** (0.023)				-0.0107 (0.0615)
R-squared	0.228	0.232	0.341	0.341	0.120		0.206	0.212
Observations (thousands)	17M	16M	2,172 K	2,172 K	147K		10 K	10 K
Educational attainment controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: Regressions at the individual level, restricted to urban prime-age males in areas with urban population of 100,000 or more. All regressions include a constant. Standard errors clustered at the area level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

The estimated coefficients when we use the 1980 value of schooling as an instrument for education in 2010 are somewhat higher but not very different than the OLS results. Using our demographically based instrument, the estimate rises without the density control and falls slightly when we control for density. Across all three specifications, the coefficients with controls consistently suggest that earnings increase by ten percent when the share of adults increases by ten percentage points, holding individual education constant.

The third and fourth columns show our results for Brazil. In five out of our six specifications, the coefficients range from 3.0 to 4.7, which is about three times higher than the coefficients for the U.S. The instrumental variables estimate that uses demographic composition is an outlier, with a coefficient close to seven. It appears clear that the estimated impact of skills on nearby earnings is higher in Brazil than in the U.S. However, the variation in education across Brazilian cities is much smaller than in the U.S. (.033 vs. .06), so the impact of a one standard deviation increase in the level of education has only a 50 percent larger impact on wages in Brazil, if we accept a coefficient of approximately 3.5. It is also quite possible that heterogeneity in unobserved human capital is larger in Brazil, which might also explain why the estimated coefficient is so large. Freguglia and Menezes-Filho (2012) estimate local wage effects with migrant data, and find that estimated local wage differences diminish significantly when they control worker fixed effects, which suggests that omitted human capital factors may be quite important.

The fifth and sixth columns show results for China. We were unable to re-access the data, so we only have the specification with recent density and with the instruments, not both together. The estimated human capital spillover coefficients range from 5.2 to 7.2, which are also much larger than in the U.S. If these coefficients are taken literally, then a ten percent increase in the share of adults with a college education in a Chinese city is associated with an over sixty percent increase in earnings. These results are roughly in line with the work of Liu (2007) who also finds large human capital externalities in China. Fu and Gabriel (2012) find that the migration patterns of high-skilled workers suggest that they are particularly responding to human capital externalities in China.

The variation in education rates across cities is also smaller in China than in the U.S., but larger than in Brazil. The impact of one standard deviation increase in area level education in China is associated with a 20 percent increase in earnings if we use the most conservative estimate. The comparable figure with the U.S. is a six percent increase in earnings.

The same pattern reappears for India in the seventh and eighth columns. With the exception of the demographically constructed instrumental variable regression, the coefficients range from 1.9 to 3.2. These are higher than in the U.S., but lower than either Brazil or China. A one standard deviation increase in area level education (.033) in India is associated with an approximately seven percent increase in earnings, if we accept a coefficient of 2.1, which is only marginally higher than in the U.S. A ten percentage point increase in the share of adults with college degrees is associated with a 21 percent increase in earnings, which is more than double the U.S. figure.

All three developing countries show a similar pattern. Education levels vary less across space than in the U.S. The estimated effect of area level education on wages is much higher than in the U.S. One class of explanations for these differences assumes that these human capital coefficients are largely spurious. A second class of explanations assumes that they are real.

For example, if we believe that the coefficients are spurious because of sorting on unobservables, then the correlation between unobservables and area level education is likely to be much higher in countries with less variation in education. If the regression was univariate, then the bias created by unobservable human capital would equal $\frac{Cov(\log(Wage), Unobservable\ Human\ Capital)}{Var(Observed\ Human\ Capital)}$. If variance of observed human capital is small, then the bias blows up proportionately.

A second type of explanation assumes that the effect is real. In that case, the high coefficients might mean that local human capital is more valuable in developing world cities than elsewhere. Local human capital might be more important in places that have lower levels of development and less such capital.

We will not try to resolve these issues here, but we believe that these extremely high

measured levels of human capital externalities especially in Brazil and China suggest that this is an important topic for future research. If these results reflect a true human capital spillover, then developing world cities success really depends on education. We now ask whether education also seems to have dynamic effects on cities in the developing world.

1.5.3 Urban Growth and Human Capital

While Gibrat's Law tells us that urban growth is not correlated with initial population levels, a long literature documents the connection between skills and subsequent growth within the United States and internationally. Glaeser *et al.* (1995) documented the correlation between the share of adults with college degrees and population growth between 1960 and 1990 within the U.S. Subsequent work (Glaeser and Saiz, 2004; ?) has shown that this correlation persists for more recent periods, and that skills also predict the growth of income at the local level within the U.S. Simon and Nardinelli (2002) show that 19th century skills predict growth over the next century within the U.S. A connection between skills and growth also holds within German (Südekum, 2010), Spain (Ramos *et al.*, 2010), and Europe more broadly (Badinger and Tondl, 2002; Rodríguez-Pose and Vilalta-Bufí, 2005). Using a sample of over 1,500 regions from 83 countries, Gennaioli *et al.* (2014) also find a connection between initial levels of human capital and regional growth.

Several hypotheses have been advanced to explain the skills-growth connection. Glaeser *et al.* (1995) originally suggested that skills sped growth by encouraging the increase of local productivity. Glaeser and Saiz (2004) find that measured human capital externalities have grown stronger over time. Alternatively, skills may have also been correlated with amenities that have gotten more desirable, or with good government. Our goal is not to distinguish between these hypotheses, but rather to look at whether the skills growth connection also exists in Brazil, China and India.

Table 1.11 shows our results when we regress population and income change between 1980 and 2010 on the share of adults with tertiary education in 1980. We present four regressions for each country: one with just education for both income and population

growth and one with added controls, which is our preferred specification. The added controls include initial income and population, as well as climate controls. Unlike other regressions in the paper, Table 1.11 uses total rather than urban populations.⁴

The first regressions in the upper panel show results for the U.S. A ten percent increase in the share of the population with a college degree in 1980 is associated with a 17 percent increase in population between 1980 and 2010. This effect is statistically and economically meaningful and it is shown in Figure 1.5. The second regression shows that the estimated coefficient increases to 2.16 when we control for initial income and population. The t-statistic is over six, which indicates the strength and robustness of the connection between skills and area population growth over this time period.

The lower two regressions show the results for income growth within the U.S. The coefficient on initial schooling is .5. With the controls, the coefficient rises to .9, which is both statistically significant and economically meaningful. Ten percentage points more college graduates in 1980 is associated with about nine percent higher income growth between 1980 and 2010. Initial income controls are particularly important because income levels typically mean revert, and higher education areas in 1980 also had higher income levels. Controlling for initial income corrects for the mean reversion of income.

Our next four regressions show the results for Brazil. The population growth pattern is much stronger than in the U.S., and the correlation with initial skills also becomes stronger when we control for other variables. A ten percentage point increase in the share of the population with college degrees in 1980 is associated with about 49 percent more population growth between 1980 and 2010. The coefficient increases to over five when we control for other variables.

In the lower panel, we show the correlation between skills in 1980 and income growth across Brazilian microregions. A one percent increase in the share of adults with college

⁴In 1980 Brazil, China and India had much lower urbanization rates than the U.S., and urbanization rates differed significantly across areas within countries. Regions with lower urbanization tend to have lower formal education levels (Young, 2013). If city growth in low-urbanization regions is disproportionately driven by own-region rural-urban migration, there can be a spurious connection between initial low area education levels and faster growth in urban population. Looking at total populations mitigates this concern, and allows us to better approximate the growth effect of human capital across areas.

Table 1.11: Human capital and growth, 1980-2010

	USA (MSAs)		Brazil (Microregions)	
	– Log change in population, 1980-2010 –			
University graduates (%) in 1980	1.702*** (0.398) N=218 R2=0.078	2.164*** (0.333) N=218 R2=0.511	4.867*** (0.696) N = 248 R2 = 0.120	5.719*** (1.099) N = 248 R2 = 0.311
	– Log change in average wages (income), 1980-2010 –			
University graduates (%) in 1980	0.493*** (0.164) N=218 R2=0.040	0.901*** (0.151) N=218 R2=0.380	13.434*** (1.871) N = 248 R2 = 0.253	12.208*** (0.850) N = 248 R2 = 0.859
	China (Cities)		India (Districts)	
	– Log change in population, 1980-2010 –			
University graduates (%) in 1980	18.99*** (6.049) N=253 R2=0.038	21.93*** (5.185) N=250 R2=0.608	0.143 (0.245) N=420 R2=0.001	0.343 (0.256) N=362 R2=0.100
	– Log change in average wages (income), 1980-2010 –			
University graduates (%) in 1980	-36.33*** (4.364) N=252 R2=0.217	-4.621 (4.471) N=249 R2=0.541		
Initial income levels control	No	Yes	No	Yes
Initial population control	No	Yes	No	Yes
Climate amenities controls	No	Yes	No	Yes

Note: All figures correspond to area-level regressions restricted to areas with total population of 100,000 or more in 1980. All regressions include a constant. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

degrees in 1980 is associated with an approximately 13 percentage point increase in income growth between 1980 and 2010. The effect falls to 12 percentage points, when we add the other controls, but this is still 13 times larger than the coefficient within the U.S. If we think in terms of the impact of one standard deviation of the skill variable, which is 1.8 times larger in U.S. than in Brazil, these differences look smaller but still sizable. Chomitz *et al.* (2005) also find significant positive effects of education on subsequent regional growth in Brazil.

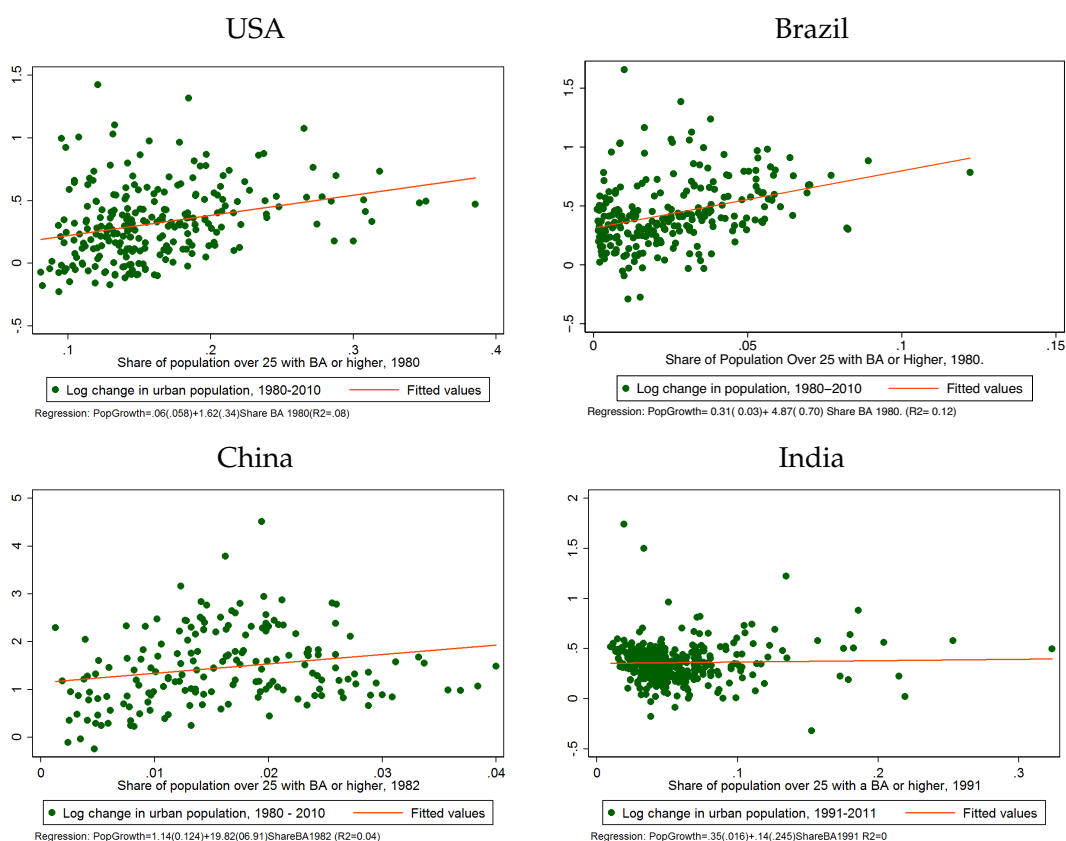


Figure 1.5: University graduates share and population growth 1980-2010

Note: Samples restricted to areas with urban population of 100,000 or more. **Sources:** See data appendix.

Higher levels of skills in 1980 is associated with a relatively larger increase in population growth within the U.S. and a relatively larger increase of income growth in Brazil. One possible explanation for this difference is greater mobility of labor and capital in the U.S. If Americans move more readily, then America will see larger population shifts and smaller

income shifts than Brazil in response to the same local productivity shocks. Greater labor mobility will smooth out the income differences.

The third panel shows results for China, where education is even more strongly associated with population growth. This result corroborates the findings of Fleisher *et al.* (2010) who show that both human capital positively impacts both output and productivity growth in China. A one percentage point increase in the share of adults with college degrees in 1980 is associated with 19 percentage points more population growth between 1980 and 2010. The impact is even larger when we control for other initial variables. A one standard deviation increase in an American area's education is predicted to increase growth by about 12 percent over thirty years. A one standard deviation increase in a Chinese area's education is predicted to increase population growth by around 52 percent. Again, the Chinese data supports the view that urban success is quite closely correlated with initial skills in the developing world.

The lower two regressions for China show a somewhat more mixed picture for income growth. Without other controls, there is a negative correlation between initial education and income growth across Chinese cities. With initial controls, the negative effect is much smaller and statistically insignificant.

Perhaps the most natural explanation for the Chinese pattern is that cities were far from an equilibrium in 1980, because of profound restrictions on mobility in China. In a world in which migration is initially forbidden and where skilled cities are more productive than unskilled cities, allowing free migration will cause the population of skilled cities to soar and the incomes of skilled cities to decline. As population flows rapidly into skilled cities after liberalization, wages in those areas will consequently decline.

The final panel shows results for India, where we only have population growth. Without other controls, the impact of initial skills on population growth is weakly positive and imprecisely measured. With controls, the estimated coefficient becomes larger but still much smaller than in the other countries, and remains statistically insignificant. As in the human capital externality regressions, the link between education and urban success seems weaker

in India than in the other developing countries.

Overall, however, this section supports the view that agglomeration and human capital are strong determinants of urban success in the developing world. Area size is more strongly correlated with income in China and India than in the U.S. and Brazil. Area education has a much stronger connection with income in the developing world than in the U.S. Area education also strongly predicts population growth in Brazil and China, if not in India, as well as income growth in Brazil. We conclude from these facts that the long literatures on agglomeration economies, human capital externalities and growth and skills in the developed world are likely to be relevant in the developing world as well.

1.6 Conclusion

In this paper, we performed three types of comparisons. First, we compared the basic patterns of urban size and growth in the four large countries. Second, we looked at whether U.S. based tests of the Rosen-Roback framework yield similar results in Brazil, China and India. Third, we looked at whether agglomeration economies and human capital externalities seem to operate similarly in the four countries.

The Zipf's law distributions were not identical across the four countries. Most notably, India, China and Brazil had too few ultra-large cities relative to what Zipf's law would predict. The natural explanation for this fact is that congestion disamenities would be too severe for cities of 50 or 100 million people. From 1980 to 2010, Gibrat's law seems to hold for the U.S. and Brazil, but not for India and China.

The most basic Rosen-Roback fact is the strong correlation between income and housing costs across metropolitan areas in the U.S., although even the U.S. quantitative relationship is too small relative to the predictions of theory. The relationship between income and rents is similar in Brazil, China and the U.S. India, surprisingly, shows no spatial relationship between income and rents, which may reflect data issues or chaotic rental markets. In the U.S., real wages decline in places with better climates. This is not true in any of the other countries. There is little relationship between income and happiness across U.S. areas, which

is also compatible with spatial equilibrium. The relationship is stronger, though imprecisely measured, in China and India.

All together, the standard spatial equilibrium framework fits the data far more poorly in India than in the U.S. and Brazil. China's data is more consistent with the spatial equilibrium predictions than India, but not as much as Brazil. One explanation for this failure is that unobserved human capital heterogeneity is far more severe in India and China than in the more urbanized, richer countries. A second explanation is that mobility is limited, either by rules, such as China's Hukou system, or by strong place-based preferences such as those related to cast-based social networks in India.

While the standard spatial equilibrium framework fares much more poorly in the less-urbanized developing world data, the urban productivity and growth relationships are far stronger in those countries. For example, the coefficient when individual income is regressed on area density is around .05 in the U.S. and .03 in Brazil. The coefficient rises to .2 for China and .08 for India. We cannot rule out that this relationship is driven by unobserved human capital, but we can say that the within-country link between density and prosperity in these places does seem remarkably strong.

Similarly, the connection between human capital and area success is also stronger in the developing world countries. For example, the core human capital externality coefficient, when log of earnings is regressed on the share of adults with a college degree or more, is approximately one in the U.S., controlling for individual human capital characteristics. In Brazil, China and India, the same coefficient ranges from two to five. This enormously strong link between area skills and area earnings, controlling for individual skill, may be driven by omitted human capital but it is certainly worth of more research.

As in the U.S., skill also predicts urban growth in Brazil and China. In our preferred specification, a one percentage point increase in the share of college graduates in 1980 is associated with a six percentage point increase in population between 1980 and 2010 in Brazil and a 22 percent point increase in population in China. The effect is much weaker in India. Human capital is also strongly associated with income growth in Brazil.

Taken together, these results suggest that the U.S. facts do matter for the developing world, but they matter more in some places than others, and they matter more in some areas of study than others. Across the board, Brazil is the most like the U.S. and China is second most like the U.S. India is different, probably because of its extremely low mobility rates and much lower income levels.

Our interpretation of these results is that skills and agglomeration impact productivity globally in rich and poor countries alike, but that a spatial equilibrium evolves over time. In the poorest places, social ties to home communities are strong. Historically, they provided safety and sustenance. As nations evolve into wealthy market economies with more homogenous human capital, a spatial equilibrium may eventually appear in countries like India, where it still has not emerged.

Developing world urbanization is among the most important social phenomena globally, but we know much more about developed world urbanization. This paper has shown that some, but not all of that developed world knowledge can be exported to Brazil, China and India. The facts from the west must now be supplemented with a robust research agenda on developing world cities.

Chapter 2

Gender-Segmented Labor Markets and the Effects of Local Demand Shocks

2.1 Introduction

For decades, researchers and policymakers have been interested in the question of how local economies react – in terms of wages, employment and real estate prices – to changes in labor demand. The answers to this question have shaped our understanding of the effectiveness and welfare consequences of local development policies (Moretti 2011). In this literature, researchers have typically assumed away gender differences in the labor market. However, gender differences exist, are large, and are likely to matter. Because male and females tend to segregate into different industries and occupations (Olivetti and Petrongolo 2014), labor demand can disproportionately favor men or women depending on which industries are growing faster. And because women are less likely than men to locate away from their families (Sorenson and Dahl 2016; Gemici 2011; Costa and Kahn 2000) and more likely to interrupt their careers at the time of marrying and having children (Goldin *et al.* 2017; Bertrand *et al.* 2010), increases in demand for male and female labor can have very different effects on migration and labor force participation. This paper incorporates gender into the analysis of labor demand shocks and studies how local economic outcomes respond depending on whether the new jobs favor male or female employment.

I first develop a framework to illustrate the theoretical mechanisms at play and generate predictions on the effects of gender-specific labor demand shocks. Specifically, I embed gender segmentation in the labor market and joint mobility frictions for couples in a standard spatial equilibrium model in the tradition of Rosen (1979) and Roback (1982). These modifications of the canonical framework yield an equilibrium in which local populations, employment, wages and housing rents can respond asymmetrically to equivalent shocks in male and female labor demand.

My model assumes that male and female workers are employed in different industries, each producing an intermediate good, and each with its own productivity shifter subject to exogenous shocks. Intermediate goods are ultimately aggregated as imperfect substitutes to produce a final generic good. On the supply side, individuals have one unit of labor, which they allocate to the workforce if the wage is weakly larger than an exogenous cost of participating in the workforce. I assume that this cost is stochastic, and that the support of its distribution starts at a higher value for women than for men, reflecting extensive work that shows that women tend to face higher opportunity costs of labor force participation (Ponthieux and Meurs 2015).

All individuals are married and each female-male pair constitutes a household. Households choose locations to optimize their combined net wages, housing rents, and amenities; and migration arbitrages away household-level welfare differences across regions. The model predicts that, due to higher cost of participation, the contribution of females to household income will be smaller in expectation, and couples choosing a joint location will be more likely to migrate in response to male rather than female work opportunities.

A key insight of the model is that, because of these gender differences in migration elasticities, local labor and housing markets respond asymmetrically to equivalent shocks to male and female labor demand. If demand for male labor increases, it leads to local population growth and to shifts in female labor supply, as migrant male workers and their spouses move in. Housing demand increases with population, pushing housing rents and compensating differentials in wages up; while the relative abundance of female labor pushes

their wages down. In contrast, shocks in the demand for female labor lead to smaller migration adjustments and effects on the housing and male labor markets, making labor force participation a relatively more important margin of adjustment in these cases.

In order to test the model's predictions I use data from Brazil during the 1991-2010 period. Using individual microdata from four editions of the population census I generate regional aggregates for 539 local labor markets with time-consistent boundaries. To measure exogenous shocks in gender-specific labor demand for each local labor market I introduce a variation of a well-known instrument proposed by Bartik (1991). My measure interacts national industry employment growth for each gender with local industry employment shares. It predicts what growth in a region's gender-specific employment would have been if the local industry shares had remained the same as in the starting year and if gender-specific employment had grown in local firms at the same rate as in same-industry firms in the rest of the country. I go beyond prior studies that use similar "shift-share" instruments assessing the plausibility of the identifying assumptions in the data, and adapting my empirical specifications to address potential identification concerns.

I find strong empirical support for the prediction that households migrate more in response to male than to female demand shocks. Male demand shocks increase the migrant population significantly more than female shocks. Joint mobility frictions seem to play an important role. Women migrate more in response to male demand shocks than to shifts in their own labor demand. The same is less true for men, whose migration responses to changes in female labor demand are much smaller. Consistent with these differential effects on population, I also find that male local demand shocks lead to growth in housing rents, and females shocks do not.

Turning to gender-specific labor market outcomes, I find that increases in male labor demand have a larger effect on own-gender employment and wages than equivalent changes in female labor demand. These differences are concentrated in the population without high-school education. In the context of the model, these findings suggests that male labor supply is more elastic than female labor supply – largely because of larger migration

responses – and that nominal wages partly reflect compensating differentials for increases in the cost of living, which are larger after positive male shocks.

The effects of shocks to the other gender's local labor demand are generally consistent with the differential migration mechanism playing an important role, but also highlight the importance of other margins of adjustment. Male shocks increased the population of employed and of non-employed females, which is in line with the presence of male-led joint migration in which tied-migrant women find disproportionately less work opportunities. However, in spite of shifting female labor supply rightwards, male shocks had a positive – though marginally significant – effect on local female wages. This could be explained, in the context of the model, by large compensating differentials for housing rents increases. In practice, it could also be driven in by family income effects on female labor supply or by changes in the skills composition of female labor, which the framework does not consider.

In the aggregate, male demand shocks increased the employment and wage gender gaps both in the 1990s and the 2000s. Meanwhile, female demand shocks reduced the gender employment gap but not the wage gap. Together, the empirical results point to two important areas for future research: the household's labor force participation decisions and the interactions between gender and education in the context of joint migration across local labor markets.

The asymmetric response of male and female labor markets to demand shocks translates into asymmetric welfare effects. While male-leaning local demand shocks are more likely benefit immigrants and landlords, female-leaning shocks are more likely to favor incumbent residents. Higher demand for female workers implies higher employment for residents because firms tap proportionally more into a labor pool already present in the region. Moreover, the smaller immigration effect limits the pressure on local housing prices, and workers are likely to receive a larger fraction of the benefits than landlords (Moretti 2011). In contrast, because higher demand for male workers leads to a larger migratory response and increase in housing demand, migrant workers and landlords receive a larger share of the economic rents.

The results also have implications for regional development policies. Initiatives that seek to create jobs and boost growth in underdeveloped regions are very popular around the globe. My findings suggest that such policies can have very different effects depending on the gender composition of the jobs that are created, and on the initial levels of male and female employment. Benefits to local populations can quickly dissipate through migration and higher local costs of living if job creation favors men.

The contributions of this paper are relevant for several literatures. First, it relates to studies of the effects of labor demand shocks on local economic outcomes including Diamond (2016), Amior and Manning (2015), Bartik (2015), Beaudry *et al.* (2014), Notowidigdo (2013), Glaeser *et al.* (2005) and earlier work by Blanchard and Katz (1992), Bartik (1991) and Topel (1986). This literature either looks at outcomes for the labor force as a whole or by skill group, aggregating male and female workers, or restricting the analysis to males. Moreover, in these works the marginal migrant is assumed to be an individual. By introducing realistic yet tractable new assumptions – a joint location constraint for married couples, and a higher opportunity cost of participating in the labor force for females – my paper shows that local labor demand shocks of equivalent size can lead to very different population, rents, employment and wage outcomes depending on whether they favor male or female jobs.

This paper also contributes to the literature on the efficiency and welfare consequences of place-based policies including Kline and Moretti (2014a), Busso *et al.* (2013), Kline (2010) and Glaeser and Gottlieb (2008) among others;¹ and related work which focuses on the extent to which immigrants – as opposed to local residents – benefit from local demand shocks and has produced contradictory results (Partridge *et al.* 2009; Bartik 2004). My work shows that the gender composition of the shocks can play an important role in determining the workers-landlords and residents-immigrants splits of welfare effects.

A closely related series of studies looks at the effects of trade shocks on local economic outcomes (Dix-Carneiro and Kovak 2017; Acemoglu *et al.* 2016; Costa *et al.* 2016; Hakobyan and McLaren 2016; Carneiro and Kovak 2015; Autor *et al.* 2013; Kovak 2013; Edmonds

¹See Neumark and Simpson 2015 for a review.

et al. 2010; among others). Trade shocks are likely to have different effects in the labor demand for men and women because local exposure depends on the industry composition of places and industries differ in the gender composition of employment. In this context, my findings suggest that gender asymmetries in the labor market could help explain regional heterogeneity in the effects of changes in import competition and export demand.

Lastly, my work contributes to the literature on the the gender wage and employment gaps (Blau 2016; Goldin 2014; Bertrand *et al.* 2010; Goldin 2006; Albrecht *et al.* 2003; Blau and Kahn 2003, 2000; Altonji and Blank 1999; Galor and Weil 1996; Lazear and Rosen 1990; among many others)² by showing how they can be exacerbated by tied migration in the context of local labor and housing markets. Male biases in labor demand, in addition to increasing the gender gap through higher male wages and employment, can also increase the relative abundance of female labor and push female wages downwards.

The remainder of the paper is organized as follows. Section 2.2 presents the model and discusses some of its predictions. Section 2.3 describes the data used in the analysis, presents relevant descriptive facts, and describes the identification strategy. Section 2.4 presents and discusses the empirical results. Section 2.5 discusses welfare and policy implications, and Section 3.6 concludes.

2.2 Spatial equilibrium with gender-segmented labor markets

In this section I develop a spatial equilibrium model that illustrates how labor demand shocks can affect employment, wages, population, and labor force participation for men and women. It incorporates the standard elements from the seminal Roback (1982) framework where local wages, housing rents, and amenities determine the geographic sorting of workers and the utility of the marginal migrant is equalized across space in equilibrium. This kind of model has been extensively used in urban economics to study the effects of local labor demand shocks in the U.S. and other high-income countries, but its use in

²See Ponthieux and Meurs 2015 for a review.

less-developed countries has not been as extensive (see Alves 2016, Morten and Oliveira 2016, and Oliveira and Pereda 2015 for recent applications in the context of Brazil).³ It has significant advantages relative to partial equilibrium approaches because it captures how aggregate labor outcomes are shaped both by the direct effects of the shock and by the endogenous adjustments of factors prices and quantities (Moretti 2011).

I depart from the standard model by incorporating segmentation by gender in the labor market and joint mobility constraints for married households.⁴ In my model the population consists of N married households indexed i , each with two members, a woman (W) and a man (M). There are a total of J regions indexed by j . Years are indexed by t . Each individual is endowed with one unit of labor. On the demand side, male and female workers are employed in different industries, each producing an intermediate good. Intermediate goods are ultimately aggregated as imperfect substitutes to produce a nationally-traded good. Each industry has its own productivity shifter, which is subject to exogenous shocks.

Individuals who sort into paid work incur a labor force participation cost φ_i , which is an exogenous stochastic variable with distribution $F(\varphi_i)$. This cost may reflect commuting (Black *et al.* 2014), childcare (Baker *et al.* 2008; Paes de Barros *et al.* 2011), or household appliances (Greenwood *et al.* 2005), among others. A key assumption of my model is that the distribution of this cost is gender-specific, such that the support starts at a value that is higher for women than for men by T_t . This is meant to reflect extensive work documenting a higher opportunity costs of labor force participation for females (Ponthieux and Meurs 2015). In this paper I abstract from the specific mechanisms that may drive this difference, and focus on its local labor market consequences.⁵

³In some cases the use of this framework may not be appropriate (Gollin *et al.*, 2017). Chapter 1 argues that in India, where geographic mobility is low and human capital heterogeneity extreme, a spatial equilibrium may not develop. However, the strong correlation between local wages and housing rents in Brazil and its higher internal mobility support the adequacy of the framework in this context.

⁴The notion that joint mobility constraints can partly explain gender differences in labor market outcomes has previously been previously explored by Gemici (2011) and Frank (1978) in the context of partial-equilibrium search models. To the best of my knowledge, this is the first paper to incorporate this constraint into a general spatial equilibrium framework.

⁵Similar simplifications can be found in Albanesi and Sahin (2017) and Garibaldi and Wasmer (2005).

Households observe local wages, housing rents and amenities, but they only learn their labor force participation costs after choosing a location. However, they know in advance $F(\varphi_i)$, and consequently their expected labor income net of participation costs in each region. After choosing a location, individuals decide whether to sort into the workplace or into domestic production by comparing wage income and the cost of participating in the labor force. For simplicity, I assume away unemployment in the model.

2.2.1 Production and labor demand

I assume that males and females sort into different industries, Y_G for $G = \{M, W\}$, where their labor is combined with traded capital K and non-traded capital \bar{Z}_j ⁶ to produce an intermediate good with the production function:⁷

$$Y_{Gjt} = \psi_{Gjt} N_{Gjt}^\beta K_{jt}^\gamma \bar{Z}_j^{1-\beta-\gamma}$$

Intermediate goods are combined with constant elasticity of substitution into a nationally-traded generic final good priced one, according to:

$$\begin{aligned} Y_{jt} &= \left(Y_{Wjt}^\sigma + Y_{Mjt}^\sigma \right)^{\frac{1}{\sigma}}, \text{ or} \\ Y_{jt} &= \left[\left(\psi_{Wjt} N_{Wjt}^\beta \right)^\sigma + \left(\psi_{Mjt} N_{Mjt}^\beta \right)^\sigma \right]^{\frac{1}{\sigma}} K_{jt}^\gamma \bar{Z}_j^{1-\beta-\gamma}. \end{aligned} \quad (2.1)$$

I assume that $0 \leq \sigma \leq 1$, which is equivalent to assuming that male and female effective labor are imperfectly substitutable factors of production with elasticity of substitution $\frac{1}{1-\sigma}$. This assumption is consistent with international empirical evidence (Olivetti and Petrongolo 2014; Johnson and Keane 2013; Acemoglu *et al.* 2004). Traded capital can be purchased in any amount at price one. The firms' problem is:

$$\max_{N_{Mjt}, N_{Wjt}, K_{jt}} \left\{ \left[\left(\psi_{Wjt} N_{Wjt}^\beta \right)^\sigma + \left(\psi_{Mjt} N_{Mjt}^\beta \right)^\sigma \right]^{\frac{1}{\sigma}} K_{jt}^\gamma \bar{Z}_j^{1-\beta-\gamma} - W_{Wjt} N_{Wjt} - W_{Mjt} N_{Mjt} - K_{jt} \right\}$$

⁶Non-traded capital is added to the production function so that there can be constant returns to scale at the firm level but decreasing returns to scale at the region level. Under these conditions it is possible to have a zero-profit condition for firms and a finite size of regions (Glaeser, 2008).

⁷I assume that regions have many homogeneous firms, so that the region-level production function is the same as the firm's. Individual firms' indexes are omitted for simplicity.

The solution yields the labor demand equations:

$$W_{Gjt} = \beta \gamma^{\frac{\gamma}{1-\gamma}} \psi_{Gjt}^{\sigma} N_{Gjt}^{\beta\sigma-1} L_{jt}^{\frac{\gamma(1-\sigma)}{(1-\gamma)}} \bar{Z}^{\frac{1-\beta-\gamma}{1-\gamma}} \quad (2.2)$$

$$L_{jt} = \left[\left(\psi_{Wjt} N_{Wjt}^{\beta} \right)^{\sigma} + \left(\psi_{Mjt} N_{Mjt}^{\beta} \right)^{\sigma} \right]^{\frac{1}{\sigma}}$$

This formulation provides insights about the effects of gender-specific productivity shocks on the local wage gap, which is given by:

$$\frac{W_{Mjt}}{W_{Wjt}} = \left(\frac{\psi_{Mjt}}{\psi_{Wjt}} \right)^{\sigma} \left(\frac{N_{Mjt}}{N_{Wjt}} \right)^{\beta\sigma-1} \quad (2.3)$$

Equation 2.3 shows that the local gender wage gap depends on the gender productivity difference, the degree of substitutability of male and female labor, and on the relative abundance of male and female workers. The direct effect of gender-specific shocks on the gap will be positive in the case of males ($\frac{\partial(w_{Mjt}/w_{Wjt})}{\partial\psi_{Mjt}} > 0$), and negative in the case of females ($\frac{\partial(w_{Mjt}/w_{Wjt})}{\partial\psi_{Wjt}} < 0$). However the total effect depends on how changes in male and female productivity affect the ratio of male to female workers. For example, if male productivity shocks generate larger migratory responses and make male labor relatively more abundant than female labor ($\frac{\partial(N_{Mjt}/N_{Wjt})}{\partial\psi_{Mjt}} > 0$), they would also have a negative partial effect on the wage gap under the imperfect substitutes assumption ($0 \geq \sigma \geq 1$). Likewise, increases in female productivity could worsen the wage gap if female migration effects are large enough to topple the wage productivity premium.

2.2.2 Household utility

Households chose locations to optimize a joint Cobb-Douglas utility function. As is standard in spatial equilibrium models following Roback (1982), they collectively derive utility from the consumption of a composite tradable good C_{ijt} priced one, housing rented at R_{jt} ,⁸ and a

⁸For simplicity, I do not include home production or leisure in this version of the the utility function. This helps to highlight the role gender-asymmetric migration responses in the model, at the cost of assuming away income effects. In a forthcoming extension of the model I introduce household production with positive utility value.

local amenities index θ_j , which I assume to be exogenous and time-invariant for simplicity.⁹

The household optimization problem is thus given by:

$$\max_{C_{ijt}, H_{ijt}} \left\{ \theta_j C_{ijt}^{1-\alpha} H_{ijt}^\alpha \right\} \text{ s.t. } W_{ijt}^{net} = C_{ijt} + R_{ijt} H_{ijt}$$

where $W_{ijt}^{net} = W_{Mijt}^{net} + W_{Wijt}^{net}$ is the household-level net labor income, and

$$W_{Gjt}^{net} = \begin{cases} W_{Gjt} - \varphi_{it} & \text{if the person sorts into the workforce} \\ 0 & \text{if the person does not} \end{cases}$$

The optimized housing consumption is therefore:

$$H_{ijt}^* = \alpha \frac{W_{ijt}^{net}}{R_{jt}} \quad (2.4)$$

Substituting the budget constraint and the optimal housing consumption into the utility function, one can express indirect utility of household i living in region j at time t as:

$$V_{ijt}(\theta_j, W_{ijt}^{net}, R_{jt}) = \alpha^\alpha (1 - \alpha)^{1-\alpha} \theta_j W_{ijt}^{net} R_{jt}^{-\alpha} \quad (2.5)$$

The spatial equilibrium assumption implies that the indirect utility is equalized across space for the marginal household, $V_{ijt}(\theta_j, W_{ijt}^{net}, R_{jt}) = \underline{U}$. Note that restricting the choice to a single location entails that household utility may be smaller than in the standard spatial equilibrium framework, where individuals are able to choose location separately. If the individual spatial equilibrium utilities for men and women are \underline{U}_M and \underline{U}_W , and the optimal combination of wages, rents and amenities are not in the same geographical region for both of them, then introducing a joint location constraint implies that at least one of the members of the household may reside in a sub-optimal location where $V_{Gijt} < \underline{U}_G$, implying $\underline{U}_{ijt} \leq \underline{U}_M + \underline{U}_W$.

⁹An emerging literature has shown the importance of endogenous amenities in shaping local economic outcomes, including Albouy and Stuart (2017); Lee and Lin (2017); Diamond (2016); and Hanlon (2015). Because in my model couples choose a single location, endogenous amenities are unlikely to be a first-order determinant of differential responses of male and female labor markets to demand shocks. They could however, affect the gender welfare gap if male and female workers differ in their preferences over amenities.

2.2.3 Labor force participation

Individuals have an exogenous and stochastic labor force participation cost, which they draw after moving to a new region from a power law with CDF $F(\varphi_i) = \left(\frac{\varphi_i}{\varphi_{min}}\right)^\iota$, $\iota \in [0, 1]$ and support $\varphi_i \in (1, \varphi_{max})$ for men and $\varphi_i \in (1 + T_t, \varphi_{max})$ for women. Individuals sort into the workplace if their wage is weakly greater than their participation cost. This implies that the participation costs that make men and women indifferent are $\varphi_{Gjt}^* = W_{Gjt}$. The female labor supply is therefore given by $N_{Wjt} = N_{jt} \left(\frac{W_{Wjt}}{1+T_t}\right)^\iota$. The implied inverse labor supply function is:

$$W_{Wjt} = (1 + T_t) \left(\frac{N_{Wjt}}{N_{jt}}\right)^{\frac{1}{\iota}} \quad (2.6)$$

Conversely, male labor supply is $N_{Mjt} = N_{jt} W_{Mjt}^\iota$, which corresponds to the inverse supply function:

$$W_{Mjt} = \left(\frac{N_{Mjt}}{N_{jt}}\right)^{\frac{1}{\iota}} \quad (2.7)$$

2.2.4 The housing market

Housing belongs to absentee landlords, who buy it from developers and rent it to local residents at R_{jt} . Profits for developers, are given by:

$$\pi_{jt} = \sum_{t=0}^{\infty} \frac{R_{jt}}{(1+r_t)^t} - CC_{jt}$$

where r_t is the national interest rate, and CC_{jt} are the local construction costs.¹⁰

There is free entry and the zero-profit condition holds, so that developers sell housing at the cost of construction, $\frac{(1+r_t)}{r_t} R_t = CC_{jt}$. For a given construction cost, there is a supply of $\bar{H} \cdot CC_{jt}^\rho$ units of housing, that is, additional units can be provided at higher construction costs with elasticity ρ . This implies that the local housing supply is given by:

$$\bar{H} \left(\frac{1+r_t}{r_t}\right)^\rho R_{jt}^\rho \quad (2.8)$$

Local housing demand is the aggregate from all N_{jt} households locating in region j at

¹⁰The housing supply component of my model follows closely Glaeser (2008).

time t . Based on equation 2.4, it can be written as:

$$\begin{aligned} H_{jt} &= \alpha \frac{\bar{W}_{jt}^{net}}{R_{jt}} N_{jt} \\ \bar{W}_{jt}^{net} &= \left(\frac{N_{Mjt}}{N_{jt}} W_{Mjt} - \bar{\varphi}_{Mjt} \right) + \left(\frac{N_{Wjt}}{N_{jt}} W_{Wjt} - \bar{\varphi}_{Wjt} \right) \end{aligned}$$

and $\bar{\varphi}_{Gjt}$ is the average participation cost for individuals of gender $G = \{M, W\}$ that sort into the workforce in region j .

In equilibrium demand and supply for housing equate, yielding the rent equation:

$$R_{jt}^* = \left(\alpha \frac{\bar{W}_{jt}^{net}}{\bar{H} \left(\frac{1+r_t}{r_t} \right)^\rho} N_{jt} \right)^{\frac{1}{1+\rho}} \quad (2.9)$$

2.2.5 Key insights and predictions

In this section, I describe the key insights and predictions provided by the model's analytical solution. Appendix B.3 describes how I close the model, and gives greater detail about the resulting expressions.

Equation 2.9 allows me to re-write the indirect utility of households living in region j (equation 2.5) only in terms of the expected net household wage, local amenity levels, the city population and exogenous parameters. The net household wage enters the utility function as an expectation because, before migration, there is uncertainty about the individuals' participation costs. Under the spatial equilibrium assumption utility is equalized for the marginal migrant household making them indifferent with across locations, $V_{jt}(\theta_j, E(W_{jt}^{net}), N_{jt}) = \underline{U}$. The spatial indifference curve can be used to express the local population in terms of the expected net household wage:

$$N_j = \left[E(W_{jt}^{net}) \right]^{\frac{1+\rho-\alpha}{\alpha}} \left(\frac{\zeta \theta_j}{\underline{U}} \right)^{\frac{1+\rho}{\alpha}} \quad (2.10)$$

where $\zeta := \frac{\alpha}{\bar{H} \left(\frac{1+r_t}{r_t} \right)^\rho}$ and $E(W_{jt}^{net}) = E(W_{Mjt}^{net}) + E(W_{Wjt}^{net})$. In turn, the gender-specific

expected net wage is given by:

$$E \left(W_{Wj}^{net} \right) = \left(\frac{W_{Wjt}}{1 + T_t} \right)^\iota \left[W_{Wjt} - \frac{\iota (1 + T_t)}{\iota + 1} \left(\left(\frac{W_{Wjt}}{1 + T_t} \right)^{\iota+1} - 1 \right) \right] \quad (2.11)$$

$$E \left(W_{Mj}^{net} \right) = W_{Mjt}^\iota \left[W_{Mjt} - \frac{\iota}{\iota + 1} \left(W_{Mjt}^{\iota+1} - 1 \right) \right] \quad (2.12)$$

where the probabilities of participating and the expected costs of participation for each gender follow from the functional form assumption on $F(\varphi_i)$ (see model solutions in Appendix B.3 for details).

Relative effects of male and female demand shocks on population and rents

I am interested in comparing the effects of equivalent shocks to the productivity to the female-intensive industry ($\Delta\psi_{Wjt}$) and the male-intensive industry ($\Delta\psi_{Mjt}$) in region j , which correspond to shifts in female and male local labor demand respectively. From the labor demand equation 2.2 it is apparent that the partial effect on gender-specific wages is positive ($\partial W_{Gjt}/\partial \psi_{Gjt} > 0$) and its size is mediated by the elasticity of substitution of male and female labor (captured by σ).

The effects of gender-specific demand shocks on migration and ultimately population will in turn depend on how migrants react to changes in the expected male and female wage. Equation 2.11 shows that, in expectation, the contribution of the female wage to the household labor income is penalized by their incremental cost of participating in the labor force, T_t . The same is not true for the expected male wage in equation 2.12. It follows that demand shocks that affect the wages for males will have a larger impact on population – through migratory adjustments – than equivalent shocks affecting female wages.

Because a larger population increases housing demand and pushes the equilibrium rent up (see equation 2.9), shocks to labor demand for males – compared to equivalent shocks for female labor – will also have a larger effect on housing rents.

Effects of male and female shocks on employment

The equilibrium under autarky, which treats the regions' population as exogenous, is useful to provide intuition of the predictions of the model and the role played by migrant households constrained to choosing a single location. In the absence of migratory adjustments, the equilibrium female and male employment in region j are given respectively by:

$$N_{Wjt}^{*aut} = N_{jt}^{\frac{1}{1-\iota\xi_1}} \left(\frac{\lambda_1}{1+T_t} \right)^{\frac{\iota}{1-\iota\xi_1}} \psi_{Wjt}^{\frac{\iota\sigma}{1-\iota\xi_1}} \Psi_{Wjt}^{\frac{\gamma(1-\sigma)}{(1-\gamma)} \frac{\iota}{1-\iota\xi_1}} \quad (2.13)$$

$$\Psi_{Wjt} = \left[\psi_{Wjt}^\sigma + \psi_{Mjt}^\sigma \left[(1+T_t) \left(\frac{\psi_{Mjt}}{\psi_{Wjt}} \right)^\sigma \right]^{\frac{\beta\iota\sigma}{1-\iota(\beta\sigma-1)}} \right]^{\frac{1}{\sigma}}, \text{ and}$$

$$N_{Mjt}^{*aut} = N_{jt}^{\frac{1}{1-\iota\xi_1}} \lambda_1^{\frac{\iota}{1-\iota\xi_1}} \psi_{Mjt}^{\frac{\iota\sigma}{1-\iota\xi_1}} \Psi_{Mjt}^{\frac{\gamma(1-\sigma)}{(1-\gamma)} \frac{\iota}{1-\iota\xi_1}} \quad (2.14)$$

$$\Psi_{Mjt} = \left[\psi_{Wjt}^\sigma \left[(1+T_t) \left(\frac{\psi_{Mjt}}{\psi_{Wjt}} \right)^\sigma \right]^{\frac{\beta\iota\sigma}{\iota(\beta\sigma-1)-1}} + \psi_{Mjt}^\sigma \right]^{\frac{1}{\sigma}}$$

with constants $\lambda_1 := \beta\gamma^{\frac{\gamma}{1-\gamma}} \bar{Z}^{\frac{1-\beta-\gamma}{1-\gamma}}$, and $\xi_1 := \frac{\beta\gamma(1-\sigma)+(1-\gamma)(\beta\sigma-1)}{(1-\gamma)}$ (see Appendix B.3 for details).

These equations show that the direct effect of shocks to own-gender labor demand on employment are positive for both men and women. The gender-specific industry productivity terms ψ_{Gjt} in equations 2.13 and 2.14 increase employment directly ($\frac{\partial \psi_{Gjt}}{\partial \psi_{Gjt}} > 0$) and dominate the substitution effect captured by the term Ψ_{Gjt} (that is, $\frac{\partial \Psi_{Gjt}}{\partial \psi_{Gjt}} > 0$). The effect, however, is larger for the case of males than females, reflecting the latter's larger labor force participation costs. While $1+T_t$ effectively scales down the constant λ_1 in equation 2.13, it does not have a similar effect in equation 2.14.¹¹ The larger employment effects of male shocks on own-employment are exacerbated in the open-region equilibrium, where population is endogenous. This is because, as discussed earlier, the own-gender migration effect is larger for men than for women. However, increases in housing rents acts as a

¹¹The term $1+T_t$ also enters the input substitution terms Ψ_{Gjt} , but its effect on the male and the female case is symmetrical.

counterbalancing force, deterring migration more in the case of male than of female shocks.

The effects of other-gender labor demand shocks on employment are also positive. In the absence of migration, this is driven primarily by the input substitution effect captured in Ψ_{Gjt} , and they are symmetric for both genders. With migration, however, an asymmetry arises. Male shocks have a disproportionately larger effect on female employment because of its larger effect on population N_{jt} .

Effects of male and female shocks on wages

In the autarkic equilibrium equilibrium female and male wages are given by:

$$W_{Wjt}^{*aut} = N_{jt}^{\frac{\xi_1}{(1-\iota\xi_1)}} \left(\frac{\lambda_1}{1+T_t} \right)^{\frac{1}{1-\iota\xi_1}} \psi_{Wjt}^{\frac{\sigma}{1-\iota\xi_1}} \Psi_{Wjt}^{\frac{\gamma(1-\sigma)}{(1-\gamma)} \frac{1}{1-\iota\xi_1}} \quad (2.15)$$

$$W_{Mjt}^{*aut} = N_{jt}^{\frac{\xi_1}{(1-\iota\xi_1)}} \lambda_1^{\frac{1}{1-\iota\xi_1}} \psi_{Mjt}^{\frac{\sigma}{1-\iota\xi_1}} \Psi_{Mjt}^{\frac{\gamma(1-\sigma)}{(1-\gamma)} \frac{1}{1-\iota\xi_1}} \quad (2.16)$$

Without household migration, the direct effect of shocks to own-gender labor demand on wages are positive for both genders and smaller for women. They enter the equation in the same structure as they do as in the employment equation, although the relative role of the input substitution term is larger.¹² The population term, however, enters the equations negatively ($\xi_1 < 0$). If the region is open to migration, the inflow of immigrants shifts the labor supply rightwards pushing wages down, but the effect is mitigated by the subsequent increase in housing rents, which deters further migration and induces firms to pay a compensating differential if they want to attract more labor. In the open region the effects on own-gender wages can be smaller for men than for women if the downward effect coming from migration, which favors female wages, dominates the differential penalty for participation costs T_t and the wage compensation for higher costs of living, which favor male wages. The net prediction on the effects of demand shocks on own-gender wages in ambiguous.

In the absence of migration, the effects on wages of other-gender labor demand shocks are

¹²The direct effect has a smaller exponent given that $\sigma > \iota\sigma$, and the input substitution component a larger exponent given that $1 > \iota$.

also positive and symmetric for both genders, and are driven entirely by input substitution. Migration introduces a negative effect of other-gender shocks because population enters the equation negatively and couples move together. And because migration responds more to male shocks, the net effect of these shocks on female wages can be negative, unless the compensating differentials for higher housing rents are high.

In sum, the model delivers clear predictions on the effects of gender-specific demand shocks on local population, housing rents, and gender-specific employment; which are all positive and larger for male than for female shocks. The predictions of the model are ambiguous with respect to wages, and because of the role of wages in the decision to sort into the workforce (equations 2.6 and 2.7), they are also ambiguous with respect to participation rates. With this framework in mind, I turn now to the empirics.

2.3 Data, Facts and Identification Strategy

In this section, I describe the data and characterize a few features of Brazilian local labor markets over the period of interest to provide context for the analysis. I also present the identification strategy, discuss its key assumptions, and address potential identification concerns.

2.3.1 Data

The data used in this analysis comes primarily from the decennial population censuses of 1980 through 2010. The Brazilian Institute for Geography and Statistics (IBGE) makes available to researchers the microdata for the long-form questionnaire sample, which corresponds to 10% of the population in 1980 and 5% in the subsequent census years. I complement this with data from other sources, including municipality areas and climate data from the Brazilian Institute of Applied Economic Research (IPEA), and GIS data from IBGE. Details of the sources and definition of the variables used in the analysis are included in the Data Appendix D. Appendix tables B.1 and B.2 present summary statistics of the main regional variables for the 1990s and the 2000s, respectively, and appendix tables B.3

and B.4 report correlations among these variables.

The definition of local labor markets used in the main specifications of the analysis is a Brazilian “microregion”. Microregions are defined by the IBGE as groupings of contiguous and economically integrated municipalities (IBGE 2002), and a growing literature acknowledges them as good approximations of the boundaries of local labor markets, and uses them in regional research (Costa *et al.* 2016; Dix-carneiro and Kovak 2016; Adão 2015; Kovak 2013).

In order to be able to compare microregions across time, it is necessary to adjust for the changes in administrative boundaries. The number of Brazilian municipalities grew dramatically over this period, going from 3,992 in 1980 to 4,491 in 1991 and to 5,565 in 2010. In a number of cases the parents of newly-created municipalities belonged to different microregions. I create time-consistent boundaries by aggregating the original the IBGE microregions that share the same family tree over this period, as in Kovak (2013). The resulting sample includes 539 regions.¹³ I use the microdata to generate regional-level aggregate measures for the different subsamples of interest (see Appendix D).

2.3.2 Descriptive facts

Brazil has been for many years among the countries with the highest economic inequality in the world. In 1991, the income share held by the top decile was 48.1%, much higher than in other large economies like India (26%), China (25.3%) or the US (26.7%) (see Chapter 1).

These disparities have major geographic and gender components. Economic opportunities are very unequally distributed over the national territory, especially across the poorer North and North East regions, and the richer South and South East regions. Figure 2.1

¹³The number of time-consistent microregions is significantly larger than that generated by Kovak (2013). This is because this paper uses as an input time-consistent municipalities (a.k.a. “Minimum Comparable Areas” – MCAs) originally produced by Reis *et al.* (2007). In this database MCAs are more aggregated than needed for accurate comparisons over the period of interest. I first re-create MCAs using the official municipalities’ family trees made available by the IBGE, and then generate time-consistent microregions using the new MCAs as input (see Data Appendix D). My empirical results are largely unchanged when I use the time-consistent microregions from Kovak (2013) to assess robustness, but in that case they are measured with less precision than in my main sample.

provides a stark illustration. The left panel contrasts the distribution of average labor income across Brazilian microregions with the same distribution across metropolitan statistical areas in the U.S.A. in the year 2000. While the average income of local labor markets in the U.S. follows a unimodal distribution, in Brazil the distribution is bi-modal. The left panel shows that, underlying this unconventional shape, are large differences in labor income among the main geographies of the country.

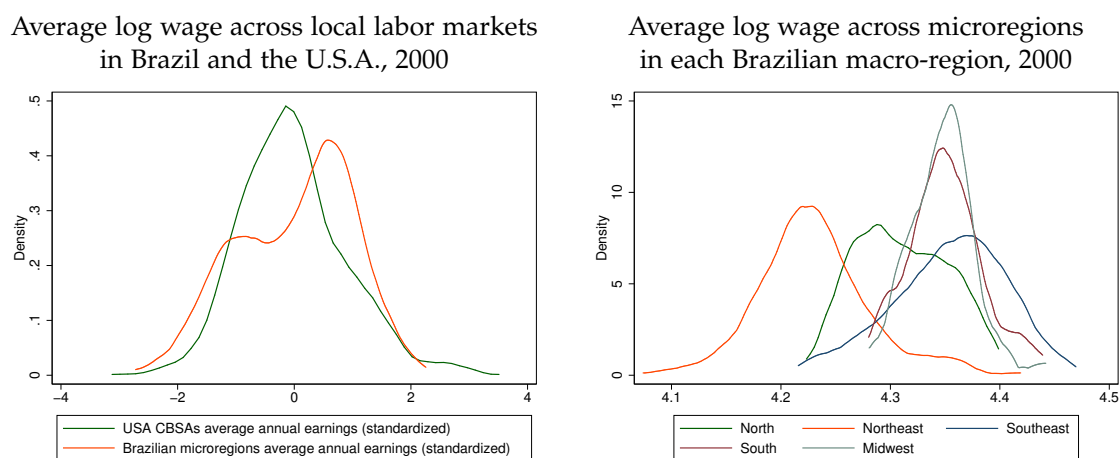


Figure 2.1: *Distribution of labor income across local labor markets in Brazil and the U.S.A.*

Inequality has also an important gender component, including a large gender wage gap and differences in labor force participation, work experience, and other correlates of labor productivity across men and women (Foguel 2016). Part of these differences are explained by the fact that, historically, females had less access to formal education. Appendix table ?? shows that in 1991 the fraction of the population not participating in the labor force was twice as large in the sample with less than high-school education than in the sample with high-school or or higher education diploma. But even among the group with more formal education, non-participation rates were much higher for women (31%) than for men (7%) at the beginning of the decade.

Gender and geographic dimensions of inequality appear to be closely intertwined. In the cross section, the participation gender gap is more acute at lower income levels. Appendix Figure B.2 shows the distribution of labor force participation across local labor

markets for the five Brazilian macro-region by gender and education group.¹⁴ Local labor markets in poorer areas tend to have lower participation rates than in richer areas. This geographic differences are significantly less pronounced among the population with high-school education.

The country had very different macroeconomic performance across the two decades covered in this study. While 1990-2000 was a period characterized by volatility and rising unemployment, 2000 to 2010 were years of consistent growth and improving economic opportunities, particularly for the lower-income population. The 1990s started with a sharp reduction of trade barriers and a major push for the privatization of state-owned enterprises. Hyperinflation, which had severely threatened the livelihood of millions of Brazilians during the 1980s and the early 1990s, was brought to a halt in 1994 by a series of economic measures known as the “Plano Real”, and a fairly stable period ensued during the second part of the decade. This stability, however, was not enough to prevent a massive jobs loss, and by the end of the decade unemployment had increased by 11 percentage points relative to 1991 levels (Appendix Table B.5). In contrast, the 2000s saw important GDP and employment growth, accompanied by progress in inequality reduction, which was driven both by a compression in the distribution of labor income, and by the expansion of transfers to low-income families (De Barros *et al.* 2006).

The sharp differences between these two decades can be seen in Figure 2.2, which shows the national employment growth by industry and gender over these two periods. There were only a handful of industries that did not see net job loss during the 1990s. Employment was lost in primary industries, manufacturing and services. Job loss did not systematically hurt men or women across industries: while sectors like agriculture, textiles manufacturing and financial services saw a disproportionate loss in male employment, female employment was more affected in other sectors like mineral mining, rubber product manufacturing and utilities.

¹⁴Macroregions are the coarser statistical division in the country, roughly equivalent to U.S. census regions.



Figure 2.2: *Employment growth by industry and gender, Brazil 1991-2010*

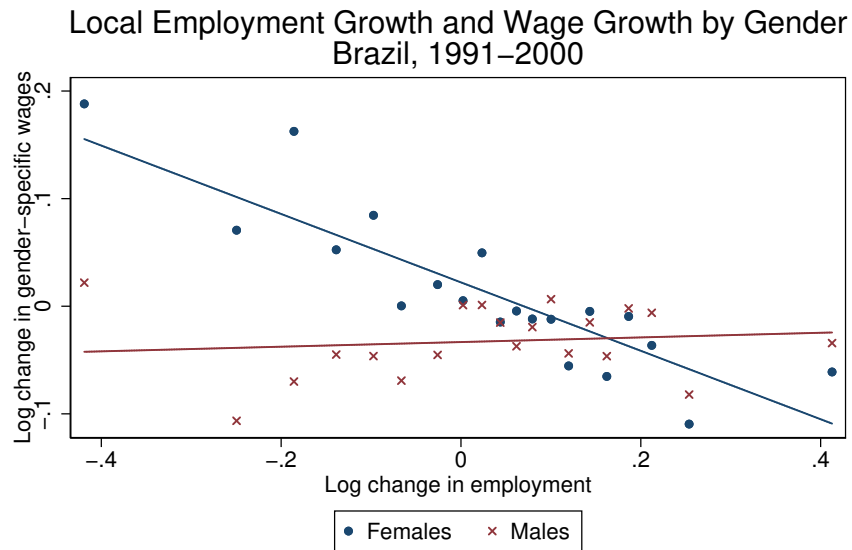
Note: Own calculations using census data. **Sources:** See data appendix.

During the 2000s, in contrast, most industries grew. Females saw a larger proportional growth in most industries, reflecting in part lower starting employment levels. But the relative growth of male and female jobs was heterogeneous across industries. My empirical strategy leverages these national gender differences to measure changes in male and female labor demand that are plausibly exogenous at the local level.

The increasing economic opportunities and shrinking inequality in the 2000s brought along a reduction of internal migration. In the 2000 census 17.41% of the population had been living in a different microregion ten years before. That number came down to 10.35% in the 2010 census. This reduction was driven by the subpopulation with lower levels of education (see Table B.6 for more on internal mobility). In terms of the framework presented in Section 2.2, this implies that asymmetric migratory responses to male and female demand shocks are likely to have played a less important role in determining local labor market outcomes in the 2000s than in the prior decade.

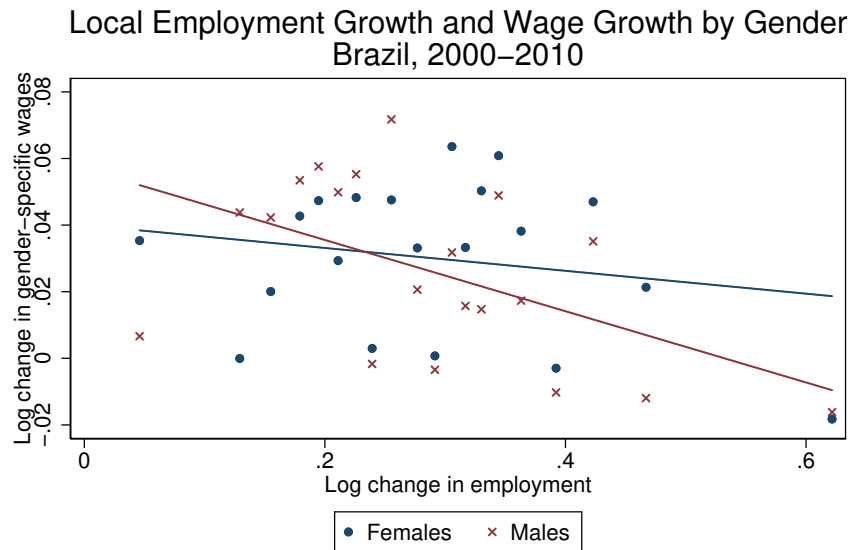
Differences in the relative role of migration may help us understand the correlation of gender wage gap and local employment growth in different points in time. The bin scatter plots in Figure 2.3 measure total employment growth (including males and females) on the horizontal axis, and gender-specific wage growth on the vertical axis. The left panel shows that in the 1990s the relationship between employment and wage growth varied significantly by gender. While Brazilian microregions that experienced employment growth also saw shrinking female wages in this decade, places that lost employment witnessed increases in wages for women. The same was not true for the relationship between the male wage and employment growth, which had a weak, positive correlation in that decade. In contrast, in the 2000s both male and female local wages decreased as employment raised. The slope of the regression line was larger for men, but the fit was much weaker than in the prior decade for both genders.

1990s



Note: Units of observation are microregions. Markers correspond to X and Y variable means in each of 20 equal-sized X-variable bins. Source: IBGE, Population Censuses of 1991 and 2000.

2000s



Note: Units of observation are microregions. Markers correspond to X and Y variable means in each of 20 equal-sized X-variable bins. Source: IBGE, Population Censuses of 1991 and 2000.

Figure 2.3: *Changes in Employment and the Gender Wage Gap*

Note: Own calculations using census data. **Sources:** See data appendix.

For the demand-side of the market to explain a pattern like the one observed in the 1990s one would have to assume that the production technology is such that the relative

demand for female labor drops in good times and increases in bad times. This would be consistent with the observed changes in female wages, but would fail to account for the relatively constant male wages. Moreover, it would not explain why the pattern changed in the following decade. The model presented in section 2.2 provides a potential supply-side explanation. If couples chose a single location and they are more responsive to male than to female job prospects, and if booms and busts in labor demand disproportionately affect men, migratory adjustments could account for the observed differences in wage growth in the 1990s. In turn, population growth and larger labor force participation of males could account for the patterns observed in the 2000s, when the migration margin was relatively less important.

2.3.3 Identification strategy

In this section I discuss the approach I use to empirically identify the effects of gender-specific changes in local labor demand on the labor and housing markets. Specifically, I am interested in measuring how migration, male and female wages and employment, and housing rents react to gender-specific labor demand shocks. The reduced-form relationship of interest for each of these outcomes is:

$$\Delta_{t-t_0} Outcome_j = \alpha + \beta_G \Delta_{t-t_0} Labor\ Demand_{jG} + \delta Controls_{j,t_0} + \Delta_{t-t_0} \epsilon_{jG} \quad (2.17)$$

where Δ_{t-t_0} denotes the log-change between the start year (t_0) and the end year (t) in region j ; subscript G denotes gender (M or W); and ϵ_{jG} is the error term.

Gender-specific Bartik shocks

In order to estimate the effect of changes in labor demand I need a measure of demand shifts that is independent from local labor supply characteristics. I introduce a variant of “shift-share” shocks, widely used in the literature studying local economies following Bartik (1991). I construct gender-specific Bartik shocks interacting the aggregate industry employment growth for each gender with each region’s start-year industry mix. Similar variations have been used in recent studies to instrument for changes in local female wages

(Bertrand *et al.* 2015; Aizer 2010). Specifically, I calculate:

$$Bartik_{jt}^G = \sum_{ind} \underbrace{\eta_{ind,j,t_0}}_{\text{Local industry shares at } t_0} \underbrace{\left(\log N_{ind,-j,t}^G - \log N_{ind,-j,t_0}^G \right)}_{\text{National change in gender G industry employment}} \quad (2.18)$$

where $N_{ind,-j,t}^G$ is the number of workers of subgroup $G = \{M, W\}$ employed in industry ind at time t nationally, excluding region j ; and η_{ind,j,t_0} is the share of employment of region j in industry ind at the start period (t_0). I use leave-one-out national employment growth, following Autor and Duggan (2003), to address concerns that the introduction of own-region employment may mechanically increase the predictive power of the shock. The gender-specific Bartik shocks in equation 2.18 predict what growth in a region's female (or male) employment would have been if the local industry shares had remained the same as in the starting year and gender-specific employment had grown in local firms at the same rate as in same-industry firms in the rest of the country. Appendix figure B.1 shows the distributions of male and female Bartik shocks for the two decades, and figure B.3 depicts the geographic distribution of these shocks.

Identifying assumptions

In spite of the widespread use of Bartik-style shocks, there was until recently little discussion on the ultimate source of the identifying variation. The standard identifying assumption is generally understood: conditional on controls, the shock should be uncorrelated with the error term ($\Delta_{t-t_0} \epsilon_{jG}$ in this context). However, the shock itself has two structural components, the local industry shares and the national industry growth rates, and it is a priori unclear which of them drives the exogenous variation.

Goldsmith-Pinkham *et al.* (2017) study this question, and conclude that identification in Bartik-style shocks comes exclusively from the local industry shares $\eta_{ind,j}$, while the national industry growth contributes only to predictive power. They show that using Bartik shocks in 2SLS estimation is numerically equivalent to using a GMM estimator where the weight matrix is constructed with the national growth rates, and the local industry shares alone are

used as the instrument.¹⁵

This implies that the underlying identifying assumption for the shock in equation 2.18 to produce causal estimates is that the vector of industry shares is uncorrelated with the decade-long changes in error term conditional on the set of controls. In the same study, the authors assess this assumption in the data in the context of existing studies that use Bartik shocks to recover the shape of the local labor supply curve. They show that local industry shares are typically correlated with observable characteristics of a place, particularly measures of education, and that estimates that do not control for these correlates may be biased.¹⁶ In addition, they find the existence of pre-trends, even after considering the mechanical autocorrelation of the Bartik shocks over time, highlighting the importance of controlling for lagged growth.

Addressing identification concerns

In order address these identification concerns, I start by implementing two tests suggested by Goldsmith-Pinkham *et al.* (2017). First, I regress the gender-specific Bartik shocks with a number of start-year microregions' characteristics, and find similarly strong correlations. The results, shown in Appendix Table B.8, show that education levels (as measured by the share of high-school educated in the adult population) is also a strong correlate of Bartik-style shocks in Brazil. They also reveal other correlates that may be specific to lower-income contexts, like urbanization rates and the demographic structure – the share of children and of prime-age adults in the population exhibit a strong connection with all the shocks.

Second, I assess the presence of pre-trends that could bias the estimates. In order to avoid capturing mechanical trends coming from the serial autocorrelation of the shocks,¹⁷ the test first obtains the residuals of a regression of gender-specific employment growth

¹⁵Preliminary work by Borusyak and Jaravel (2017) suggests that identification can in some cases also come from randomness in the national industry growth rates.

¹⁶Specifically, they find that IV estimates of the inverse elasticity of labor supply attenuate by over 25% after including base-year controls that are found to be correlated with the Bartik shock.

¹⁷Amior and Manning (2015) also highlight the role of serial correlation of demand shocks, showing that it can explain a large variation in local joblessness.

on the corresponding shock. It then regresses the Bartik shocks from one decade in the future on these residuals. I repeat the exercise using growth in wages as an outcome. If future shocks are able to predict the fraction of the lagged outcomes that is not explained by contemporary shocks, it is taken as evidence of the existence of pre-trends. The results of these tests are shown in Table B.9. I find no statistically-significant evidence of pre-trends in the 1990s, but strong evidence in the 2000s.

In order to address the concerns raised by the correlation with start-year variables and the presence of pre-trends, I include in all my regressions a set of base-year and lagged controls. Part of these controls are directly informed by the tests described above. The base-year controls include population density, average log wages, average log housing rents, share of adults with high-school education or higher, shares of the population in six different age groups to account for demographic differences across localities, urbanization rate, formal and informal employment shares in the population, unemployment rate, and winter temperatures as a proxy of climate amenities.¹⁸ The lagged-growth controls include the changes in the decade preceding the start year in population, wages, informal and formal employment, unemployment and urbanization rates. In my preferred specification I also use controls that seek to prevent comparing local economies that are structurally very dissimilar. These include employment shares of three broadly-defined industries: agriculture, manufacturing and government, and state fixed effects. My results reflect, therefore, comparisons of microregions within states, which have broadly similar industry structures.

The coefficients on the gender-specific Bartik shocks can be given a causal interpretation under the selection-on-observables assumption. While by definition one can't rule-out the presence of potential unobservable confounders, looking at the correlation with other regional characteristics not included in the set of controls can be informative. If after

¹⁸Temperature appears to be a good proxy for time-invariant amenities that affect location decisions in Brazil. Oliveira and Pereda (2015) find that non-agricultural workers in Brazil have high willingness to pay for more temperate climate, and Chapter 1 in this dissertation finds that housing rents are larger in places with better climate amenities.

controlling for the variables described above the shocks are still correlated with other start-year or lagged variables, identification would be put into question. I perform this exercise with a number of variables and find in all cases that the remaining variation of the shock is uncorrelated with characteristics not included as controls.

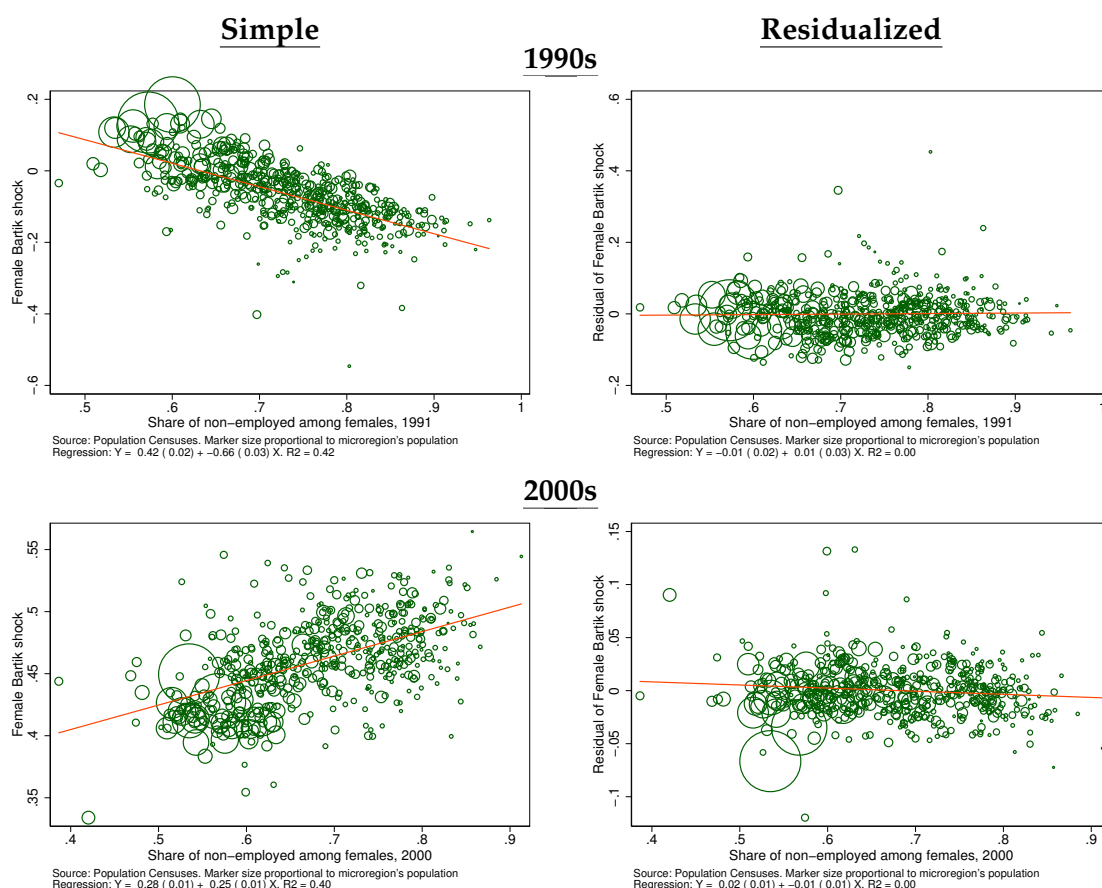


Figure 2.4: *Female Bartik shocks and base year female non-participation*

Note: Own calculations using census data. **Sources:** See data appendix.

Figure 2.4 presents one example of these tests. Here I compare the Bartik shocks before and after controls with the share of non-employed individuals (i.e. non-participant or unemployed) in the adult female population for both decades. While the left column uses the Bartik shock measures without modifications, the right column uses the residuals of a regressions of the shocks on the controls. The evidence suggests that the included variables

are effective at controlling for other potential confounders.¹⁹

2.4 Results

This section presents the empirical results of the paper. I estimate reduced-form regressions of the form described in equation 2.17, using decade-long changes for the 1990s and the 2000s. My units of observation are Brazilian microregions, and I use gender-specific Bartik shocks as measures of exogenous shifts in male and female labor demand.

I first look at the migration elasticity of households with respect to male and female labor demand shocks, and establish the existence of asymmetric responses. Second, I assess the effects of male and female shocks on housing rents, finding that male shocks lead to faster growth in local costs of living. Third, I evaluate the effects of male and female labor demand shocks on employment by gender, finding that male shocks tend to increase the employment gender gap and females shocks to decrease it. Fourth, I look at the net effect of gender-specific shocks on wages, finding that male shocks worsened the wage gap in both decades, and female shocks also worsened it during the 1990s. Finally, I look at the effects on the population of males and females that do not participate in the local labor force. I find that male shocks reduce non-participating male population, and increase non-participating female population. Conversely female shocks fail to reduce, or even have a positive effect the non-participant female population.

Both the shocks and the outcome variables are calculated for adults ages 15 through 64 who are not enrolled in an educational institution as students. In my preferred specification I exclude groups for whom wage determination is likely to follow a logic different from the standard market forces, including employers, career public servants, and member of security forces. In the robustness checks I relax these restrictions and all the key results are preserved. In all regressions I include the same set of controls described in Section 2.4 and cluster the standard errors at the mesoregion level (groupings of economically-related

¹⁹In this, as in most cases I tested, using only a small set of controls (population density, wages, education, informality and urbanization rates) already eliminates the correlation with the non-included control.

adjacent microregions) to address spatial autocorrelation concerns.

2.4.1 Gender-specific demand shocks and household migration

I begin by looking at the effects of gender-specific labor demand shocks on male and female migration. Table 2.1 shows the coefficients on gender-specific Bartik shocks in the regression model described in equation 2.17. The outcome variable is the log of the population of each subgroup that declared at the end of the decade (census year t) that they were living in a different microregion at the beginning of the decade (census year $t - 10$). Columns 1 and 2 present the effects of female and male shocks on own-gender migration. The next two columns report the effects of other-gender shocks. Column 3 presents the effect of male shocks on female migration, and column 4 the effects of female shocks on male migration. The last four columns present the test statistic and the p-value of Wald chi-square tests of the null hypothesis that the corresponding male and female coefficients are the same, based on seemingly unrelated regression models that include the correspondent female and male regressions. The difference that is being tested is indicated in the title of each column (for instance, the figures in the fifth column refer to tests of the difference of the coefficients in columns 1 and 2). For each outcome I also calculate the effect separately for two education subgroups: adult population with high-school degree or higher, and adult population without high-school degree. This serves as a reference to assess whether the aggregate gender effects could be affected by human capital heterogeneity within gender.²⁰ Unless otherwise specified, the tables of results for other outcomes follow the same layout.

The results provide strong support for the assumption that couples tend to migrate together, and the prediction that they are more likely to migrate in response to changes in labor demand for males. First, the effect of male Bartik shocks on own-gender migration is significantly larger than the equivalent effect of female shocks. This is true in both decades, although the 2000s results are measured less precisely, consistent with the lower

²⁰Note that the population without high-school degree is still a significant majority of employees during this period, and variations in the Bartik shocks is disproportionally driven by variation in employment opportunity of this subpopulation.

Table 2.1: *Effects of gender-specific demand shocks on migrant population*

	Own-gender shocks		Other-gender shocks		Hypothesis tests (χ^2 and p-val.)			
	Females (1)	Males (2)	Females (3)	Males (4)	(1)-(2)	(3)-(4)	(1)-(4)	(2)-(3)
<u>Panel A: 1991-2000</u>								
All observations	3.90*** (0.70)	6.43*** (1.30)	6.61*** (1.31)	3.74*** (0.69)	8.17 0.00	10.83 0.00	2.25 0.13	1.11 0.29
Less than high school	3.74*** (0.72)	6.51*** (1.28)	6.57*** (1.29)	3.63*** (0.69)	10.03 0.00	11.65 0.00	0.89 0.35	0.12 0.73
High-school or higher	4.01*** (0.82)	7.69*** (1.42)	7.31*** (1.32)	3.95*** (0.79)	10.73 0.00	15.26 0.00	0.03 0.87	0.49 0.48
<u>Panel B: 2000-2010</u>								
All observations	1.48 (1.69)	3.19** (1.30)	2.82** (1.27)	2.00 (1.75)	0.53 0.47	0.12 0.73	3.44 0.06	2.98 0.08
Less than high school	1.90 (1.78)	3.09** (1.38)	2.78** (1.31)	2.31 (1.86)	0.22 0.64	0.03 0.85	1.10 0.29	1.00 0.32
High-school or higher	0.82 (1.72)	3.91*** (1.30)	3.48*** (1.33)	1.01 (1.80)	1.76 0.18	1.02 0.31	0.07 0.79	0.83 0.36

Note: Outcomes measured restricting the sample to individuals aged 15 through 64, excluding individuals in school, employers, civil servants, and public security. All regressions include a constant. Robust standard errors clustered at the mesoregion level in parentheses, except for the hypothesis tests. The hypothesis tests are Wald chi-square tests of the hypotheses of the type $H_0: \beta_{males} - \beta_{females} = 0$ on SUR models including the respective female and male regressions.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

aggregate migration of that decade. Second, male shocks also have a larger effect on other-gender migration than female shocks do. And third, the size of the response of the migrant population of men and women to the same shock (e.g. male and female migration in response to male shocks) is very similar.

The difference in male-female migration elasticity were more pronounced at younger ages, strongest between ages 15 and 34 (see Appendix Figure B.4). Moreover, the disproportionately larger responses of females to male shocks were much smaller and statistically insignificant in the 2000s, the decade when migration dropped in the country as a whole. This suggests that the migration mechanisms highlighted in Section 2.2 were more prevalent in the 1990s, which should be kept in mind when interpreting the rest of the results.

The composition of the migrant population is also consistent with the joint mobility assumption and the asymmetric response to male and female shocks. Appendix Table B.6 shows that while 57% of the adult population is married, 62% of the migrant adults

are.²¹ Among the population with less than high-school education, 69% is married in the aggregate, and 71% is married among the migrants. More importantly, if females staying behind was the norm in Brazilian internal economic migration, this would reflect on a disproportionately high share of males in the migrant population, particularly among the married migrants. But as shown in Table B.6, if anything, females are a larger share in the married migrant population. Males are a larger share in the migrant population only among singles. The fact that women migrated more than men over this period in spite relatively smaller own-gender migration elasticity suggests both that couples tend to move together, and that they tend to disproportionately follow males' work opportunities.

Moreover, it seems to be the case that females migrated more than males not because of, but in spite of their employment prospects. Appendix table B.7 reports aggregate economic outcomes separately for all individuals and for migrant individuals by gender and educational attainment category in 2000. While migrant women exhibit labor force participation rates and employment rates similar to the average, migrant males have lower non-participation and higher employment rates than the region as a whole. Migrant women without high-school education participate more in the labor force than the average, with a higher rate of failure at finding jobs as suggested by higher unemployment figures. Migrant women with high-school education tend to participate in the labor force less than the average, and still have higher unemployment rates. Migrant men participate more and have lower unemployment rates than the average in both educational attainment categories. In addition, migrant men tend to be disproportionately employed in the formal sector, and migrant women disproportionately employed in the informal sector.

²¹An important caveat of looking at differences between married and single populations is that the census only provides contemporaneous information on marital status, that is, marital status at the beginning of the period is not observed. It is likely that marital status is endogenous to labor market shocks, since individuals' economic situation tend to affect their propensity to marry.

Table 2.2: *Effects of gender-specific demand shocks on population*

	Own-gender shocks		Other-gender shocks		Hypothesis tests (χ^2 and p-val.)			
	Females (1)	Males (2)	Females (3)	Males (4)	(1)-(2)	(3)-(4)	(1)-(4)	(2)-(3)
Panel A: 1991-2000								
All observations	0.29 (0.20)	0.71*** (0.25)	0.65** (0.29)	0.33** (0.17)	13.44 0.00	3.84 0.05	0.50 0.48	0.39 0.53
Less than high school	0.16 (0.24)	0.97*** (0.24)	0.88*** (0.30)	0.20 (0.19)	21.38 0.00	13.57 0.00	0.29 0.59	0.63 0.43
High-school or higher	0.56*** (0.18)	0.42 (0.40)	0.45 (0.29)	0.36 (0.24)	0.12 0.73	0.13 0.72	0.86 0.35	0.01 0.92
Panel B: 2000-2010								
All observations	0.09 (0.23)	0.76*** (0.20)	0.75*** (0.18)	0.10 (0.26)	4.23 0.04	3.40 0.07	0.02 0.90	0.01 0.92
Less than high school	-0.28 (0.29)	0.81*** (0.22)	0.75*** (0.21)	-0.56* (0.30)	5.85 0.02	8.72 0.00	7.77 0.01	0.36 0.55
High-school or higher	0.78** (0.40)	1.24*** (0.39)	0.60 (0.37)	0.65 (0.51)	0.59 0.44	0.01 0.93	0.15 0.70	7.50 0.01

Note: Outcomes measured restricting the sample to individuals aged 15 through 64, excluding individuals in school, employers, civil servants, and public security. All regressions include a constant. Robust standard errors clustered at the mesoregion level in parentheses, except for the hypothesis tests. The hypothesis tests are Wald chi-square tests of the hypotheses of the type $H_0 : \beta_{males} - \beta_{females} = 0$ on SUR models including the respective female and male regressions.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

2.4.2 Effects on population and housing rents

The same gender asymmetries can be seen using log changes in population as the dependent variable (Table 2.2). A ten percent predicted increase in employment for men was associated with a 7.1 percent increase in male population in the 1990s and a 7.6 percent increase in the 2000s. Conversely, a ten percent predicted increase in employment for women corresponded to increases of 2.9 percent in 1990s and 0.9 percent in 2000s, both statistically insignificant. Migration responses appear to have been an important mechanism of adjustment to geographically heterogeneous changes in demand, particularly in the 1991-2000 period, a finding that is at odds with Dix-carneiro and Kovak (2016), who find little evidence of interregional migration in Brazil in response to a trade liberalization shock but consistent with Morten and Oliveira (2016), who find strong migration responses to changes in road infrastructure.

The effects on population by schooling category also reveal the presence of potential

Table 2.3: *Effects of gender-specific demand shocks on rents*

Dependent Variable: Δ Avg. Log Rent Residuals			
Aggregate Shock (1)	Female Shock (2)	Male Shock (3)	Diff. Test (χ^2 and p-val.) (4)
0.41 (0.31)	0.01 (0.27)	0.63** (0.31)	4.25 0.04

Note: Outcomes measured restricting the sample to households with rent data. All regressions include a constant. Robust standard errors clustered at the mesoregion level in parentheses, except for the hypothesis tests. Tests are Wald chi-square scores of the hypothesis $H_0 : \beta_{males} - \beta_{females} = 0$ on SUR models.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. **Sources:** See data appendix.

composition effects. In spite of the lower migratory response of women, the population of females with high-school education or higher grew significantly in response to female labor demand shocks. Factors like increases in female education, reduction in family size, and faster urbanization rates contributed to significant growth in female labor force participation during this period (Scorzafave and Menezes-Filho 2005). A possible explanation for the coefficients by education group in column 1 of Table 2.2 is that females endogenously acquired more education in localities with better female labor prospects. Another is that, following negative shocks to local employment, the female population that left (following their husband's or looking for better opportunities for themselves) were disproportionately low-educated, so that the female employment left behind was positively selected. In my model workers are only heterogenous in gender but not in skills, so it is not informative of this margin of adjustment. How education and gender interact in the context of local labor markets and joint mobility frictions is a promising area for future research.

Given that male labor demand shocks have a larger effect on migration and population, the model predicts that it should also have a larger effect on housing demand and ultimately housing rents. The results for this outcomes are reported in Table 2.3. I only observe

housing rents in the 1991 and 2010 census, so the coefficients on the table corresponds to regressions on differences over a 20-year period (in contrast with decade-long changes in all other tables). The dependent variable is the change in average log rent controlling for dwelling characteristics. Specifically, I run individual regressions of the log housing rent on a vector of characteristics of the property (see Appendix D), obtain the residuals, and average them at the microregion level to obtain the regional housing rents for each period.²²

I find that, while male shocks had a significant, positive, and large effect on housing rents, the effects of female shocks was not distinguishable from zero. A ten percent expected increase in male employment was associated with a 6.3 percent increase in housing rents. Compared to female labor demand shocks, male shocks made Brazilian regions more expensive over this period.

2.4.3 Employment effects

I turn now to the effects on employment. Recall that, in the context of the model, the effect on own-employment is expected to be positive for both men and women but larger for men. This is true even in the absence of migration effects, because female employment is restricted by their larger labor force participation costs. The model also predicts the effects of other-gender shocks to be positive and larger for males.

The employment predictions of the model are generally supported by the data, as reported on Table 2.4. A 10 percent increase in predicted male employment leads to a 14.2 percent increase in actual employment in the 1990s, and 12.7 percent in the 2000s. The effects are driven by the low-education population. Meanwhile, a 10 percent increase in predicted female employment leads to a 6.9 percent increase in actual employment in the 1990s, and a 3.2, not statistically significant percent in the 2000s. Among women, the effects are driven by the population with high-school degree or higher.

A deviation from the the model's predictions is the observed negative effect of female

²²All monetary variables in this paper are expressed in 2010 Reais, using INPC deflators published by the IBGE, and corrected using the method suggested in Corseuil and Foguel (2002).

Table 2.4: *Effects of gender-specific demand shocks on employment*

	Own-gender shocks		Other-gender shocks		Hypothesis tests (χ^2 and p-val.)			
	Females (1)	Males (2)	Females (3)	Males (4)	(1)-(2)	(3)-(4)	(1)-(4)	(2)-(3)
Panel A: 1991-2000								
All observations	0.69*** (0.20)	1.42*** (0.27)	0.85** (0.39)	0.55*** (0.20)	12.84 0.00	0.87 0.35	0.97 0.32	4.43 0.04
Less than high school	0.69*** (0.25)	1.66*** (0.27)	1.05** (0.45)	0.44** (0.23)	14.19 0.00	2.52 0.11	2.08 0.15	2.71 0.10
High-school or higher	0.91*** (0.24)	0.80* (0.48)	1.01*** (0.38)	0.43 (0.32)	0.04 0.84	1.60 0.21	1.55 0.21	0.19 0.67
Panel B: 2000-2010								
All observations	0.32 (0.36)	1.27*** (0.24)	0.44 (0.36)	-0.72* (0.39)	3.75 0.05	4.22 0.04	14.85 0.00	8.46 0.00
Less than high school	0.09 (0.39)	1.27*** (0.25)	0.44 (0.40)	-1.23*** (0.39)	4.52 0.03	7.02 0.01	22.06 0.00	5.73 0.02
High-school or higher	1.19** (0.47)	1.20*** (0.43)	0.35 (0.42)	0.70 (0.61)	0.00 0.99	0.21 0.65	1.22 0.27	7.61 0.01

Note: Outcomes measured restricting the sample to individuals aged 15 through 64, excluding individuals in school, employers, civil servants, and public security. All regressions include a constant. Robust standard errors clustered at the mesoregion level in parentheses, except for the hypothesis tests. The hypothesis tests are Wald chi-square tests of the hypotheses of the type $H_0 : \beta_{males} - \beta_{females} = 0$ on SUR models including the respective female and male regressions.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

shocks on male employment in the 2000s, which was concentrated in the low-education population. In the context of low migratory responses, this effect is ruled by input substitution in production in the theory. The negative coefficient suggest that low-education male labor may have been complementary to high-education female labor in this period. The labor force polarization literature in the U.S. has highlighted similar complementarities between high- and low-skilled workers (Autor and Dorn 2013; Autor *et al.* 2009). These potential interactions are not captured in a model that assumes away skills heterogeneity.

In the aggregate, the gender employment gap was exacerbated by male-leaning local demand shocks and improved by female-leaning shocks during the period of analysis. This can be seen in Table 2.5, which directly measures the effects of the shocks on the log differences across decades in the gap, defined as the ratio of male to female employment rates. These effects are largely driven by the low-education population, with no statistically significant effects on the gap among individuals with high school education or higher.

Table 2.5: *Effects on the employment gap*

	1991-2000			2000-2010		
	Females (1)	Males (2)	Diff. Test (χ^2 and p-val.) (3)	Females (4)	Males (5)	Diff. Test (χ^2 and p-val.) (6)
All observations	-2.37* (1.23)	2.88* (1.61)	13.13 0.00	-4.23*** (0.90)	3.45*** (0.92)	26.05 0.00
Less than high school	-2.78** (1.27)	3.35* (2.00)	12.58 0.00	-4.46*** (0.90)	3.47*** (1.02)	26.65 0.00
High-school or higher	-0.32 (0.40)	-0.08 (0.59)	0.56 0.45	-0.49 (0.32)	0.35 (0.24)	3.42 0.06

Note: Outcomes measured restricting the sample to individuals aged 15 through 64, excluding individuals in school, employers, civil servants, and public security. All regressions include a constant. Robust standard errors clustered at the mesoregion level in parentheses, except for the hypothesis tests. The hypothesis tests are Wald chi-square tests of the hypotheses of the type $H_0 : \beta_{males} - \beta_{females} = 0$ on SUR models including the respective female and male regressions.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

Male and female Bartik shocks affect the gender employment gap at different margins. Appendix Figure B.5 depicts the predictive margins at different levels of the shocks, that is, the predicted effects on the gap if every microregion had had the same intensity of the shock holding all other actual characteristics as they were in practice. In both decades, male shock tend to reduce the employment gaps only at lower levels. Shocks above median intensity appear to have no effect on the gap. For females, the negative effects on the employment gap are concentrated on the higher levels. The gap appears to have little sensitivity to female shocks below median intensity.

2.4.4 Wages and participation effects

I turn now to the effects of gender-specific local demand shocks on wages and labor force participation. It is useful to look at both outcomes together in the context of the model, where wages are the key endogenous driver of labor force participation decisions. Decisions to sort into the workplace are also affected by the opportunity cost of participating, which is exogenously determined in the model.

My preferred wage measure controls for observable individual characteristics like educa-

Table 2.6: *Effects of gender-specific demand shocks on wages*

	Own-gender shocks		Other-gender shocks		Hypothesis tests (χ^2 and p-val.)			
	Females (1)	Males (2)	Females (3)	Males (4)	(1)-(2)	(3)-(4)	(1)-(4)	(2)-(3)
<u>Panel A: 1991-2000</u>								
All observations	0.03 (0.14)	0.53*** (0.18)	0.35* (0.20)	0.38*** (0.09)	8.35 0.00	0.04 0.85	5.54 0.02	1.33 0.25
Less than high school	0.03 (0.15)	0.50*** (0.19)	0.35 (0.23)	0.37*** (0.09)	6.81 0.01	0.01 0.93	4.84 0.03	0.76 0.38
High-school or higher	-0.25 (0.19)	0.07 (0.30)	-0.12 (0.27)	0.08 (0.20)	0.93 0.34	0.37 0.54	1.58 0.21	0.26 0.61
<u>Panel B: 2000-2010</u>								
All observations	0.37 (0.26)	0.56*** (0.21)	0.26 (0.20)	0.33 (0.27)	0.27 0.60	0.04 0.84	0.04 0.85	2.15 0.14
Less than high school	0.46 (0.35)	0.47** (0.22)	0.17 (0.25)	0.27 (0.27)	0.00 0.99	0.05 0.82	0.38 0.54	1.71 0.19
High-school or higher	0.18 (0.29)	0.59*** (0.22)	0.33 (0.23)	0.58* (0.30)	1.20 0.27	0.39 0.53	2.02 0.16	1.07 0.30

Note: Outcomes measured restricting the sample to individuals aged 15 through 64, excluding individuals in school, employers, civil servants, and public security. All regressions include a constant. Robust standard errors clustered at the mesoregion level in parentheses, except for the hypothesis tests. The hypothesis tests are Wald chi-square tests of the hypotheses of the type $H_0 : \beta_{males} - \beta_{females} = 0$ on SUR models including the respective female and male regressions.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. **Sources:** See data appendix.

tion levels, age as a proxy of work experience, and race. Specifically, I calculate the residuals of individual-level Mincer-style regressions of log wages on individual characteristics (Mincer 1974). I then use the regional averages of these wage residuals for the sub-population of interest. The use of this formulation is fairly standard in the urban literature (e.g. Glaeser and Gottlieb 2009b; Chapter 1 in this dissertation).

I find that, on average, local male wages increase more than local female wages in response to an equivalent shock in demand (Table 2.6). In the aggregate sample, which includes both education groups, the effect of own-gender shocks on female wages is not statistically different from zero, while the equivalent effect on male wages is. The difference across genders is only significant in the 1990s, when the migration effects are more likely to be present.

The fact that male shocks simultaneously lead to larger employment and wage effects than equivalent female shocks is hard to explain in an partial equilibrium setting. Larger

Table 2.7: *Effects of gender-specific demand shocks on non-participant population*

	Own-gender shocks		Other-gender shocks		Hypothesis tests (χ^2 and p-val.)			
	Females (1)	Males (2)	Females (3)	Males (4)	(1)-(2)	(3)-(4)	(1)-(4)	(2)-(3)
<u>Panel A: 1991-2000</u>								
All observations	0.26 (0.22)	-1.45*** (0.40)	0.58* (0.33)	-0.43* (0.22)	25.71 0.00	12.22 0.00	9.53 0.00	39.45 0.00
Less than high school	0.06 (0.26)	-1.18*** (0.39)	0.80** (0.34)	-0.64*** (0.24)	11.73 0.00	25.22 0.00	10.15 0.00	33.90 0.00
High-school or higher	0.18 (0.32)	-1.48 (1.02)	0.01 (0.56)	-0.37 (0.50)	2.62 0.11	0.31 0.58	0.99 0.32	1.95 0.16
<u>Panel B: 2000-2010</u>								
All observations	0.45** (0.20)	-0.49 (0.37)	0.26 (0.19)	1.03*** (0.39)	4.65 0.03	2.95 0.09	2.85 0.09	5.42 0.02
Less than high school	-0.03 (0.29)	-0.32 (0.42)	0.31 (0.26)	0.23 (0.48)	0.25 0.62	0.02 0.88	0.50 0.48	3.49 0.06
High-school or higher	0.31 (0.59)	1.47* (0.76)	0.56 (0.47)	0.31 (0.93)	1.28 0.26	0.05 0.82	0.00 1.00	2.36 0.12

Note: Outcomes measured restricting the sample to individuals aged 15 through 64, excluding individuals in school, employers, civil servants, and public security. All regressions include a constant. Robust standard errors clustered at the mesoregion level in parentheses, except for the hypothesis tests. The hypothesis tests are Wald chi-square tests of the hypotheses of the type $H_0 : \beta_{males} - \beta_{females} = 0$ on SUR models including the respective female and male regressions.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

employment effects suggest a more elastic labor supply curve, which in turn should imply relatively smaller wage effects. In general equilibrium, however, male shocks could have larger wage and employment effects than female shocks because of large compensating differentials for local costs of living. The fact that housing rents react to changes in male labor demand but not to changes in female demand supports this interpretation.

The combined wages and participation effects, however, cannot be fully accounted for by the assumptions of the model. Male shocks led to large immigration not only of males, but also of females, and this shift in supply should have generated a downwards pressure on the female wage. One explanation for why male shocks would not drive down female wages, still within the model, is again compensating differentials, since increases in male labor demand make the region more expensive for both members of the household. However, that explanation is at odds with the fact that the non-participant population of females increased in response to male shocks, particularly during the 1990s (Table 2.7). In the model, higher

wages should have led to higher participation.

Other potential explanations requires me to step out the current assumptions of the model. One is the role of income effects. Indeed, an added worker effect, whereby married females reduce their labor force participation in response to increased wages or employment of married men, has been well-documented in the literature at the individual level (Fernandes and de Felicio 2005; Soares and Izaki 2002). This effect could potentially account for both the reduced participation rates and the increased female wages, as the female supply shifts backwards. A natural extension of the model is therefore to incorporate leisure or home production to the utility function to allow for income effects.

Other non-exclusive explanations for the observed results involve skill composition effects. As the population results in Table 2.2 suggest, women could have endogenously sorted into education in regions with higher female labor demand, pushing up female wages. If the added-worker effect was disproportionately prevalent among the low-education population, it could result in the females that remain in the labor force being positively selected (Olivetti and Petrongolo 2008). For example, Hunt (2002) showed that in East Germany following reunification the average wage gap fell significantly, but a large portion of this change was explained by involuntary exit from the labor force of low-skilled workers, who were disproportionately women.

The introduction of an additional dimension of heterogeneity – education – makes the effects of the shocks on participation and wages potentially non-monotonic. For instance, while workers with the same education could be imperfect substitutes, high- and low-education workers could be complements (Moretti 2004). If skilled female workers and unskilled male workers complement each other in production, increases in demand for the former will result in higher wages and participation of low-skilled males, like those observed in the 1990s in Brazil. These effects could account for the fact that shocks to female labor demand affect the gender wage gap differently in different decades (Table 2.8), or that the predictive margins of these effects are non-linear and non-monotonic (as illustrated in Appendix Figure B.6). Further research is needed to better understand the way in which

Table 2.8: *Effects on wage gaps*

	1991-2000			2000-2010		
	Females (1)	Males (2)	Diff. Test (χ^2 and p-val.) (3)	Females (4)	Males (5)	Diff. Test (χ^2 and p-val.) (6)
Panel A: 1991-2000						
All observations	0.34*** (0.13)	0.09 (0.20)	2.22 0.14	-0.22 (0.27)	0.48** (0.23)	3.16 0.08
Less than high school	0.33** (0.15)	0.03 (0.19)	3.15 0.08	-0.10 (0.34)	0.39 (0.24)	1.17 0.28
High-school or higher	0.35 (0.27)	0.10 (0.42)	0.99 0.32	0.34 (0.36)	0.21 (0.29)	0.08 0.78

Note: Outcomes measured restricting the sample to individuals aged 15 through 64, excluding individuals in school, employers, civil servants, and public security. All regressions include a constant. Robust standard errors clustered at the mesoregion level in parentheses, except for the hypothesis tests. The hypothesis tests are Wald chi-square tests of the hypotheses of the type $H_0 : \beta_{males} - \beta_{females} = 0$ on SUR models including the respective female and male regressions. *** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

gender and education differences interact to shape outcomes in local labor markets.

2.4.5 Robustness

I perform multiple robustness checks of the results. First, I evaluate if the results vary without the exclusions from my main sample. Second, I test if the results are sensitive to the definitions of local labor markets and the definitions of industries. Third, I evaluate the sensitivity of the results to the inclusion of different sets of controls.

Results may be sensitive to the definition of labor markets, among other things due to spatial autocorrelation. Regions close to each other may have similar results, or outcomes may spillover to neighboring local markets. Concerns about potential geographic correlation of shocks is reduced by clustering all standard errors at the level of mesoregions, the next higher level of geography. However, this does not solve the issue of neighbors' spillovers. I re-run the analysis using the definition of minimum comparable microregions from Dix-carneiro and Kovak (2016). This involves a coarser aggregation of municipalities as noted in Section 2.3.3, with the final sample including 411 microregions as opposed to 539 in my

sample. The results are less precisely measured but still statistically significant in all the relevant cases, and all the findings discussed above hold.

My main specifications include all regions of Brazil in unweighted region-level regressions. Given the distribution of regional characteristics this implies that relatively small, frequently less urbanized regions, drive the results. To assess the extent to which the findings also hold in big urban centers I replicate the analysis restricting it to large urban conglomerates. Specifically, I use “Arranjos Populacionais” (IBGE 2016), which are groupings of core urban centers with municipalities closely integrated to them through work and education daily commuting. I consider both urban agglomerations and self-standing municipalities with large urban populations, and use the same correction described in Section 2.3.3 to account for changing administrative boundaries. Most of the key results of the analysis are also found in this urban centers sample. In this sample, non-employment for women also decreases in the 1990s following a male demand shock, although this reduction is still significantly smaller than the reduction in male non-employment. This is consistent with the view that local employment shocks in urban centers are less distortionary than shock in less urbanized regions, since tied-migrant females are more likely to find jobs in denser agglomerations.

I also verify that other smaller sources of that the industry definition used does not alter my results significantly. I take the industry definitions from Dix-carneiro and Kovak (2016) and replicate the analysis. The Bartik shocks calculated in with this industry definition are highly correlated with the ones calculated with the census definitions, and the regression results are unchanged in all the main specifications of interest.

The restrictions included in the sample do not seem to affect the results in any significant way either. The main conclusions are preserved if the sample includes only adult population aged 25 to 64 (that is, dropping the population aged 15 through 24 relative to my main sample), or include self-employed, government workers and domestic workers.

Adding and subtracting plausibly relevant controls – beyond the core baseline controls which include starting income levels, age structure, urbanization rate, and share of high-

school educated population – also leave the results largely unchanged. An exception is the inclusion of base year share of non-employed men and women. After controlling for initial non-employment the discrepancy in the wage effects loses statistical significance. However, the differences in employment, population, and non-employment effects remain, and the non-employment effects among females are unambiguously positive and statistically significant. Consistent with the theory, low labor force participation among local residents does seem to be driving part of the wage effects. However, in the aggregate, I still find local labor supply of male labor to be more elastic than female labor due to larger male migration elasticity.

2.5 Welfare and policy implications

The results discussed above imply that male and female labor demand shocks can have very different welfare consequences. When local demand for female labor increases, firms on average are able to tap into a labor pool readily available in the region, which translates into higher employment for incumbent residents. In this context, local residents are able capture a larger share of the economic rents generated by the shock than workers living outside of the region and potential migrants. Because females have a more modest immigration effect than men, the pressure on housing prices, is not as high. As discussed by (Moretti 2011), in this kind of situation housing prices see limited increases and workers are likely to receive a larger fraction of the benefits than landlords.

The evidence presented in this paper also raises the possibility that female labor is less efficiently allocated across space than male labor. This implies that local labor demand shocks for women could potentially bring about aggregate efficiency gains to the national economy, reducing misallocation. This is a promising area for future research.

Welfare consequences of male labor demand shocks are very different. When local demand for male worker increases, there is a larger migratory response. The framework suggests that in this situation local workers are likely to share larger fractions of the economic rents with migrant workers and landlords. .

These patterns are important in many policy contexts, and in particular for regional development policies. These policies typically have as one of their main goals the generation of jobs for locals in underdeveloped regions. They are widespread throughout the globe (Kline and Moretti 2014a) and have been used in Brazil since at least the 1940s (Resende 2013; Cavalcanti Ferreira 2004). My findings suggest that the same policy can have very different effects depending on whether job growth favors male or female employment. If policies favor job creation for men over job creation for women, benefits to local residents are more likely to dissipate through migration and higher local costs of living. Moreover, the initial employment rates for men and women are likely to matter. “Place-making” policies may be more effective in improving the economic conditions of locals in places with lower levels of female employment.

2.6 Conclusion

This paper shows that the effects of local labor demand shocks can differ significantly by gender. I compare shifts in local labor demand for males and for females in the context of Brazil during the period 1991-2010. Male employment shocks, relative to equivalent female shocks, lead to larger increases in population, rents, and the gender economic gap.

I interpret these results in light of a spatial equilibrium model with gender-segmented labor markets. In this framework, the gender differences in population and employment effects are related to joint mobility constraints of married couples. Because men have in expectation lower opportunity costs of participating in the labor force than women, male jobs prospects carry a larger weight in household location decisions than female job prospects. As a consequence, the migration elasticity of households is larger with respect to male than with respect to female demand shocks, and the former have larger effects on local population and prices than the latter. Because of tied migration, shocks in labor demand of one gender also affect the local labor supply of the other gender, with larger effects in male than in female shocks. The empirical results are largely consistent with the presence of this mechanism. Other non-exclusive margins of adjustments that appear to be at play

are composition effects – related to females becoming more educated and supplying more labor over time – and income effects – which lead individuals to supply less labor if their partners’ improve their work conditions. These are important areas for future research in local labor markets.

The presence of gender-differentiated migratory adjustments have important welfare implications. Because increases in local employment prospects for men place larger pressures on housing rents than equivalent shocks for women, part of the male wage effects captures compensating differentials for higher costs of living. It follows that while male shocks are more likely to benefit migrants and landlords and exacerbate the gender economic gap, female shocks are more likely to benefit local residents and reduce economic inequalities across genders. In addition, tied migration may lead to geographic misallocation of female labor, as tied-migrant women locate in regions that are not necessarily their individually optimal choice, and tied-stayer women take less advantage of jobs opportunities outside their place of residence than men do.

These findings have important policy consequences. Regional development and other “place-making” policies can lead to very different outcomes depending on how they affect labor demand for men and for women. In contexts with large gender and geographic disparities like Brazil during the period of study, this paper points towards significant advantages of expanding local female job opportunities.

Chapter 3

Education and Local Labor Market Outcomes.

Evidence from a Large Federal Program in Brazil

3.1 Introduction

Policymakers often turn to education as a strategy to promote economic development in their localities in the medium and long run. However, the existing literature gives us reasons to be skeptical of this approach. While researchers have indeed established that more educated cities tend to grow faster (Glaeser *et al.* 1995; Glaeser and Shapiro 2003; Shapiro 2006; Gennaioli *et al.* 2014), it is unclear that investing in local education will necessarily result in higher local education levels, because educated individuals may leave the city if they find better work opportunities elsewhere (Abel and Deitz, 2012). This local “brain drain” may be more pronounced in low and middle-income countries, where differences in economic opportunities across rich and poor cities are oftentimes larger than those among rich and poor countries (Acemoglu and Dell 2010). Even if frictions prevent the educated population from migrating, the labor market effects of local education expansion remain ambiguous. Education can both increase the individual productivity and generate productivity spillovers, boosting labor demand (Moretti 2004.) But if the growth in the supply of educated workers outpaces demand growth, it can also push down the equilibrium local wage for this group. This paper studies empirically the effects of expanding local

education on the economic outcomes of individuals and of places in the context of Brazil.

In order to capture exogenous expansions of local education investment, I use FUNDEF, a large federal policy enacted in the late 1990s. FUNDEF effectively redistributed sizable resources earmarked for primary education and middle school across municipalities within states. Because of the resource allocation rules and the timing of the policy announcement and implementation, the municipality-level changes in education resources produced by the policy in its first year were non-predictable and uncorrelated with the local policy preferences (Estevan 2015; Menezes-Filho and Pazello 2006.)

I start by showing that FUNDEF led to an increase in the educational attainment of the individuals exposed to the policy. Following Duflo (2001), I take advantage of the fact that age at the time of implementation mediates individual's exposure to the program. Those who were of middle-school age or younger in 1998 were potentially exposed, whereas individuals that were older than middle-school age were not. I demonstrate the "FUNDEF shock" - i.e. the size of the policy-related changes in local public education budgets - led to higher educational attainments among the cohorts that were in principle exposed to the policy, compared to the cohorts that were not. A policy-mandated one percent increase in the baseline education budget in the individual's municipality of education was associated with a 2.4 higher likelihood of completing at least primary school. The equivalent figures are 1 for middle-school and 0.4 for high-school.¹

Individuals exposed to FUNDEF were also more likely to migrate after finishing their education. One percentage point education budget increase was associated with a 1.2 percentage points increase in the likelihood of being a migrant in 2010 for this group, relative to those who were not exposed.

Higher local public education expenditures also led to higher wages for the beneficiaries of the policy. A one percent larger shock was associated with an average 1.7% increase in hourly wages for individuals who were in principle exposed to the program, relative to

¹FUNDEF may have also increased the probability of graduating from college, but the period between the year in which the policy was implemented (1998) and the year in which the outcomes are measured (2010) is insufficient to have precise measures of this effect.

those who were not. The effects was completely driven by male workers, for whom hourly wages increased by 2.8%. The effect on female wages was statistically non-distinguishable from zero. The program also increased informality rates and unemployment, mainly among women. Its effects on labor force participation was mixed: positive along the intensive margin, and negative along the extensive margin, with minor differences across genders.

The program's impact on individual wages appears to be more related to migration than to productivity effects. I estimate a direct effect of the program on hourly wages of 5.2% for migrants, but only of 0.8% for non-migrants. The difference is highly statistically significant. Using exposure to FUNDEF as an instrument for individual educational attainment, I find average returns of 1.9% for middle school attainment, and of 0.8% for high school attainment. However, when I control for region-of-work fixed effects, these estimates become *negative* 0.6 and 1.01, respectively. These findings are consistent with prior research showing that a large fraction of the wage effects of migration are explained by the characteristics of the destination place, rather than of the individual (Clemens 2013; De la Roca and Puga 2017.) Interestingly, the gender differences in the effects on labor market outcomes are not explained by differential effects on the probability of migrating, suggesting that other mechanisms -such as the male-biased joint mobility decisions studied in Chapter 2- may be at play.

To study the effects of education expansion at the local labor market (microregion) level, I use a standard difference in differences regression set-up. After conditioning on long-term trends in local labor market outcomes, regional program intensity is uncorrelated with 1990s trends in the share of primary school educated in the working-age population, suggesting that the approach is valid in this context. This is not the case for measures of higher education attainment (i.e. middle school and high school.)

I find that FUNDEF had a positive impact on aggregate educational attainment, particularly at the primary school level. A program-induced one percentage point in the local public education budgets (corresponding to 1.63 standard deviations) was associated with a 7.5 percentage points increase in the the share of (at least) primary-educated in the adult

population (corresponding to 0.44 standard deviations.)

In spite of increasing the share of primary-educated, FUNDEF was associated with worsening average local wages, labor force participation, formality rates, and unemployment. The evidence suggests that this is because local supply of educated labor outpaced local demand. Using program exposure to instrument for changes in the share of primary educated in the adult population, I find that a one percentage point higher share was associated with a 0.3 decrease in the average hourly wage net of observable individual characteristics and a positive -although not statistically significant- change in employment.

This chapter contributes to the literature on the effects of school spending on educational and labor market outcomes. Recent work has found that increases in education investments lead to higher educational attainment (Hyman 2017) and better labor market outcomes (Jackson *et al.* 2016) in the U.S. context. The connection between education resources and learning outcomes is empirically weaker (Hanushek 2003.) This paper highlights an important mechanism mediating the connection between education investments and labor market outcomes, namely, the effect of these investments on the individual likelihood of migrating to more productive regions.

A related literature studies the geographic sorting of workers by skills, and how it affects econometric estimates of returns to schooling. More educated workers in the U.S. tend to migrate to places where the returns to education are larger (Heckman *et al.* 1996; Dahl 2002), and where better amenities can be found (Dahl 2002; Diamond 2016.) This generates an upward bias in OLS estimates of returns to education in local labor markets (Dahl 2002). In addition to documenting similar patterns in a developing country context, my work points to large gender differences in selection. Increases in local public schools budgets in Brazil raised both educational attainment and the probability of migrating for both men and women, but while males obtained significantly better labor market outcomes, females did not.

My work also relates to the literature on the effectiveness of place-based policies. Economists have been skeptical of growth-promotion investments targeting specific cities

or regions, because mobility responses may undermine potential benefits of these policies for locals (Glaeser and Gottlieb 2008; Kline and Moretti 2014b). In recent work Austin *et al.* (2018) take a more favorable view specifically with respect to place-based policies targeting local labor demand, arguing that joblessness is a more acute social problem than low income in the U.S., and these policies are likely more effective at alleviating it than “people-based” policies. My paper treats local education investments as place-based policy targeting labor supply, and shows that in Brazil these investments led to better outcomes for individuals but not for places. Migration was a key mediator, as individuals that obtained a higher education left to places with better economic opportunities. Existing work in the U.S. context has also found a prominent role of migration adjustments in determining the local effects of academic R&D activities (Abel and Deitz, 2012) and the establishment of new colleges (Andrews 2017.)

Finally, I make a contribution to the literature on education as a driver of economic growth. While at the country level multiple studies have failed to find a connection between human capital and growth (Pritchett, 2006), or have found it only in a subset of countries (Krueger and Lindahl, 2001), at the local level the literature has documented a strong connection between initial schooling levels and subsequent growth in population and/or wages (Glaeser *et al.*, 1995; Shapiro, 2006; Da Mata *et al.*, 2007; Gennaioli *et al.*, 2014; Chapter 1 in this dissertation). Local governments around the world motivate education expenditures as long-run development strategies. This paper shows that these investments not only can be ineffective at improving the labor market conditions of residents, but they can lead to worsened outcomes if the supply of qualified labor is not met by corresponding demand increases. Education investments are likely justified given their positive effects on multiple other outcomes including crime rates (Lochner and Moretti 2004), health and mortality (Lleras-Muney 2005), fertility rates, and the stability of marriages (Oreopoulos and Salvanes 2011), to name just a few. But their prospective effects on local economic development are not unambiguously positive.

The remainder of this chapter proceeds as follows. Section 3.2 describes the FUNDEF

program and related facts about the context in which the policy was implemented. Section 3.3 discusses the data and how I use the variation introduced by the program to identify the effects of increases in education attainment on individual and on local labor market outcomes. Section 3.4 presents and discusses the evidence of the effects of FUNDEF on individual educational attainment, migration, and labor market outcomes. Section 3.5 focuses on the effect of the program at the regional level. Section 3.6 concludes.

3.2 The FUNDEF program and its context

The Fund for Sustainment and Development of Fundamental Education and Appreciation of Teaching - FUNDEF (Fundo de Manutenção e Desenvolvimento do Ensino Fundamental e de Valorização do Magistério), was enacted in July of 1998 with the goal of improving the distribution and spending efficiency for basic and middle-school education within states. The 1988 Constitution had mandated that state and municipal governments invest at least 25% of their total revenues in public education. This rule brought about large differences in the public education budget and the per-student education expenditure across high-revenue and low-revenue subnational governments (Gordon and Vegas 2005; Estevan 2015), which the reform aimed to correct.

The reform targeted school years 1 through 8, of which years 1 through 4 were considered primary education (educação básica) and years 5 through 8 middle school education (ensino médio).² It kept in place the 25% minimum requirement, but introduced the mandate that three-fifths of these resources (i.e. 15% of total revenues) were to be transferred to a state-level fund, which then redistributed it to the municipal and state school systems according to their share in state-level enrollment for schooling years 1 through 8 (Menezes-Filho and Pazello 2006.) The reform also introduced a minimum level of spending per student. States with insufficient education budget became entitled to receive federal transfers to be able to

²Primary education was extended from 4 to 5 years to include kindergarten in 2003.

meet this benchmark.³ In addition, the reform mandated that 60% of the resources were to be spent in teachers' wages, while the remaining funds could be used for eligible operation and maintenance activities (De Mello and Hoppe 2005.)

The introduction of FUNDEF increased both the total resources locally spent on education and the share of municipal systems in these spendings. The program had a "decentralization" effect, in that it transferred resources from state to municipal public education systems because municipalities had higher enrollment relative to revenues than the states did (Menezes-Filho and Pazello 2006.) In spite of this, per-student transfers increased in real terms (De Mello and Hoppe 2005), and total municipal expenditure in education increased by about 8% (Menezes-Filho and Pazello 2006.) The program does not appear to have crowded out resources from other sources of financing (Gordon and Vegas 2005.)

The program had a relatively minor impact on the level of education in which the funds were invested. Most municipalities were already spending 60% or more of their mandated education budget (equivalent to 15% of their total budget) in Fundamental Education. The program did lead to a small initial reduction in expenditures in pre-school education (Menezes-Filho and Pazello 2006.) In 2006 FUNDEF was replaced by FUNDEB (Fund for the Development of *Basic* Education and Appreciation of the Teaching Profession"), which expanded the coverage of the fund to high-school education.

At the time of the introduction of the policy, the vast majority of students were enrolled in public education. Table 3.1 provides a break down of enrollment in the year 1997 for the grades affected by the program by school system. In that year, over 34 million students were enrolled in Fundamental Education. Around 90% of students were enrolled in public schools, either state or municipal (the share of federal schools in enrollment was negligible).⁴ Within public school enrollment, about 40% was in municipal systems and 60% in state systems.

³The impact of these transfers in the overall policy was relatively small. In 1998, a total of 8 out of 26 states received federal top-up transfers, which amounted to 3.7% of the total balance of the funds. By 2002 there were only 5 recipient states, with transfers accounting to 1.8% of the total funds (De Mello and Hoppe 2005.)

⁴Brazil had 5,507 municipalities and 26 states at the time the policy was implemented.

Table 3.1: *Enrollment in Fundamental Education in Brazil in 1997, by system*

Grades	Total enrollment	School System			
		Municipal	State	Federal	Private
1	6,575,734	58.2%	33.9%	0.0%	7.9%
2	5,154,094	46.7%	43.6%	0.1%	9.6%
3	4,724,389	41.6%	48.1%	0.1%	10.3%
4	4,113,911	38.9%	49.8%	0.1%	11.2%
5	4,510,872	21.5%	68.0%	0.1%	10.4%
6	3,630,218	19.8%	68.0%	0.1%	12.0%
7	2,993,337	18.0%	68.2%	0.2%	13.6%
8	2,526,833	16.4%	68.1%	0.2%	15.4%
Primary Education (1-4)	20,568,128	47.6%	42.8%	0.1%	9.5%
Middle School (5-8)	13,661,260	19.3%	68.1%	0.1%	12.5%
Fundamental Education (1-8)	34,229,388	36.3%	52.9%	0.1%	10.7%

Source: Brazilian Education Census of 1997.

In terms of net enrollment rates, even though Brazil was lagging behind relative to other middle-income countries by the beginning of the 2000s (De Mello and Hoppe 2005), it had experienced an unprecedented expansion in education at all levels starting in the early 1990s (Menezes-Filho 2001; De Barros *et al.* 2006.) Figure 3.1 shows the percentage of the adult population in each educational attainment category at the beginning and at the end of the decade. The share of the population with primary education or less went from 60% to 42%. Meanwhile, the share with high-school education increased from 20% to 29%, and the share with college education or higher from 7% to 13%. Females expanded their favorable schooling gap relative to males. By the end of the decade, 44% of adult women had achieved at least high-school education, compared to 39% of adult men. Existing research has shown that FUNDEF played a role in the increases in enrollment, particularly at the primary and middle school levels (Gordon and Vegas 2005; De Mello and Hoppe 2005; Menezes-Filho and Pazello 2006; Cruz and Rocha 2018).⁵

⁵Other social programs introduced during the 2000s, and in particular conditional cash transfers that required low-income families to enroll their children in school (Bolsa Escola and Bolsa Família), could also have had a role. Existing evaluations suggest that their contribution to enrollment in fundamental education was negligible, largely because the beneficiaries of these programs already had their children enrolled in school (Schwartzman 2005). However, they may well have had an impact -starting in the mid-2000s- at the high-school

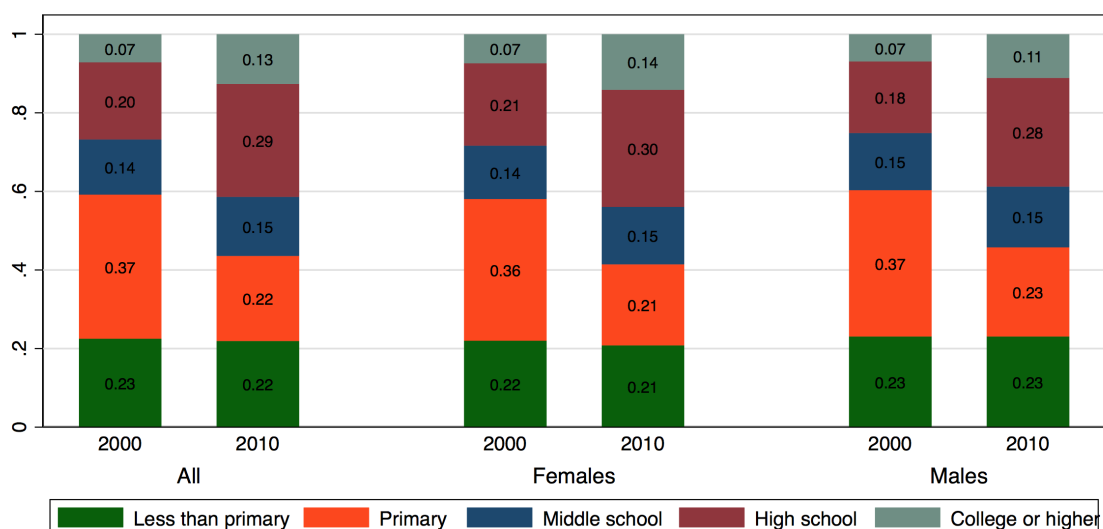


Figure 3.1: Percentage of Brazilian population age 23 or older in each educational attainment category

Note: Own calculations using census data. **Source:** Population censuses of 2000 and 2010.

The 2000s was also a decade of improving labor market conditions, especially for the *least educated* population. Figure 3.2 shows decade-long changes in labor market outcomes for adults wage earners aged 23 through 64 in five different education groups. During this decade, employment rates increased by 5 percentage points for workers with less than primary education. The increases were less pronounced at higher educational categories and were only of 1.5 percentage points for workers with a college degree or higher. This aggregate pattern is largely driven by females (Appendix Figure C.1.) Among males, the increase in employment rates was generally smaller and was more pronounced among higher education groups.

While part of the low-education employment growth reflects a recovery from unusually high unemployment rates in the 1990s,⁶ a dominating force behind employment growth comes from changes in female labor force participation (Corseuil *et al.* 2010.)⁷ Prior studies

level, where enrollment was smaller. On the demand side, Bolsa Família included stipends for youth aged 15 to 17 to attend school, while simultaneously FUNDEB expanded coverage of supply-side subsidies to high-school (OECD 2011.)

⁶During that decade, national unemployment rates grew sharply in Brazil. These trends were particularly severe among semi-skilled and low-skilled workers (Reis 2006.)

⁷Prior studies have also argued that the increase in female labor force participation was partly a response to

have found a strong positive association between education levels and female labor force participation in Brazil during the 1980s and 1990s (Scorzafave and Menezes-Filho 2001; Soares and Izaki 2002). During the 2000s, participation also increased significantly among low-education females (Appendix Figure C.1.) Participation rates actually decreased among males over this period, particularly at lower education levels. These changes were apparent mostly along the extensive margin. The intensive margin of participation (average weekly hours worked) decreased over the period, at rates that were similar for both genders and more pronounced at lower education levels.

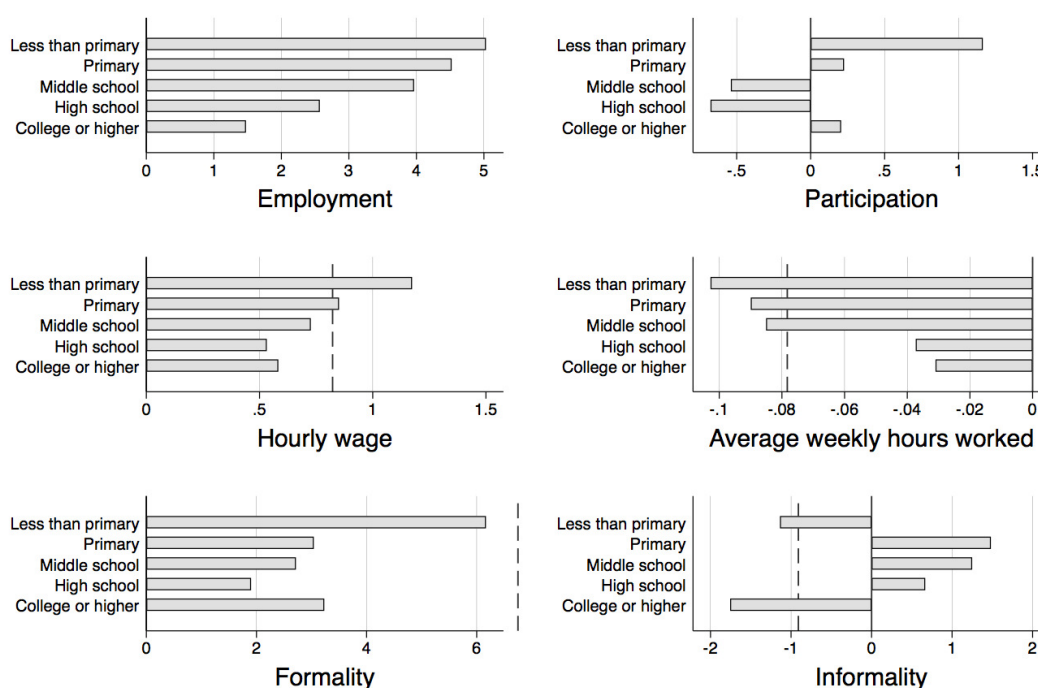


Figure 3.2: *Changes in labor market outcomes 2000-2010 by educational attainment category*

Note: Restricted to wage-earning population aged 23 through 64. Dashed lines denote population averages. All estimates are own calculations from microdata using sample weights. See the data appendix D for details on the measurement of each variable. **Source:** Population censuses of 2000 and 2010.

Much of the employment growth in this decade went towards formal jobs, but an important share of the increase in female labor force participation was also absorbed by the informal sector. Changes in formality rates were most pronounced at both the lowest and the

increased unemployment among male heads of household (Fernandes and de Felicio 2005.)

highest education categories. For both of these groups, aggregate formality rates increased *and* aggregate informality rates decreased over the period (Figure 3.2.) These aggregates, however, mask important gender heterogeneity. While formal employment increased for both males and females, informal employment dropped for males (particularly among the least educated) but *increased* for females (Appendix Figure C.1). Gender differences are more pronounced at middle-education levels (primary, middle school, and high school), where growth in female informal employment drives the overall increase in informality rates over the period.

Lower-education workers also saw notable wage increases in the 2000s. Adult wage-earners with less than primary education saw their average hourly wage more than double over the 2000s in real terms. For workers with a college education, the increase was close to 60%. The accelerated growth in both wages and employment during this decade is consistent with a net increase in labor demand, particularly among low-education groups.⁸

3.3 Measures, data and identification

This paper's empirical strategy relies on the redistribution of public education finance across municipalities introduced by FUNDEF. Because the program led to an increase in enrollment (De Mello and Hoppe 2005; Cruz and Rocha 2018) the program should also have had a measurable effect on the average educational attainment of individuals after school age and can potentially be used as a source of exogenous variation.

I study the effect of FUNDEF on educational attainment and on labor market outcomes at the individual and at the regional level. This section starts by describing the measures used and their data sources. It then turns to discussing the baseline empirical specifications employed, as well as the identifying assumptions.

⁸Prior studies have documented a relative increase in demand of low-education workers during the preceding decade, linking it with the national trade liberalization policy. Gonzaga *et al.* (2006) show that, following the early 1990s liberalization, employment shifted from high-skilled to low-skilled sectors (although the share of high-skilled in both sectors increased.) The skills wage differential dropped over this period. Dix-Carneiro and Kovak (2017) find that employment loss in regions that were hit the hardest by trade liberalization in the 1990s became more severe -rather than mean-reverting- during the 2000s.

3.3.1 Measures and data

To capture the impact of FUNDEF on local education finance I use the program-induced proportional change in local educational budget, following Estevan (2015). This municipality-level variable measures, for all education systems operating in the jurisdiction (municipal or state-level), the gap between the funds received from FUNDEF in the first year of implementation of the program⁹ and the funds contributed to the program (15% of the total revenues) in the same year.¹⁰ This gap is expressed as a share of the funds contributed to the program.

Formally, the municipality-level “FUNDEF Shock” measure is defined, for municipality j , as:

$$FS_j = \sum_{e \in \{m,s\}} \eta_{j,97}^e \left(\frac{I_{j,98}^e - O_{j,98}^e}{O_{j,98}^e} \right) \quad (3.1)$$

where the units of observations are school systems, denoted by the superscript $e = \{m, s\}$, which can be municipal (m) or state-level (s). The main weight is the share of the municipal system e located in municipality j in the state-level enrollment in public education in 1997 ($\eta_{j,97}^e$). The term in parenthesis is the program-induced percentage change in education transfers, where $I_{m,98}^e$ is the money that the municipal system received from FUNDEF in 1998, and $O_{m,98}^e$ the money that it contributed to the program’s state-level fund.¹¹

⁹I use only the variation of the first year of the program to address potential distortions related to municipalities inflating enrollment figures to capture additional FUNDEF funds. There is evidence showing that some municipalities did engage in this behavior in subsequent years. However, the 1998 transfers were based on data collected in 1997, before the allocation rules of the program had been announced (Estevan 2015).

¹⁰Local education budgets come from four taxes and transfers (FPM/FPE, IPIExp, LC87/96 and ICMS), as determined in the constitution of 1988. The FUNDEF policy applies to the money related to these sources.

¹¹A concern raised by Kosec (2014) is that the 1998 revenues may be affected by omitted variables (e.g. macroeconomic fluctuations) that also affect directly the outcome variables. To address this concern she employs a measure of the FUNDEF shock based on revenue data from 1997, the year prior to the start of the program. Estevan (2015) uses a similar approach to estimate a “predicted” impact of FUNDEF, specifically:

$$FS_j^{pred} = \sum_{e \in \{m,s\}} \eta_{j,97}^e \left(\frac{I_{j,97}^e - O_{j,97}^e}{O_{j,97}^e} \right)$$

where $O_{j,97}^e$ corresponds to 15% of the actual 1997 revenues, and $I_{j,97}^e$ is a simulated FUNDEF transfer, based on enrollment shares and simulated total value of each state-level FUNDEF fund in 1997. I replicate all the analyses using this alternative measure, and I obtain virtually identical results.

While this measure is useful to capture the exposure of a particular individual to FUNDEF, it doesn't adequately capture the incidence of the program in a particular local economy. This is because local labor markets in Brazil oftentimes incorporate two or more geographically proximate municipalities. Thus, in order to study the effects of FUNDEF on the aggregate outcomes of local economies - which I refer to as "regions" throughout the paper - I use a regional-level shock, namely:

$$FS_r = \sum_{j \in r} \varsigma_{j,97} \times FS_j \quad (3.2)$$

where $\varsigma_{j,97}$ is the share of municipality j in region r 's school-age population.

To approximate the boundaries of local labor markets I use "microregions". These are groupings of contiguous and economically integrated municipalities defined by the Brazilian Institute of Statistics (IBGE 2002). I use the time-consistent boundary definition from Chapter 2, which corrects for municipality-level boundary changes over the period of interest, following the method proposed by Kovak (2013).

The data used in this analysis comes from multiple sources. The enrollment data comes from the Brazilian School Census. The data on taxes and transfers used to calculate the resources contributed to and received from FUNDEF are from the National and State Treasuries (Secretaria do Tesouro Nacional, STN) and were compiled by Estevan (2015). The school-age population shares, as well as most of the outcome variables and controls, are constructed from the microdata of the decennial population censuses published by the IBGE. Appendix Figure C.2 shows the distribution of the FUNDEF shock measured at the municipal level and at the regional level. Appendix tables C.1 through C.5 report summary statistics and correlations for individual and regional-level variables. Appendix D offers a detailed description of each variable and their sources.

3.3.2 Identification of individual effects

The first part of the analysis focuses on the effect of education on individual's educational attainment and labor market outcomes. An important limitation is that the Brazilian

population census of 2010, the year in which outcomes are measured, did not record the exact number of years of schooling for individuals. Therefore my analysis is based on educational attainment categories.

To capture the direct effect of FUNDEF on individual outcomes, I follow Duflo (2001), and take advantage of the fact that the exposure to the program varies by year-of-birth cohort and by how the program affected resources for public education in the municipalities where the individual went to school. Specifically, in my baseline specification I estimate:

$$Y_{ijb} = \beta_0 + \sum_{a=l}^h (FS_j \times d_{ia}) \beta_{1,a} + \beta_2 FS_j + \sum_{a=l}^h (E_{j,97} \times d_{ia}) \beta_{3,a} + \beta_4 C_{j,97} + \beta_k + \beta_r + \epsilon_{ijb} \quad (3.3)$$

where the dependent variable Y_{ijb} is the outcome of interest measured in 2010 for individual i , educated in municipality j and born in year b . FS_j is the FUNDEF shock in municipality j (equation 3.1), d_{ia} is a dummy that takes the value one if individual i was age $a \in [l, h]$ in 1998, $E_{j,97}$ are fundamental education enrollment rates in municipality j in 1997, C_j is a vector of municipality of origin controls (ten age-group shares in the total population of municipality j in 1997), β_k is a cohort of birth fixed-effect, and β_r a region of work fixed effect (only used in of the some specifications in which Y_{ijb} are individual labor market outcomes.)

The set of cohorts included in each regression ($a \in [l, h]$) vary depending on the outcome variable. The youngest cohort l is chosen to ensure that the individuals included in the analysis were old-enough in 2010 for the outcome variable to be adequately measured. For instance, if the outcome variable is a dummy for having attained high-school education or higher, I use $l = 6$ so that the youngest cohort included was age 18 in 2010 (age 17 corresponds, in theory, to the last year of high school). For labor market outcomes I use $l = 3$ to ensure that all cohorts included were of working age in 2010. In all individual specifications I restrict the analysis to younger cohorts (up to age 40 in 2010). My baseline specification uses $h = 27$, and uses the cohort aged 28 in 1998 as the reference group.

For the coefficients $\beta_{1,a}$ to be given a causal interpretation, a given cohort's exposure to the program should be independent of the error term ϵ_{ijb} conditional on the controls.

Exposure to the program is a function of the individual's year of birth and the individual's municipality of education. Year of birth is exogenous. Municipality of education may, in principle, be endogenous if families with school-aged children selectively migrated towards beneficiary reasons. In practice, this appears unlikely because FUNDEF benefited regions where education opportunities were lagging relative to others. However, I replicate the analysis using region of birth¹² instead of region of education as a robustness check.

Given that this is a difference-in differences set up, the causal interpretation relies also on the assumption that, in the absence of the program, the changes in Y_{ijb} would not have been systematically different between individuals who studied in regions with high program incidence and individuals who studied in regions with low program incidence.

3.3.3 Identification of regional effects

The second part of the analysis turns to the effects of investments in public education on aggregate local education attainment levels and labor market outcomes. The unit of observation is the microregion. My preferred specification for all reduced-form analysis is the standard difference-in differences regression set up, namely:

$$Y_r = \alpha_0 + \alpha_1 Post + \alpha_2 FS_r + \alpha_3 (Post \times FS_r) + \alpha_4 (Post \times C_r) + \epsilon_r \quad (3.4)$$

where Y_r is the outcome of interest in region r , $Post$ is a dummy that takes the value 1 for the year 2010 (post treatment) and zero for the year 2000 (pre treatment), FS_r is the regional-level shock from equation 3.2, and C_r is a vector of regional-level lagged changes in local labor market conditions (measured in the 1980-1991 decade.)¹³

The key identifying assumption is that of “parallel trends”, namely, that in the absence of the program changes in Y_r would not have been systematically different between high-incidence and low-incidence regions, conditional on the controls. The use of lagged

¹²I do not observe region of birth directly in the data, but I can infer it for most individuals using their migration and residence data.

¹³Use 1980-1991 to measure pre-trends because the alternative (1991-2000) includes three years in which the program was already in place.

trends controls in this case is important, because the program targeted low-enrollment municipalities, and having low enrollment is likely correlated with pre-existing trends in the share of educated workers in the labor force, and in labor market outcomes.

3.4 Individual-level results

This section focuses on the effects of FUNDEF at the individual level. I start by exploring the effects on educational attainment. Second, I turn to the effects on the likelihood of migrating. Third, I look at labor market outcomes effects, and how they differ among men and women, and among migrants and non-migrants. Finally, I explore to what extent labor market outcomes can be explained by characteristics of the place of work, as opposed to individual-level characteristics.

3.4.1 Effects on individual educational attainment

I first turn to the effects of FUNDEF on individual educational attainment. Figure 3.3 plots the estimated coefficients $\hat{\beta}_{1,a}$ for cohorts $a \in [l, h]$ from a linear probability estimation of equation 3.3. These measure the effect of the exposure to the program on the likelihood of having a given educational attainment in 2010 for each cohort relative to the cohort aged 28 in 1998. The figure looks at four left-hand-side dichotomous measures of education attainment, namely, primary, middle school, high school and college. All measures take a value of one if the individual has achieved *at least* that attainment level in 2010.

If the program had an effect on individual educational attainment, we should observe it in the cohorts that were exposed to the program, and see no effects in the cohorts that were not exposed. Figure 3.3 shows the education level that corresponded to each cohort's age in 1998. Recall that the program targeted primary and middle school. This implies that the cohorts that were enrolled in these educational levels, as well as younger cohorts, were in theory exposed to the program, and older cohorts were not. Consequently, the x axes of the graphs in Figure 3.3 capture exposure to the program, the younger the individual, the greater the exposure. Individuals aged 6 or less in 1997 entered primary school when the

program was already in place, and were in principle fully exposed.

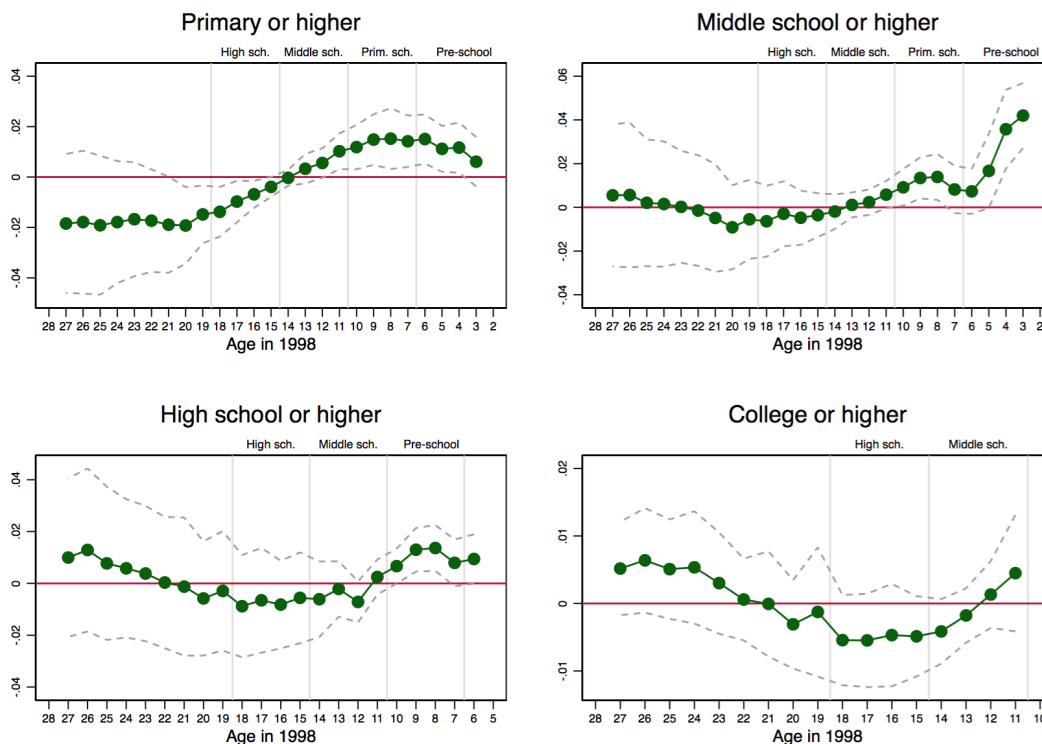


Figure 3.3: *Effects of FUNDEF on probability of reaching a specific educational attainment in 2010 by cohort*

Note: The markers represent the coefficient on the interaction of the FUNDEF treatment variable and each cohort dummy in equation 3.3. Dashed lines are 95% confidence intervals, with standard errors clustered at the municipality of education level. **Sources:** See data appendix.

According to these estimates, FUNDEF had a positive effect on primary school attainment, particularly among younger cohorts. A program-induced one percent higher local education budget in the municipality of education led to an increase of around 2 percentage points in the likelihood of having completed at least primary education for most cohorts that were at least partially exposed to the program, relative to the cohort aged 28 in 1998. I also find positive effects in individuals whose age in 1998 corresponded to the first two years of middle-school. This could be potentially explained by late school entrance and high repetition rates.¹⁴

¹⁴In Brazil, as in many other developing countries, the incidence of late school entry and the repetition rates are high (Estevan 2015). This implies that a subset of individuals in the cohorts that were old enough to have

FUNDEF also had a positive effect on middle school and high school completion. This is consistent with existing work showing that investment at lower education levels can increase enrollment and attainment at higher levels (Hyman 2017.) The effect is present and statistically significant for the cohorts that were, in theory, enrolled in primary school at the time that the program started. In the case of middle school attainment, it is noticeably larger and significant for the cohorts that were yet to enter primary school in 1998. In both cases, the effect is lower and not statistically significant for individuals that were ages 6 and 7 in 1998, suggesting that these cohorts had more difficulty in completing high-school than others. A potential explanation is that these cohorts faced a more challenging economic environment than others, given that they would have been in the last and second-to-last years of high school during the 2009 recession.

The trend of the coefficients across cohorts suggest that the program may have also increased the probability of graduating from college, but the available sample is insufficient to estimate this effect precisely. Individuals that were old-enough to have left college in 2010 were already enrolled in middle school in 1998, which implies that I can use for estimation only the four cohorts that were, in principle, the least exposed to the policy.

The figures show that the identification strategy is reasonable in this context. Both in the case of middle school and of high school, the cohorts that were in theory not exposed to the program had an educational attainment that was not statistically different than that of the comparison group in 2010.¹⁵

I obtain measures of the average effect of the program on the individual educational attainment by comparing all the cohorts that were, in principle, exposed to the program, to all the cohorts that were not. Specifically, I estimate the following variant of equation 3.3:

finished middle school by 1998 were still eligible to attend school and thus could have benefited from FUNDEF.

¹⁵In the case of the primary education measure, I do find a significant negative effect for cohorts aged 16 through 20 in 1998. This may reflect an omitted variable that was correlated with municipal FUNDEF exposure and negatively affected primary school enrollment in these cohorts in prior years.

$$Y_{ijb} = \beta_0 + \beta_1 (FS_j \times T_i) + \beta_2 FS_j + \sum_{a=l}^h (E_{j,97} \times d_{ia}) \beta_{3,a} + \beta_4 C_{j,97} + \beta_k + \epsilon_{ijb} \quad (3.5)$$

where T_i is a dummy that takes the value one if individual i was age 14 or younger in 1998, and the value zero if they were age 15 or older. Estimates of $\hat{\beta}_1$ are reported in Table 3.2.

The results suggest that, on average, one percent increase in the education budget in the individual's municipality of education led to a 2.4 percentage points higher likelihood of having completed at least primary school in 2010 for individuals who were in theory exposed to the program, relative to those who were not. The effects on the probability of completing at least middle and high school were of 1 and 0.4 percentage points, respectively. I don't find a significant average effect on the probability of completing college. Finding smaller effects at higher education levels is what we would expect given the program's target groups. By means of comparison, Duflo (2001) finds that a large school construction program in Indonesia induced about 6% of the population to complete at least primary education. She also finds a smaller effect on middle-school completion, and a negative effect on high-school completion.

FUNDEF appears to have had a stronger educational attainment effects among males than among females. Columns 2 and 3 in Table 3.2 report separate estimations for each gender from a seemingly unrelated regressions (SUR) model. Column 4 presents the results of tests of differences between the male and female coefficients. The point estimates are larger for males than for females in all attainment categories, although the difference is not statistically significant for middle school. Interestingly, in the gender-specific regressions I obtain statistically significant effects for college education attainment, which are positive for men, and negative for women. Appendix Figure C.3 presents estimations of cohort-specific effects by gender.

Table 3.2: *Effects of FUNDEF on probability of reaching of reaching a specific educational attainment*

	All	By gender		
	(1)	Male (2)	Female (3)	Test (4)
<i>(F-stat and p-val.)</i>				
<u>Panel A: Lowest education attained</u>				
Primary school or higher	0.024*** (0.001)	0.026*** (0.001)	0.022*** (0.001)	17.67 0.000
Middle school or higher	0.010*** (0.001)	0.010*** (0.001)	0.009*** (0.001)	1.61 0.204
High school or higher	0.004*** (0.001)	0.005*** (0.001)	0.002** (0.001)	4.69 0.030
College or higher	0.000 (0.001)	0.003*** (0.001)	-0.003*** (0.001)	29.14 0.000
Cohort of birth dummies	Yes	Yes	Yes	Yes
Enrollment rates times cohort of birth dummies	Yes	Yes	Yes	Yes
Demographic structure controls	Yes	Yes	Yes	Yes

Note: The table reports the coefficients on the treatment variable in equation 3.5. Regressions are at the individual level. Robust standard errors clustered at the municipality of education level in parentheses. Column 4 reports results of adjusted Wald tests of hypotheses of the type $H_0 : \beta_{males} - \beta_{females} = 0$ on SUR models' coefficients in columns 2 and 3.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

3.4.2 Effects on likelihood of migrating

Increases in local provision of education may lead to increased out-migration. If labor demand for educated workers is unevenly distributed in the national geography, newly-educated workers in places with low demand for skills will have the incentive to leave looking for opportunities that better match their qualifications. Moreover, if migration is costly (Morten and Oliveira 2016), individuals may be closer to the margin of migrating as they get educated and their potential income increases. Austin *et al.* (2018) document that, in the U.S., prime-age male migrants are on average more educated than the non-migrant population in their place of origin.

The effects of education on migration will, in turn, mediate the effects of education expansion on aggregate local labor market outcomes. Higher levels of education may make individuals more productive and give them access to higher paying jobs, but if the supply of

educated workers grows faster than the demand, the local private returns to education may be minimal or even negative. In that case, migration may allow individuals to obtain higher returns from their education, and potentially improve the returns of local non-migrant educated workers by alleviating excess supply.

If the effect of education on local labor market outcomes comes primarily from productivity spillovers, then improving local levels of education can lead to in-migration through increased labor demand. And the education profile of immigrants will, in turn, shape the aggregate education levels of the local economy further.

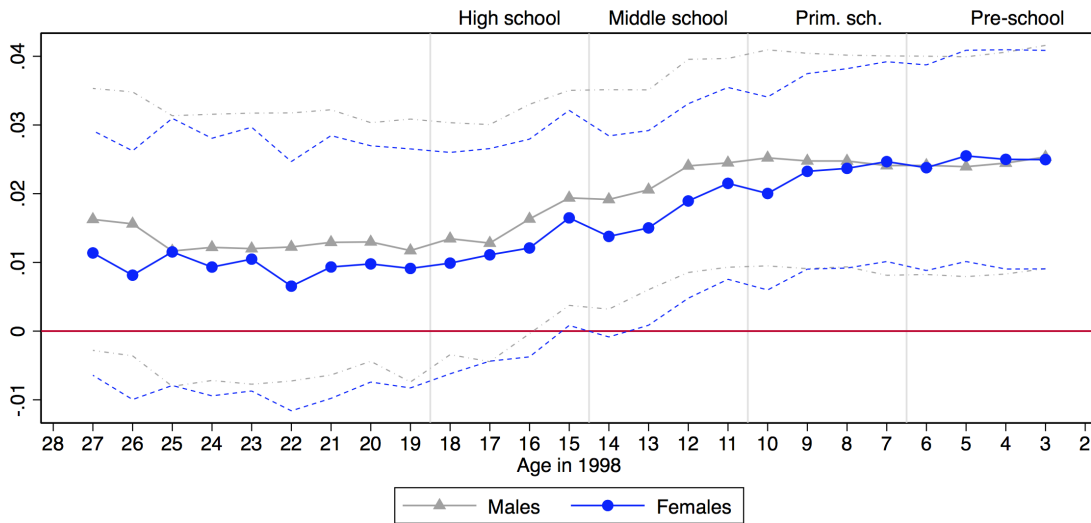


Figure 3.4: *Effects of FUNDEF on probability of being a migrant in 2010 by cohort and gender*

Note: The markers represent the coefficient on the interaction of the FUNDEF treatment variable and each cohort dummy in equation 3.3. Dashed lines are 95% confidence intervals, with standard errors clustered at the municipality of education level. **Sources:** See data appendix.

My empirical results suggest that individuals who were exposed to FUNDEF were more likely to migrate than the reference group. Figure 3.4 depicts the estimates for $\hat{\beta}_{1,a}$ in a linear probability estimation of equation 3.3, where the outcome of interest is a dummy for being a migrant. Here, migrant is defined as a person that in 2010 was living in a region different than the one where their municipality of education was located. The coefficients are statistically significant only for the cohorts that were, in theory, at least partially exposed to the program. Some of the point estimates appear to be higher, on average, for males, but

gender differences are not statistically significant in this specification.

Table 3.3 reports measures of the effect of the program on two measures of migration. The first measure defines migrant as someone that in 2010 was living in a municipality different than their municipality of education (even if it was in the same microregion). The second measure includes only migrants that in 2010 lived in a different microregion (the same definition as in Figure 3.4). Columns 1 through 3 report estimates of the reduced-form effect of the program ($\hat{\beta}_1$ in equation 3.5.) On average, one percentage point increase in the FUNDEF shock was associated with a 1.2 percentage points higher probability of migrating to a different municipality, and a 0.8 percentage points higher probability of migrating to a different microregion for the beneficiaries of the policy. These average (across cohorts) results also confirm that there are no measurable gender differences in the program's migration effects.

To explore the extent to which the effects of the program on migration operates through its effect on individual educational attainment, I estimate the following model using 2SLS:

$$Y_{ijb} = \beta_0 + \beta_1 \times d_{i,sch} + \sum_{a=1}^h (E_{j,97} \times d_{ia}) \beta_{3,a} + \beta_4 C_i + \beta_5 C_{j,97} + \beta_k + \beta_r + \epsilon_{ijb} \quad (3.6)$$

where $d_{i,sch}$ is a dummy that takes the value one if individual i has attained the level of schooling $sch = \{p, ms, hs, c\}$ in 2010,¹⁶ and C_i is a vector of individual-level characteristics, including sex and race (variation on individual age is already captured by the cohort of birth dummies). I instrument for $d_{i,sch}$ using the interactions of the FUNDEF shock with the cohort of birth identifiers, $FS_j \times d_{ia}$. The estimates are reported in columns 5 through 7 of Table 3.3.

The results suggest that increases in the probability of a higher education attainment lead to increases in the likelihood of migrating. The estimated effects are larger for primary education, and when explaining migration across municipalities (as opposed to migration across microregions). When using higher measures of educational attainment (middle school and high school) as the instrumented explanatory variable, I also find large and significant

¹⁶The levels of schooling are defined as having at least a given educational attainment, where attainment can be primary (p), middle school (ms), high school (hs) or college (c).

Table 3.3: *Individual effects of FUNDEF on migration*

	Reduced-form effects of FUNDEF on migration				2SLS estimates of effects of education on migration		
	All	By gender			By educational attainment		
		Males	Females	Test	Prim.	Mid-sch.	High-sch.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	<i>(F-stat and p-val.)</i>						
Probability of migrating to a different municipality	0.012*** (0.000)	0.012*** (0.000)	0.012*** (0.000)	0.42 0.51	0.461*** (0.121)	0.531** (0.240)	0.226** (0.110)
Probability of migrating to a different microregion	0.008*** (0.001)	0.007*** (0.001)	0.009*** (0.003)	0.18 0.42	0.326*** (0.076)	0.357** (0.148)	0.139* (0.077)

Note: Columns 1 through 3 reports linear probability estimates of $\hat{\beta}_1$ in equation 3.5. Column 4 reports results of adjusted Wald tests of the hypothesis $H_0 : \beta_{males} - \beta_{females} = 0$ on SUR models' coefficients in columns 2 and 3. Columns 5 through 7 report estimates of $\hat{\beta}_1$ in equation 3.6 where the instruments are the interactions of the FUNDEF shock in the municipality of education and the cohort fixed-effects. All regressions use weighting based on sample design. Robust standard errors in parentheses.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

effects on migrating to a different municipality, and relatively smaller effects on migrating to a different microregion.

3.4.3 Effects on individual labor outcomes

I turn now to the analysis of the effects of FUNDEF on individual labor market outcomes. Figure 3.5 shows estimates of the cohort-specific effects ($\hat{\beta}_{1,a}$ in equation 3.3) on wages and labor force participation. The figure at the top presents results for hourly wages net of observable individual characteristics. The two figures at the bottom measure labor force participation. To capture the extensive margin of labor force participation, I use a dummy that takes a value one if the individual is either formally employed, informally employed, or unemployed in 2010. To measure the intensive margin, I use the average number of paid hours worked per week. All regressions are based on a sample that includes individuals aged 15 through 40 in 2010, except for those that were enrolled in school in that year. Appendix D provides further details on measurement.

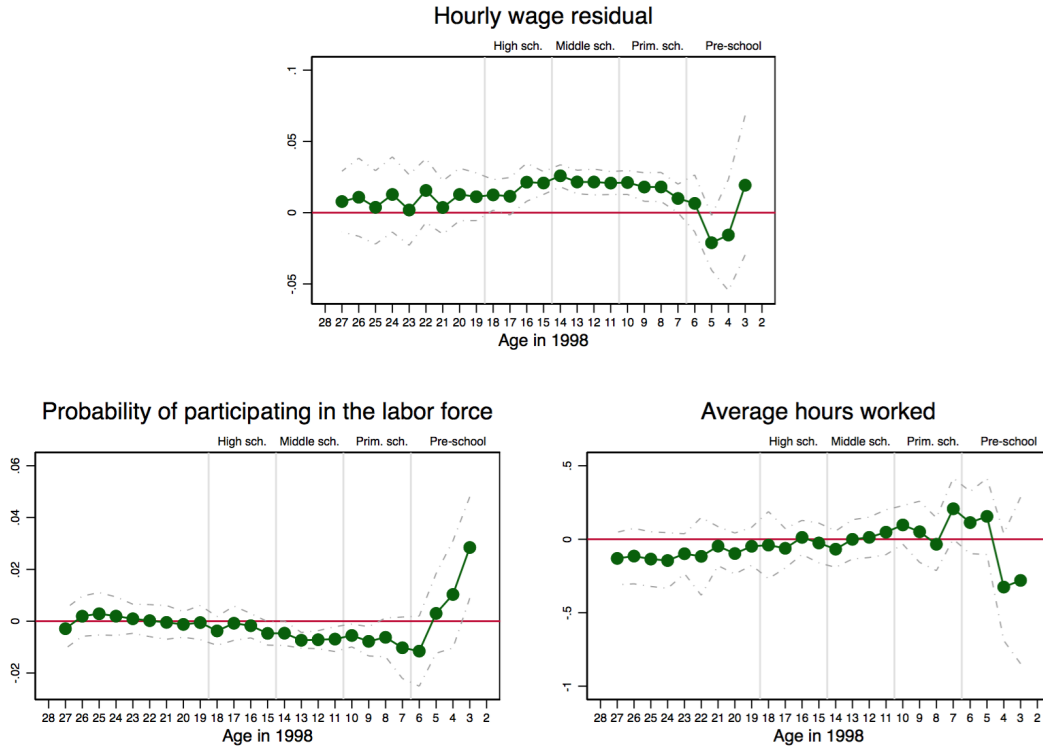


Figure 3.5: *Effects of FUNDEF on wages and labor force participation*

Note: The markers represent the coefficient on the interaction of the FUNDEF treatment variable and each cohort dummy in equation 3.3. Dashed lines are 95% confidence intervals, with standard errors clustered at the municipality of education level. **Sources:** See data appendix.

In contrast with the first-stage estimates discussed in section 3.4, when looking at labor market outcomes there is no expectation that the coefficients for cohorts that were in theory not exposed to FUNDEF should necessarily be zero. This is because all cohorts work in the same labor markets. Workers from different cohorts may be complements or substitutes in production, or increases in education in a subset of workers may lead to positive or negative spillovers on others.

FUNDEF had a positive impact on wages for most cohorts exposed to the program. As shown in the top graph in Figure 3.5, the only exceptions are the cohorts aged 4 through 7 in 1998. The cohort aged 4 in 1998 turned working age in 2009, a recession year, which may explain this pattern. Significant effects on individuals older than middle-school age in 1998 could be due to spillovers or, more plausibly, by late school entrance and high

Table 3.4: *Effects of FUNDEF on individual labor market outcomes*

	All	By gender			By migrant		
	(1)	Male (2)	Female (3)	Test (4)	Non-mig. (5)	Migrant (6)	Test (7)
				(F-stat and p-val.)			(F-stat and p-val.)
Hourly wage	0.017*** (0.002)	0.028*** (0.003)	0.003 (0.004)	56.44 0.000	0.008*** (0.002)	0.052*** (0.004)	117.59 0.000
Monthly wage	0.019*** (0.001)	0.031*** (0.003)	0.003 (0.004)	88.32 0.000	0.009*** (0.002)	0.055*** (0.003)	182.03 0.000
Labor force participation	-0.005*** (0.001)	-0.004*** (0.001)	-0.006*** (0.001)	3.53 0.060	-0.007*** (0.001)	0.004*** (0.002)	35.88 0.000
Weekly hours worked	0.116*** (0.028)	0.156*** (0.036)	0.115** (0.049)	0.52 0.469	0.090*** (0.031)	0.109 (0.068)	0.07 0.796
Formality	-0.004*** (0.001)	-0.002 (0.001)	-0.009*** (0.002)	16.83 0.000	-0.006*** (0.001)	0.003 (0.002)	16.80 0.000
Informality	0.003*** (0.001)	0.002 (0.001)	0.005*** (0.001)	3.81 0.051	0.005*** (0.001)	-0.003* (0.002)	12.86 0.000
Unemployment	0.001* (0.001)	-0.000 (0.001)	0.004*** (0.001)	11.05 0.001	0.002* (0.001)	0.001 (0.001)	0.35 0.552
Cohort of birth dummies	Yes	Yes	Yes		Yes	Yes	
Enrollment x cohort controls	Yes	Yes	Yes		Yes	Yes	
Demographic structure controls	Yes	Yes	Yes		Yes	Yes	
Region of work fixed effect	Yes	Yes	Yes		Yes	Yes	

Note: Columns 1 through 3 reports linear probability estimates of $\hat{\beta}_1$ in equation 3.5. Column 4 reports results of adjusted Wald tests of the hypothesis $H_0 : \beta_{males} - \beta_{females} = 0$ on SUR models' coefficients in columns 2 and 3. Columns 5 through 7 report estimates of $\hat{\beta}_1$ in equation 3.6 where the instruments are the interactions of the FUNDEF shock in the municipality of education and the cohort fixed-effects. All regressions use weighting based on sample design. Robust standard errors in parentheses.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

repetition rates. Similar results are obtained for alternative wage measures, and are reported in Appendix Figure C.4.

Table 3.4 provides measures of the average effects (across exposed cohorts) based on equation 3.5. One percent increase in the education budget in the individual's municipality of education led, on average, to a 1.7% increase in hourly wages and a 1.9% increase in monthly wages for individuals who were in theory exposed to the program, relative to those who were not. Jackson *et al.* (2016) also find positive effects of K-12 education spending on wages in the U.S. context.

I find a sharp gender difference in the wage effect. While a one percent FUNDEF shock led to a 2.8% increase in male hourly wage, the effect on female wages was statistically non-distinguishable from zero. Such gender differences are only significant in the generations that experienced a positive wage effect, as shown in Appendix Figure C.5. These results are consistent with those of Chapter 2, where I find that over the 1990s and 2000s, males wages grew when local labor demand increased while female wages did not.

FUNDEF had, on average, a negative effect on labor force participation. A one percent FUNDEF shock was associated with an average 0.5 percentage points reduction in the probability of participating in the labor force for the exposed cohorts. The effect was notoriously different for the cohorts that were fully exposed to the program (ages 6 and younger in 1998), for whom exposure is associated with increased labor force participation. The gender differences in the effects also vary across cohorts. While, on average, women exposed to the program decreased their participation more than men exposed to the program (Table 3.4), among the younger cohorts women had a larger participation increase (Appendix Figure C.5).

In contrast, I find a significant positive effect on the intensive margin of participation (average hours worked) for both men and women. Again, this net positive effect has a very different explanation for each gender. As depicted in Appendix Figure C.5, the positive effect among males reflects a decline in participation of the older generation (those not affected by the program). Among females, it reflects an increase in the number of hours worked of the generations affected by FUNDEF.

Conditional on participating in the labor force, the program decreased the probability of becoming formally employed. A one percentage point higher FUNDEF shock was associated with an average 0.4 percentage points reduction in the probability of formal employment. Part of this change was absorbed by an increased probability of informal employment (0.3 percentage points), and part by an increased probability of unemployment (0.1 percentage point.)

Virtually all of these other negative effects on employment outcomes are driven by

women. I find no measurable effect on males' formality or unemployment rates. In the women-only sample, the program is associated with a 0.9 drop in the probability of formal employment, matched with a 0.5 percentage points increase in the probability of informal employment and a 0.4 increase in the probability of unemployment.

3.4.4 Mechanisms

A possible explanation for the positive effect of FUNDEF on average wages and other labor market outcomes is that higher education attainment made workers more productive, leading to higher incomes. Prior research has found large returns to education in developing countries. Duflo (2001) showed that a large school construction program in Indonesia in the 1970s led to an increase in the average years of schooling of the population exposed to the program, and wage returns to a year of schooling in the range of 6.8 to 10.6 percent in a sample restricted to males.¹⁷

An alternative explanation is that educated workers obtained higher incomes because they became able to move to more productive places. Internal mobility in Brazil is relatively high,¹⁸ and we have already seen that the program led to out-migration among its beneficiaries (Section 3.4.4.) Moreover, multiple studies have attributed at least part of the wage premium of migrants to the characteristics of their destination place. Glaeser and Maré (2001) showed that moving to cities gives workers both a static and a dynamic wage effect, so that there urban wage premium accrues over time for workers who live in MSAs, and remains with them after they leave. De la Roca and Puga (2017) found similar results in Spain, where workers who move to larger cities have a discrete increase in wages upon migrating, and accumulate human capital at a faster pace than workers that stayed in smaller cities. Clemens (2013), using the U.S. visa lottery as a source of exogenous location

¹⁷Although Indonesia was not as close to achieving universal primary education in the 1970s as Brazil was in the 1990s, the INPRES school construction program allocation rule also prioritized regions that had the highest non-enrolled school-age population (Duflo 2001.)

¹⁸Even though in Brazil internal mobility had slowed down relative to the prior three decades, it was still high over the period of interest. Between 2000 and 2010, 10.35% of the adult population changed microregions of residence (Chapter 1)

for employees of a software firm, found large wage differences between programmers that stay in India and those who migrate to the U.S., which seems to be derived exclusively from the location.

In an effort to tell apart these alternative explanations, I start by estimating the effects of FUNDEF on labor market outcomes separately for migrants and non-migrants. Figure 3.6 displays the results of the cohort-specific regression (equation 3.3) for wages and labor force participation. I find that, among the cohorts exposed to the program, wages increased for migrants and decreased for non-migrants. The gap widens in younger generations, who were in principle more exposed to the program. I find very similar patterns using alternative measures of wages (Appendix Figure C.6.) In most cohorts, I also find a larger effect on the extensive margin of labor force participation among migrants than among non-migrants.

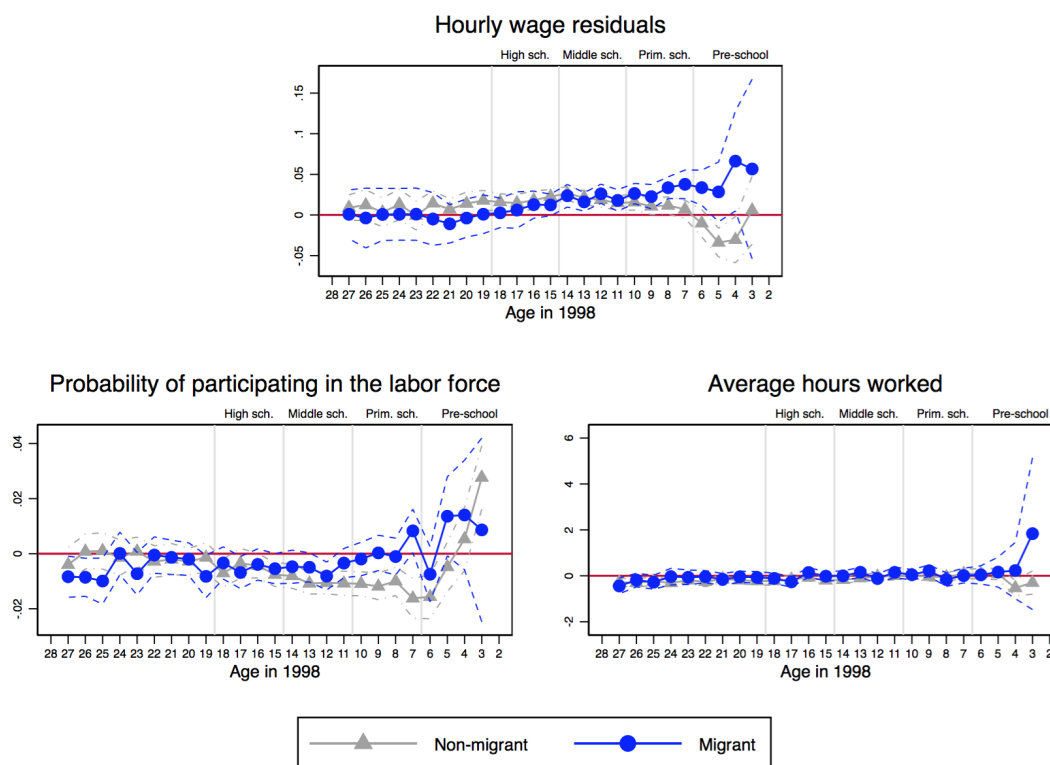


Figure 3.6: *Effects of FUNDEF on wages and labor force participation by migrant status*

Note: The markers represent the coefficient on the interaction of the FUNDEF treatment variable and each cohort dummy in equation 3.3. Dashed lines are 95% confidence intervals, with standard errors clustered at the municipality of education level. **Sources:** See data appendix.

Across cohorts, I find that the effects of the program on labor market outcomes were systematically better for migrants than for non-migrants. Columns 5 and 6 in Table 3.4 report estimates of the average treatment effects on the exposed cohorts from equation 3.5, calculated separately for migrants and non-migrants. Column 7 in the same table reports tests of differences of the coefficients of the two groups. I find that a one percent increase in the education budget in the municipality of education led, on average, to a 0.8% increase in individual hourly wages for non-migrants, and to a 5.2% increase for migrants. The difference in the coefficients is highly statically significant. Migrants also had a positive effect on labor force participation (as opposed to a negative effect for non-migrants) and a negative effect on informality (which contrasts with a positive effect for non-migrants.)

Interestingly, the sharp gender differences in labor market outcomes' effects discussed in Section 3.4.3 appear to be largely orthogonal to the differences between migrants and non-migrants. For females and non-migrants FUNDEF implied worse labor market outcomes relative to males and migrants, respectively. But as discussed in Section 3.4.2, the effect of the program on the likelihood to migrate was not statistically different between men and women. In other words, it appears that females obtained lower labor market effects from FUNDEF not *because* they migrated less, but *in spite* of migrating at similar rates. A possible explanation for this result is the presence of male-biased joint mobility decisions. In Chapter 2 I find that married couples in Brazil during this period were more likely to migrate in response to better labor market prospects for men than for women. Tied-migrant women, consequently, were more likely than men to locate in regions with weak job prospects for their human capital levels. This may be an important hurdle for the ability of women to turn their increasing education levels into better job market outcomes, specially given that the majority of married women in Brazil have a partner that has a lower educational attainment (Ganguli *et al.* 2014.)

The fact that I find significantly better effects on labor market outcomes for migrants could, in turn, be explained by differences in individual characteristics or by differences in characteristics of the place of work of migrants and non-migrants.

A large literature has studied the connection between education and the geographic sorting of workers. Educated individuals are relatively more mobile (Notowidigdo 2013), and when they migrate they are more likely to go to larger (Combes *et al.* 2008; Glaeser and Resseger 2010), more distant (Wozniak 2010), and more educated places (Berry and Glaeser 2005; Diamond 2016.) In addition, return migrants (De la Roca 2017) and migrants to smaller cities (Combes *et al.* 2012b) tend to be negatively selected.

I find that, during the period of interest, migrants in Brazil did have higher observable human capital characteristics. In 2010, the migrant population had, on average, higher educational attainment than the non-migrant population (Appendix Figure C.7.) While among the former 59% had middle school or higher education, this number was 56% among the latter. I also find a larger migrant-non migrant gap in the wage effect when I use wage measures that do not control for observable individual characteristics (Figure C.6). These patterns are consistent with previous literature documenting that internal migrants in Brazil, as in many other context, are positively selected (Dos Santos Júnior *et al.* 2005; Freguglia and Menezes-Filho 2012.)

In addition to this selection on observables, migrants could also be selected on characteristics that are hard to observe or unobservable.¹⁹ It is possible that the observable and unobservable characteristics that drive sorting are strongly correlated and accounted for by the controls. Prior literature has shown that, while migrants to larger cities in the U.S. and Spain are positively selected on schooling and other observed characteristics, there is little evidence of sorting on unobserved characteristics, as captured by individual fixed effects (Baum-Snow and Pavan 2012; De la Roca and Puga 2017). However, if these unobserved characteristics are not accounted for by my controls, my estimates of the effects of FUNDEF on the wage of migrants may have an upward bias.

¹⁹For instance, looking at data from Project STAR, a well-studied experiment that randomly assigned kindergarten students in Tennessee to classrooms with different characteristics in the mid 1980s, Chetty *et al.* (2011) find that the likelihood of living out of state as adults was positively associated with kindergarten test scores.

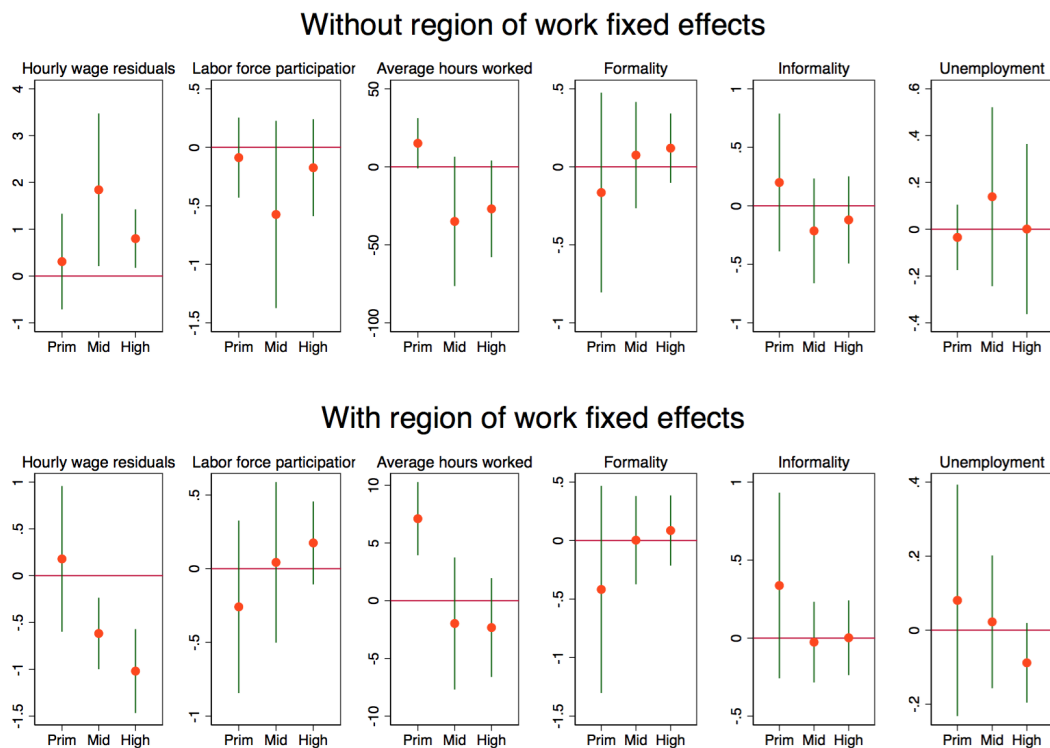


Figure 3.7: *Effects of education attainment on labor market outcomes (2SLS estimates)*

Note: The figure reports 2SLS coefficients on different levels of education attainment using the interaction of cohort fixed-effects and the intensity of FUNDEF transfers in the municipality of education as instruments. Markers denote coefficients and lines denote 95% confidence intervals. **Sources:** See data appendix.

The census data does not allow me to control for individual-level fixed effects. Instead, I look at what happens to the returns to education when I account for characteristics of the place of work of migrants. To this effect, I produce 2SLS estimates of the returns to educational attainment (coefficient $\hat{\beta}_1$ in equation 3.6) using two different specifications. In the first specification I do not control for place of work fixed effects. The returns captured by these estimates reflect both any increase in individual productivity and any gains from relocating to more productive places. In the second specification, I control for region of work fixed effects, shutting down the variation coming from potential productivity differences across localities. Figure 3.7 reports the results of these estimations for the three attainment levels affected by FUNDEF. In addition to wages, it includes 2SLS estimates to other labor market outcomes of interest.

The results are consistent with the interpretation that the positive connection between exposure to FUNDEF and individual wages is derived from the productivity of the places where the beneficiaries worked in 2010, rather than from increases in the productivity of the individual workers. In the specification without fixed effects I estimate positive and significant returns to middle school attainment and to high school attainment. However, when I control for time-invariant characteristics of the place of work, I obtain *negative* and statistically significant estimates. If unobserved individual characteristics were the key drivers of the positive wage effects of education attainment, it would be hard to explain why the estimates of these effects turn negative with the introduction of region of work fixed effects.

The only other outcome for which I observe a statistically significant effect of educational attainment is the average hours worked per week. Achieving primary school or higher attainment is associated with an weekly increase of 15 work hours worked in the specification that does not control for region of work characteristics, and of 7 hours in the specification that does. I do not find equivalent effects for middle school and high school achievement.

Finding negative returns to education in a period where employment was increasing suggests that the growth in the supply of educated workers outpaced demand growth (Pritchett 2001). This is in line with Andrade and Menezes-Filho (2005) who find that, during the 1980s and 1990s, the increase in the relative supply of middle-education workers in Brazil outpaced growth in their relative demand, while demand for high-education workers remained stable, and the relative supply of the least educated workers decreased, driving the relative wage increase.²⁰

My returns to education estimates could be biased if the program affected not only the quantity but also the quality of local education. A long-standing literature has documented that school quality can affect both returns to education and educational attainment levels (Card and Krueger 1992; Heckman *et al.* 1996; Deming *et al.* 2014.) Hanushek and Woessmann

²⁰In this work low-education workers as defined as having less than primary school, the middle-education workers as having at least primary and up to high-school, and the high-education workers as having at least one year of college.

(2012) find that differences in quality of education explain why Latin America trailed other world regions in terms of economic development, in spite of having higher initial attainment levels.

Whether the effects of local public education investments on the outcomes of interest come from changes in quantity or quality of education is much harder to identify. The literature has failed to find a systematic relationship between additional resources and the quality of schooling (Hanushek 1997, 2003.) In theory, the program could have deteriorated educational quality, introducing an downward bias. In its initial years of implementation, FUNDEF was associated with both increases in total enrollment and decreases in the total number of schools -as state-run schools closed- leading do higher average class size. Moreover, municipalities had some discretion on the nature of their education investments, and whether these emphasized quantity or quality may be endogenous.²¹

However, existing evaluations of the effects of FUNDEF find that, on the net, the program had a *positive* effect on quality through increasing the total number of public teachers, their wages, the availability of funds for their training (Menezes-Filho and Pazello 2007).²² These findings suggest that quality is unlikely to be behind the negative returns estimates.

3.5 Regional-level results

I turn now to the analysis of regional-level effects. I start by assessing the effect of FUNDEF on aggregate local educational attainment. Second, I discuss the effects of the program on migration, and how in turn they may shape the education composition of the local labor force. Finally, I explore the effects of the program on aggregate local labor market outcomes.

²¹Katrina Kosec (2014) finds that municipalities with higher median income and higher inequality spent less of the program's revenues in expanding public school enrollment. Rather, they were more likely to invest in public infrastructure with general-public use (e.g. roads and parks).

²²This contrasts with recent experimental evidence from Indonesia, which finds that increases in teacher wages led to higher teacher satisfaction but had no impact on learning outcomes of students (de Ree *et al.* 2018.)

3.5.1 Effects on regional educational attainment levels

The fact that FUNDEF did increase educational attainment among individuals that were exposed to the policy does not necessarily imply that we will observe an increase in the local education levels of the places that benefited from the program. In a context where there is free internal mobility, as in Brazil, individual beneficiaries may choose to migrate to other locations in search of better economic opportunities (Andrews 2017; Abel and Deitz 2012). Moreover, if the program did increase local education levels, and that in turn increased local productivity and labor demand (as in Moretti 2004) it could have attracted workers from different regions. In that case the education attainment of immigrants may have, in turn, contributed to shaping the aggregate education levels of the local economy.

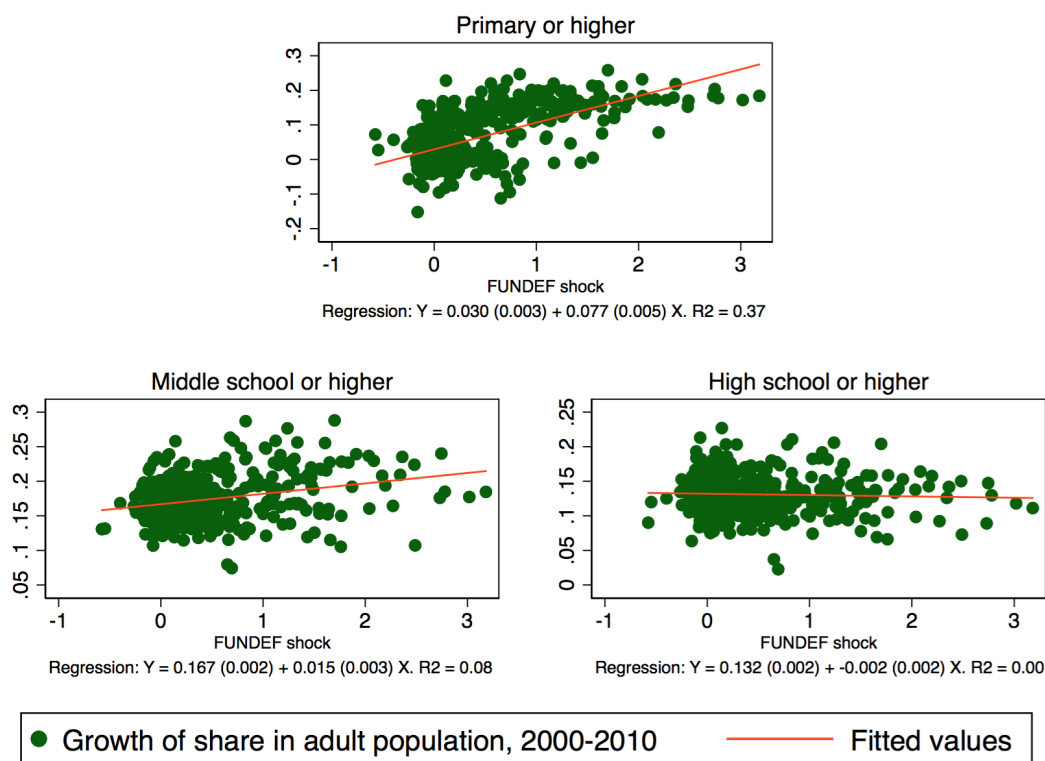


Figure 3.8: *Effects of FUNDEF on growth of the share of educated people among adults*

Note: Observations are microregions in which all municipalities have data on FUNDEF shock (N=456). **Sources:** See data appendix.

Regions with higher incidence of the program did see rising regional education levels

in the 2000s. Figure 3.8 shows the simple correlation between the regional-level FUNDEF shock and the growth of aggregate education levels. It uses as local education measures the share of the “educated” in the adult population for three different categories of educational attainment: primary school or higher, middle-school or higher, and high-school or higher. The program appears to have been particularly effective at increasing the share of adults completing primary education. On average, a one percentage point increase in the education budget -which corresponds to 1.63 standard deviations- was associated with a 7.7 percentage points increase in the share of individuals with primary (or higher) education in the adult population -equivalent to a 0.44 standard deviations reduction-. The program had a weaker correlation with higher education attainment measures.

In order to explore to what extent this relationship can be interpreted as causal, I turn to a difference-in-differences regression set-up. Table 3.5 reports the coefficients on the interaction between the FUNDEF shock and the “after” period (2010) dummy in equation 3.4. The difference in differences technique identifies the average treatment effect on the treated. In this context, assuming that the parallel trends assumptions holds conditional on controls, the coefficient α_3 in equation 3.4 is an estimate of the average treatment effect of increased local public education investments on the beneficiary regions’ outcomes of interest. Note that this estimate reflects both the direct effect of increased relative supply of local educated labor -i.e. the program’s effect on moving a share of the local population from a low education category to a high education category- and any general equilibrium effects -e.g. effects of local education levels on labor demand and subsequent migratory adjustments (Moretti, 2011)-.

The difference in differences estimation yields estimates that are very close to the coefficients of the simple OLS regression. A one percentage point larger increase in FUNDEF transfers was associated with a 7.5 percentage points increase in the share of individuals with primary education or higher. In this case, the point estimates are fairly similar for the sample restricted to males than for the sample restricted to females. The estimates for higher education levels are smaller, specifically of 1.3 percentage points for the share of

Table 3.5: *Effects of FUNDEF on local education attainment*

	Change in share of educated in adult population 2000-2010			1991-2000 (placebo test)		
	Primary (1)	Mid-school (2)	High-school (3)	Primary (4)	Mid-school (5)	High-school (6)
All individuals	0.075*** (0.008)	0.013** (0.006)	-0.004 (0.004)	0.004 (0.005)	-0.030*** (0.004)	-0.023*** (0.003)
Males only	0.081*** (0.009)	0.005 (0.006)	-0.010** (0.004)	0.002 (0.005)	-0.030*** (0.004)	-0.022*** (0.003)
Females only	0.070*** (0.007)	0.021*** (0.007)	0.003 (0.005)	0.006 (0.005)	-0.030*** (0.004)	-0.025*** (0.003)
Effective F statistic	117.31	9.75	0.52			
Weak instrument test critical value	23.11	23.11	23.11			
Changes in formality rates and wages in the 1980s	Yes	Yes	Yes	Yes	Yes	Yes

Note: The table reports the coefficients on the *Post* \times *Treatment* interaction in equation 3.4. Regressions are at the microregion level (N=456). Robust standard errors clustered at the mesoregion level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

middle school or higher, and of a non-significant negative 0.4 percentage points for the share of high school or higher. At the middle-school margin, the effect is driven by the female population, and at the high-school margin the negative effect is driven by the male population.

To explore the validity of the parallel trends assumption, replicate the same analysis using 2000 as the “post” period and 1991 as the “pre-period”. The results are reported in columns 4 through 6 of Table 3.5. After controlling for local labor market trends in the 1980s, the FUNDEF shock appears to be largely uncorrelated with the 1990s trends in the share of the primary educated in the population. However, the same is not true in the case of the 1990s trends in the shares of middle school and high school educated population. Regions with high FUNDEF transfers were also regions in where the shares of adults with higher education levels was declining during the 1990s.

The fact that the parallel trends assumption appears to hold - conditional on controls - for primary education but not for higher education levels is puzzling. As discussed in Section 3.3, FUNDEF targeted low-enrollment regions, and it is reasonable to expect that program intensity correlates with prior local trends on educational attainment. Many

Brazilian regions that saw deteriorating labor market conditions during the 1980s and 1990s, experienced an economic recovery during the 2000s. Deteriorating conditions in the prior decades may have in turn led to lower enrollment rates in 1998. This motivates the use of 1980s trends controls in my preferred specification. But while conditioning on these variables accounts for the correlation of the program intensity with pre-trends in the share of primary education, it fails to do the same for the cases of middle school and high school education. A possible explanation for this difference relates to the effects of the program on the immigration of educated workers, which I explore next.

3.5.2 FUNDEF, migration, and the educational attainment levels of regions

A likely source of endogeneity of the shares of middle-school and high-school educated is the potential effect of the program on the *demand* for qualified workers. Andrabi *et al.* (2013) argue that regions with initially low education levels face subsequent low supply of local population qualified to teach. The authors document, in the context of Pakistan, that the construction of government girls' secondary schools was associated with a higher likelihood of private schools presence in the following years. The introduction of FUNDEF in 1998 increased the availability of funds specifically earmarked for teacher wages in beneficiary regions, and migrants may have filled at least part of the unmet demand. Controlling for the volume and education of composition of migrants during the 1990s, makes the FUNDEF shock uncorrelated with pre-trends in shares of middle school and high-school educated (Appendix Table C.7). Moreover, in Brazil - as in Pakistan - women play a prominent role as teachers, which may explain why the correlation of the program with the growth in the shares of middle- and high school educated in the 2000s is noticeably larger for females than for males.

To further explore the role of migration in the composition of local human capital following FUNDEF, Table 3.6 reports difference in differences estimates of the effect of the program on population growth for different education attainment groups. A long-standing literature has documented a strong connection between initial education *levels* and subsequent population growth in U.S. cities (Glaeser *et al.* 1995; Glaeser and Shapiro 2003; Shapiro

Table 3.6: *Effects of FUNDEF on regional population*

	Log of population			Shares of education group in population		
	All (1)	Males (2)	Females (3)	All (4)	Males (5)	Females (6)
<u>Panel A: Effects on aggregate population</u>						
All adult population	0.015 (0.028)	0.013 (0.027)	0.018 (0.029)			
<u>Panel B: Effects on population by education group</u>						
Less than primary	-0.085*** -0.025	-0.087*** -0.024	-0.085*** -0.026	-0.026*** -0.006	-0.025*** -0.006	-0.026*** -0.006
Primary or higher	0.191*** (0.037)	0.218*** (0.038)	0.168*** (0.036)	0.075*** (0.008)	0.081*** (0.009)	0.070*** (0.007)
Middle school or higher	0.256*** (0.046)	0.266*** (0.048)	0.250*** (0.045)	0.013** (0.006)	0.005 (0.006)	0.021*** (0.007)
High school or higher	0.247*** (0.050)	0.258*** (0.052)	0.242*** (0.049)	-0.004 (0.004)	-0.010** (0.004)	0.242*** (0.049)
College or higher	0.354*** (0.066)	0.230*** (0.068)	0.419*** (0.073)	-0.011*** (0.001)	-0.009*** (0.001)	-0.013*** (0.001)

Note: The table reports the coefficients on the *Post* \times *Treatment* interaction in equation 3.4. Regressions are at the microregion level (N=456). Robust standard errors clustered at the mesoregion level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

2006). Chomitz *et al.* (2005) and Chapter 1 find a similar correlation in Brazil. However, I find a small and statistically non-significant connection between FUNDEF-induced local public education investments and aggregate population growth in microregions during the 2000s.

But while overall mobility towards FUNDEF-intensive regions was not systematically different than regions with low program intensity, the changes in the education composition were. The overall population with less than primary education shrank in beneficiary regions, while the population in the other educational categories increased. The increases in primary and middle-school educated population - the schooling levels targeted by the program - outpaced growth in other education categories, and the share of these education groups grew following the program, while the shares of the other groups shrank. The fact that the

share of primary-educated grew the most in spite of the fact that FUNDEF was associated with a *negative* net migration in this education category (Appendix Table C.6), is consistent with the individual-level findings showing that the program's largest impact on educational attainment was at this level (Section 3.4).

Gender differences by education group are consistent with the interpretation that FUNDEF had a direct impact on the *demand* of workers with intermediate education. While program intensity in the region was associated with a significant increase in the share of middle school and high school educated among women, the effect was small and non-significant for middle school, and negative for high-school in the case of men. Furthermore, the effects of the program on migration is positive and significant for females, and close to zero and non-significant for males in these two education categories (Appendix Table C.6.)

3.5.3 FUNDEF and regional labor market outcomes

I turn now to the effects of FUNDEF on labor market outcomes at the regional level. Table 3.7 summarizes the results of regional-level regressions exploring six aggregate labor market outcomes. Panel A reports difference in differences estimates of the reduced-form effects of the program on the outcomes (equation 3.4.) Panels B and C explore the effects of changes in local education attainment levels on local labor market outcomes, estimating regressions of the form:

$$\Delta_{2000s}Y_r = \gamma_0 + \gamma_1\Delta_{2000s}Prim_r + \gamma_2(\Delta_{1980s}C_r) + \epsilon_r \quad (3.7)$$

where Δ denote decade-long changes, Y_r is the regional-level outcome of interest, $Prim_r$ is the share of primary-educated in region r , and $\Delta_{1980s}C_r$ are the same lagged trends controls used in my estimates of equation 3.4. Panel B reports OLS estimates of $\hat{\gamma}_1$, and Panel C reports 2SLS estimates of the same coefficient using the regional-level FUNDEF shock (equation 3.2) as instrument.

My choice of explanatory variable y informed by the findings in prior sections. Changes in local education levels at the primary education margin capture the level at which the program had the strongest impact, and are uncorrelated with observable pre-trends

Table 3.7: *Effects of FUNDEF on local labor market outcomes*

	Employed population (1)	Hourly wage res. (2)	Particip. (3)	Hours worked (4)	Formal (5)	Informal (6)	Unemp. (7)
<u>Panel A: Reduced-form relationship</u>							
All individuals	-0.029 (0.034)	-0.031* (0.018)	-0.018*** (0.005)	-0.486*** (0.176)	-0.024*** (0.005)	0.005 (0.006)	0.019*** (0.003)
Males only	-0.123*** (0.040)	-0.048** (0.020)	-0.012** (0.005)	-0.439*** (0.161)	-0.019*** (0.006)	0.002 (0.007)	0.017*** (0.003)
Females only	-0.009 (0.042)	-0.007 (0.016)	-0.023*** (0.006)	-0.487** (0.229)	-0.038*** (0.005)	0.021*** (0.007)	0.016*** (0.005)
<u>Panel B: Effects of changes in the share of primary educated , OLS</u>							
All individuals	-0.005 (0.104)	-0.490*** (0.097)	-0.211*** (0.030)	-0.106*** (0.032)	-0.235*** (0.037)	0.098** (0.046)	0.138*** (0.027)
Males only	0.115 (0.104)	-0.573*** (0.097)	-0.073** (0.033)	-0.070*** (0.026)	-0.153*** (0.033)	0.044 (0.043)	0.109*** (0.025)
Females only	0.093 (0.151)	-0.290*** (0.102)	-0.336*** (0.047)	-0.119** (0.055)	-0.396*** (0.057)	0.290*** (0.069)	0.106*** (0.037)
<u>Panel C: Effects of changes in the share of primary educated , 2SLS</u>							
All individuals	0.079 (0.195)	-0.380* (0.219)	-0.226*** (0.052)	-0.225*** (0.060)	-0.293*** (0.054)	0.049 (0.072)	0.244*** (0.051)
Males only	0.090 (0.186)	-0.522** (0.215)	-0.146** (0.065)	-0.195*** (0.049)	-0.217*** (0.056)	0.006 (0.073)	0.211*** (0.046)
Females only	0.400 (0.270)	-0.031 (0.226)	-0.294*** (0.057)	-0.226** (0.094)	-0.502*** (0.072)	0.276*** (0.093)	0.226*** (0.068)

Note: Panel A reports the coefficients on the *Post* \times *Treatment* interaction in equation 3.4. Panel B reports OLS estimates of the coefficient on the change in the share of individuals with at least primary school in the adult population in equation 3.7. Panel C reports 2SLS estimates of the same coefficient using the regional-level FUNDEF shock (equation 3.2) as instrument. Regressions are at the microregion level (N=456). Robust standard errors clustered at the mesoregion level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

conditional on controls. In Table 3.5, where I report the first-stage results discussed in Section 3.5.1, I also include the results of the test for weak instruments of Montiel Olea and Pflueger (2013). The regional FUNDEF shock is a strong instruments for this explanatory variable, but not for measures of changes in local education at higher levels.

The results show that, on average, labor market outcomes worsened in regions that benefited to FUNDEF. A program-induced one percentage point larger public education budget was associated with a 3.1% reduction of the average hourly wage, after controlling for individual characteristics. This is in spite of the fact that the program had a direct positive effect through its mandated increases in teachers' wages.²³ The program was also associated with lower participation (in the extensive and the intensive margins), higher informality rates, and higher unemployment. The decline in wages was stronger among males, and the increase in informal employment among females.

The evidence is consistent with the interpretation that local supply of educated labor outpaced local demand. Worsened labor market outcomes could be partially explained by negative selection among the non-migrants. However, while the IV estimates show a negative 0.3 percentage points effect on the hourly wage, the point estimate on employment is positive (although not statistically significant), suggesting a downward-sloping relatively inelastic demand. Local employment grew primarily in the informal sector and among women.

3.6 Conclusion

This paper explores the effects of public education investments on individual and regional labor market outcomes. Using Brazil's FUNDEF as a source of exogenous variation in local public education budgets I find generally positive effects at the individual level and negative effects at the regional level.

FUNDEF had a positive effect on individual educational attainment. Cohorts that were in

²³Following the introduction of the FUNDEF, teacher's salaries rose by an average of 13%, and in the poor north east increases were as high as 60% (OECD 2011.)

principle exposed to the program had, on average a 2.4 percentage points higher likelihood of attaining at least primary education relative to cohorts that were not exposed. The reform was less effective in increasing education attainment at other margins, with a 1 percentage point effect on middle school attainment, and a 0.4 percentage point effect on high school attainment.

The program also had a positive effect on individual wages, which was concentrated among individuals that migrated outside their region of education. One percent increase in the education budget in the individual's municipality of education led, on average, to a 1.7% increase in hourly wages and a 1.2 percentage points increase in the likelihood of migrating of individuals who were (in theory) exposed to the program. The wage effect was 5.2% for migrants, and only 0.8% for non-migrants. I estimate positive average returns to educational attainment for middle school and high school in the order of 1.9% and 0.8%, respectively, but these estimates become *negative* when I control for region of work fixed effects, suggesting that the bulk of the wage effect comes from characteristics of migrants' destination regions.

The results unveil large gender differences in the individual effects of local education spending. While the average wage effect for males was of 2.8%, the equivalent for women was close to zero. This gap is not explained by gender differences in migration elasticities. Joint location decisions that favor male over female labor market prospects may account for these patterns, but further research is required to better understand the mechanisms at play.

FUNDEF also led to higher educational attainment at the regional level, specially at the primary education education margin. But the increase in the share of educated workers was associated with worsening local labor market outcomes. The results on wages and employment suggest that growth in demand for educated labor was not large enough to absorb the program-related supply shifts.

References

- ABEL, J. R. and DEITZ, R. (2012). Do colleges and universities increase their region's human capital? *Journal of Economic Geography*, **12** (3), 667–691.
- ACEMOGLU, D. and ANGRIST, J. (2001). How Large are Human-Capital Externalities? Evidence from Compulsory Schooling Laws. *NBER Macroeconomics Annual 2000, Volume 15*, **15** (January), 9–74.
- , AUTOR, D., DORN, D., HANSON, G. H. and PRICE, B. (2016). Import Competition and the Great US Employment Sag of the 2000s. *Journal of Labor Economics*, **34** (S1), S141–S198.
- , — and LYLE, D. (2004). Midcentury Women, War, and Wages: The Effect of Female Labor Supply on the Wage Structure at Midcentury. *Journal of Political Economy*, **112** (3), 497–551.
- and DELL, M. (2010). Productivity Differences Between and Within Countries. *American Economic Journal: Macroeconomics*, **2**, 169–188.
- ADÃO, R. (2015). Worker Heterogeneity, Wage Inequality, and International Trade: Theory and Evidence from Brazil.
- AIZER, A. (2010). The Gender Wage Gap and Domestic Violence. *The American economic review*, **100** (4), 1847–1859.
- ALBANESI, S. and SAHIN, A. (2017). The Gender Unemployment Gap.
- ALBOUY, D. (2015). What Are Cities Worth? Land Rents, Local Productivity, and the Total Value of Amenities.
- and STUART, B. (2017). Urban Population and Amenities: The Neoclassical Model of Location.
- ALBRECHT, J., BJÖRKLUND, A. and VROMAN, S. (2003). Is There a Glass Ceiling in Sweden? *Journal of Labor Economics*, **21** (1), 145–177.
- ALONSO, W. (1964). *Location and Land Use*. Cambridge, MA: Harvard University Press, Cambridge, Mass.
- ALTONJI, J. G. and BLANK, R. M. (1999). Race and gender in the labor market. *Handbook of Labor Economics*, **3** (3), 3143–3259.
- ALVES, G. (2016). Slum Growth in Brazilian Cities.

- AMIOR, M. and MANNING, A. (2015). The Persistence of Local Joblessness.
- ANDERSON, G. and GE, Y. (2005). The size distribution of Chinese cities. *Regional Science and Urban Economics*, **35** (6), 756–776.
- ANDRABI, T., DAS, J. and KHWAJA, A. I. (2013). Students today, teachers tomorrow: Identifying constraints on the provision of education. *Journal of Public Economics*, **100**, 1–14.
- ANDRADE, A. A. S. D. and MENEZES-FILHO, N. A. (2005). O Papel da Oferta de Trabalho no Comportamento dos Retornos à Educação no Brasil. *Pesquisa e Planejamento Econômico*, **35** (2), 189–225.
- ANDREWS, M. (2017). The Role of Universities in Local Invention: Evidence from the Establishment of U.S. Colleges. *Mimeo*.
- ARRIAGA, E. (1968). Components of city growth in selected Latin American countries. *The Milbank Memorial Fund Quarterly*, **46** (2), 237–252.
- AU, C. C. and HENDERSON, J. V. (2006). How migration restrictions limit agglomeration and productivity in China. *Journal of Development Economics*, **80** (2), 350–388.
- AUSTIN, B., GLAESER, E. and SUMMERS, L. L. (2018). Saving the heartland: Place-based policies in 21st century America. *BPEA Conference Drafts*, pp. 1–96.
- AUTOR, D. H. and DORN, D. (2013). The Growth of Low Skill Service Jobs and the Polarization of the U . S . Labor Market. *American Economic Review*, **103** (5), 1553–97.
- , — and HANSON, G. H. (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review*, **103**, 2121–2168.
- and DUGGAN, M. G. (2003). The Rise in the Disability Rolls and the Decline in Unemployment. *Quarterly Journal of Economics*, **118** (1), 157–206.
- , KATZ, L. F. and KEARNEY, M. S. (2009). The Polarization of the U.S. Labor Market. *The American Economic Review*, **96** (2), 189–194.
- AZZONI, C. R., MENEZES-FILHO, N. and SILVEIRA-NETO, R. (2000). Geography and income convergence among Brazilian states. (May), 30.
- and SERVO, L. M. S. (2002). Education, cost of living and regional wage inequality in Brazil. *Papers in Regional Science*, **81** (2), 157–175.
- BADINGER, H. and TONDL, G. (2002). Trade, human capital and innovation: The engines of European regional growth in the 1990s. In *42nd Congress of the European Regional Science Association: "From Industry to Advanced Services - Perspectives of European Metropolitan Regions"*, Dortmund.
- BAKER, M., GRUBER, J. and MILLIGAN, K. (2008). Universal Childcare, Maternal Labor Supply, and Family Well-Being. *Journal of Political Economy*, **116** (4), 709–745.
- BARTIK, T. J. (1991). *Who Benefits from State and Local Economic Development Policies?* Kalamazoo: Books from Upjohn Press.

- (2004). Evaluating the Impacts of Local Economic Development Policies on Local Economic Outcomes: What has been done and what is doable? In *Evaluating local economic and employment development: how to assess what works among programmes and policies*, 4, Paris: OECD Publishing, pp. 113–142.
- (2015). How Effects of Local Labor Demand Shocks Vary with the Initial Local Unemployment Rate. *Growth and Change*, **46** (4), 529–557.
- BAUM-SNOW, N. and PAVAN, R. (2012). Understanding the city size wage gap. *Review of Economic Studies*, **79** (1), 88–127.
- BEAUDRY, P., GREEN, D. A. and SAND, B. M. (2014). Spatial equilibrium with unemployment and wage bargaining: Theory and estimation. *Journal of Urban Economics*, **79**, 2–19.
- BELL, M. and MUHIDIN, S. (2009). Cross-National Comparisons of Internal Migration.
- BERRY, C. and GLAESER, E. (2005). The Divergence of Human Capital Levels Across Cities. *Papers in Regional Science*, **84** (3), 407–444.
- BERTRAND, M., GOLDIN, C. and KATZ, L. F. (2010). Dynamics of the gender gap for young professionals in the financial and corporate sectors. *American Economic Journal: Applied Economics*, **2** (3), 228–255.
- , KAMENICA, E. and PAN, J. (2015). Gender Identity and Relative Income Within Households. *The Quarterly Journal of Economics*, **130** (2), 571–614.
- BLACK, D. A., KOLESNIKOVA, N. and TAYLOR, L. J. (2014). Why do so few women work in New York (and so many in Minneapolis)? Labor supply of married women across US cities. *Journal of Urban Economics*, **79**, 59–71.
- BLANCHARD, O. J. and KATZ, L. F. (1992). Regional Evolutions. *Brookings Papers on Economic Activity*, **1992** (1).
- BLAU, F. D. (2016). *Gender, Inequality, and Wages*. Oxford University Press.
- and KAHN, L. M. (2000). Gender Differences in Pay. *Journal of Economic Perspectives*, **14** (4), 75–100.
- and — (2003). Understanding International Differences in the Gender Pay Gap. **21** (1).
- BORUSYAK, K. and JARAVEL, X. (2017). Consistency and Inference in Bartik Research Designs.
- BOSKER, M., BRAKMAN, S., GARRETSEN, H. and SCHRAMM, M. (2008). A century of shocks: The evolution of the German city size distribution 1925–1999. *Regional Science and Urban Economics*, **38** (4), 330–347.
- BRÜLHART, M. and MATHYS, N. A. (2008). Sectoral agglomeration economies in a panel of European regions. *Regional Science and Urban Economics*, **38** (4), 348–362.
- BUETTNER, T. and EBERTZ, A. (2009). Quality of life in the regions: results from German Counties. *Annals of Regional Science*, **43**, 89–112.

- BUSO, M., GREGORY, J., KLINE, P. and BERKELEY, U. C. (2013). Assessing the Incidence and Efficiency of a Prominent Place Based Policy. *American Economic Review*, **103** (2), 897–947.
- CARD, D. and KRUEGER, A. B. (1992). Does School Quality Matter? Returns to Education and the Characteristics of Public Schools in the United States. *Journal of Political Economy*, **100** (1), 1–40.
- CARNEIRO, R. D. and KOVAK, B. K. (2015). Trade liberalization and the skill premium: A local labor markets approach. In *American Economic Review*, vol. 105, pp. 551–557.
- CAVAILHES, J., JOLY, D., CARDOT, H., HILAL, M., WAVRESKY, P. and BROSSARD, T. (2009). The price of climate: revealed preferences of French consumers. In *5th Urban research Symposium, Cities and climate change: responding to an urgent agenda*, Marseille.
- CAVALCANTI FERREIRA, P. (2004). Regional Policy in Brazil : A Review. *World Bank, commissioned paper*.
- CHAMPERNOWNE, D. (1953). A Model of Income Distribution. *The Economic Journal*, **63** (250), 318–351.
- CHAN, K. (2007). Misconceptions and Complexities in the Study of China's Cities: Definitions, Statistics, and Implications. *Eurasian Geography and Economics*, **48** (4), 383–412.
- CHEN, W.-c. (2012). How Education Enhances Happiness: Comparison of Mediating Factors in Four East Asian Countries. *Social Indicators Research*, **106** (1), 117–131.
- CHENG, Z., WANG, H. and SMYTH, R. (2014). Happiness and job satisfaction in urban China: A comparative study of two generations of migrants and urban locals. *Urban Studies*, **51** (10), 2160–2184.
- CHETTY, R., FRIEDMAN, J. N., HILGER, N., SAEZ, E., SCHANZENBACH, D. W. and YAGAN, D. (2011). How does your kindergarten classroom affect your earnings? Evidence from project star. *Quarterly Journal of Economics*, **126** (4), 1593–1660.
- CHOMITZ, K. M., DA MATA, D., DE CARVALHO, A. Y. and MAGALHÃES, J. C. (2005). Spatial dynamics of labor markets in Brazil. *World Bank Policy Research Working Paper*, (3752).
- CICCONI, A. (2002). Agglomeration effects in Europe. *European Economic Review*, **46** (2), 213–227.
- and HALL, R. E. (1996). Productivity and the Density of Economic Activity. *The American Economic Review*, **86** (1), 54–70.
- CLEMENS, M. A. (2013). Why do programmers earn more in Houston than Hyderabad? Evidence from randomized processing of US visas. *American Economic Review*, **103** (3), 198–202.
- COLLIER, P. and VENABLES, A. J. (2016). Urban infrastructure for development. *Oxford Review of Economic Policy*, **32** (3), 391–409.

- COMBES, P.-P., DÉMURGER, S. and LI, S. (2013). Urbanisation and Migration Externalities in China. *Centre for Economic Policy Research Discussion Paper*, (9352).
- COMBES, P. P., DURANTON, G. and GOBILLON, L. (2008). Spatial wage disparities: Sorting matters! *Journal of Urban Economics*, **63** (2), 723–742.
- COMBES, P.-P., DURANTON, G., GOBILLON, L., PUGA, D. and ROUX, S. (2012a). The productivity advantages of large cities: distinguishing agglomeration from firm selection. *Econometrica*, **80** (6), 2543–2594.
- , —, — and ROUX, S. (2010). Estimating agglomeration economies with history, geology, and worker effects. In E. L. Glaeser (ed.), *The Economics of Agglomeration*, vol. I, 1, The University of Chicago Press, pp. 15–65.
- , —, — and — (2012b). Sorting and local wage and skill distributions in France. *Regional Science and Urban Economics*, **42** (6), 913–930.
- and GOBILLON, L. (2015). The Empirics of Agglomeration Economies. In G. Duranton, J. Henderson and W. C. Strange (eds.), *Handbook of Regional and Urban Economics*, Volume 5, 5, North Holland, pp. 247–348.
- CORBI, R. B. and MENEZES-FILHO, N. A. (2006). Os determinantes empíricos da felicidade no Brasil. *Revista de Economia Política*, **26** (4), 518–536.
- CORSEUIL, C. H. and FOGUEL, M. N. (2002). Uma Sugestão de Deflatores para Rendas Obtidas a Partir de Algumas Pesquisas Domiciliares do IBGE. *Mercado de Trabalho: conjuntura e Análise - Nota Técnica*, (897), 13.
- , RAMOS, L., ARAÚJO, F. V. D. S. and POLOPONSKY, K. (2010). Um Panorama do Mercado de Trabalho Brasileiro em 2009. *IPEA Mercado de Trabalho*, **45**, 17–27.
- COSTA, D. L. and KAHN, M. E. (2000). Power Couples: Changes in the Locational Choice of the College Educated, 1940-1990. *The Quarterly Journal of Economics*, **115** (November), 1287–1315.
- COSTA, F., GARRED, J. and PESSOA, J. P. (2016). Winners and losers from a commodities-for-manufactures trade boom. *Journal of International Economics*, **102**, 50–69.
- CRUZ, G. and ROCHA, R. (2018). Efeitos do FUNDEF/B sobre Frequência Escolar, Fluxo Escolar e Trabalho Infantil: Uma Análise com Base nos Censos de 2000 e 2010. *Estudos Econômicos (São Paulo)*, **48** (1), 39–75.
- CUÑADO, J. and PÉREZ DE GRACIA, F. (2012). Does Education Affect Happiness ? Evidence for Spain. *Social Indicators Research*, **108** (1), 185–196.
- D' COSTA, S. and OVERMAN, H. G. (2014). The urban wage growth premium: Sorting or learning? *Regional Science and Urban Economics*, **48**, 168–179.
- DA MATA, D., DEICHMANN, U., HENDERSON, J. V., LALL, S. V. and WANG, H. G. (2007). Determinants of city growth in Brazil. *Journal of Urban Economics*, **62** (2), 252–272.

- DAHL, G. B. (2002). Mobility and the Return to Education: Testing a Roy Model with Multiple Markets. *Econometrica*, **70** (6), 2367–2420.
- DALMAZZO, A. and DE BLASIO, G. (2007). Production and consumption externalities of human capital: An empirical study for Italy. *Journal of Population Economics*, **20** (2), 359–382.
- DE BARROS, R. P., CARVALHO, M. D., FRANCO, S. and MENDONÇA, R. (2006). Uma análise das principais causas da queda recente na desigualdade de renda brasileira. *Texto para Discussão - IPEA*, **1203**.
- DE LA ROCA, J. (2017). Selection in initial and return migration: Evidence from moves across Spanish cities. *Journal of Urban Economics*, **100**, 33–53.
- DE LA ROCA, J. and PUGA, D. (2017). Learning by Working in Big Cities. *Review of Economic Studies*, **84** (1), 106–142.
- DE MELLO, L. and HOPPE, M. (2005). Education Attainment in Brazil: The Experience of FUNDEF. *OECD Economics Department Working papers*, (2011), 27.
- DE REE, J., MURALIDHARAN, K., PRADHAN, M. and ROGERS, H. (2018). Double for Nothing? Experimental Evidence on an Unconditional Teacher Salary Increase in Indonesia. *Quarterly Journal of Economics*, pp. 993–1039.
- DEMING, D. J., HASTINGS, J. S., KANE, T. J. and STAIGER, D. O. (2014). School Choice, School Quality, and Postsecondary Attainment. *American Economic Review*, **104** (3), 991–1013.
- DIAMOND, R. (2016). The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000. *American Economic Review*, **106** (3), 479–524.
- DIX-CARNEIRO, R. and KOVAK, B. K. (2016). Trade Liberalization and Regional Dynamics.
- DIX-CARNEIRO, R. and KOVAK, B. K. (2017). Trade Liberalization and Regional Dynamics. *American Economic Review* (forthcoming), pp. 1–102.
- DOS SANTOS JÚNIOR, E. D. R., MENEZES FILHO, N. A. and CAVALCANTI FERREIRA, P. (2005). Migração, Seleção e Diferenças Regionais de Renda no Brasil. *Pesquisa e Planejamento Econômico*, **35** (3), 299–331.
- DUFLO, E. (2001). Schooling and labor market consequences of school construction in Indonesia: Evidence from an unusual policy experiment. *American Economic Review*, **91** (4), 795–813.
- DURAND, J. D. and PELAEZ, C. (1965). Patterns of Urbanization in Latin America. In *Proceedings of the Sixtieth Anniversary Conference of the Milbank Memorial Fund. Part 2: Components of Population Change in Latin America.*, The Milbank Memorial Fund Quarterly, vol. 43, pp. 166–196.
- EDMONDS, E. V., PAVCNİK, N. and TOPALOVA, P. (2010). Trade adjustment and human capital investments: Evidence from indian tariff reform. *American Economic Journal: Applied Economics*, **2** (4), 42–75.

- ESTEVAN, F. (2015). Public education expenditures and private school enrollment. *Canadian Journal of Economics*, **48** (2).
- FALCK, O., FRITSCH, M. and HEBLICH, S. (2011). The phantom of the opera: Cultural amenities, human capital, and regional economic growth. *Labour Economics*, **18** (6), 755–766.
- FELER, L. and HENDERSON, J. V. (2011). Exclusionary policies in urban development: Under-servicing migrant households in Brazilian cities. *Journal of Urban Economics*, **69** (3), 253–272.
- FERNANDES, R. and DE FELICIO, F. (2005). The Entry of the Wife into the Labor Force in Response to the Husband's Unemployment: A Study of the Added Worker Effect in Brazilian Metropolitan Areas. *Economic Development and Cultural Change*, **53** (4), 887–911.
- FIESS, N. M. and VERNER, D. (2003). *Migration and human capital in Brazil during the 1990s*. Tech. Rep. 3093.
- FLEISHER, B., LI, H. and ZHAO, M. Q. (2010). Human capital, economic growth, and regional inequality in China. *Journal of Development Economics*, **92** (2), 215–231.
- FOGUEL, M. N. (2016). The Effects of Gender Segregation at the Occupation, Industry, Establishment, and Job-Cell Levels on the Male-Female Wage Gap.
- FONTES, G. G., SIMOES, R. F., DE OLIVEIRA, A. M. and HERMETO, C. (2009). Urban Attributes and Wage Disparities in Brazil: A Multilevel Hierarchical Model. *Regional Studies*, **44** (5), 596–607.
- FOSTER, N. and STEHRER, R. (2009). Sectoral Productivity, Density and Agglomeration in the Wider Europe. *Spatial Economic Analysis*, **4** (4), 427–446.
- FRANK, R. H. (1978). Why Women Earn Less: The Theory and Estimation of Differential Overqualification. *The American Economic Review*, **68** (3), 360–373.
- FREGUGLIA, R. S. and MENEZES-FILHO, N. A. (2012). Inter-regional wage differentials with individual heterogeneity: evidence from Brazil. *The Annals of Regional Science*, **49** (1), 17–34.
- FU, Y. and GABRIEL, S. A. (2012). Labor migration, human capital agglomeration and regional development in China. *Regional Science and Urban Economics*, **42** (3), 473–484.
- FULLER, B. and ROMER, P. (2014). Urbanization as Opportunity. *Marron Institute on Cities and the Urban Environment*, NYU.
- GABAIX, X. (1999). Zipf's Law for Cities: An Explanation. *The Quarterly Journal of Economics*, (August), 739–767.
- and IBRAGIMOV, R. (2011). Rank-1/2: A Simple Way to Improve the OLS Estimation of Tail Exponents. *Journal of Business & Economic Statistics*, **29** (1), 24–39.
- GALOR, O. and WEIL, D. N. (1996). The Gender Gap, Fertility, and Growth. *The American Economic Review*, **86** (3), 374–387.

- GANGULI, I., HAUSMANN, R. and VIARENGO, M. (2014). Marriage, education and assortative mating in Latin America. *Applied Economics Letters*, **21** (12), 806–811.
- GARIBALDI, P. and WASMER, E. (2005). Equilibrium Search Unemployment, Endogenous Participation, and Labor Market Flows. *Journal of the European Economic Association*, **3** (4), 851–882.
- GEMICI, A. (2011). Family Migration and LaborMarket Outcomes.
- GENNAIOLI, N., LA PORTA, R., LOPEZ DE SILANES, F. and SHLEIFER, A. (2014). Growth in regions. *Journal of Economic Growth*, **19**, 259–309.
- GERDTHAM, U.-G. and JOHANNESSON, M. (2001). The relationship between happiness, health, and socio- economic factors: results based on Swedish microdata. *The Journal of Socio-Economics*, **30** (6), 553–557.
- GIESEN, K. and SUDEKUM, J. (2011). Zipf’s law for cities in the regions and the country. *Journal of Economic Geography*, **11** (4), 667–686.
- GLAESER, E. and GOTTLIEB, J. D. (2006). Urban resurgence and the consumer city. *Urban Studies*, **43** (8), 1275–1299.
- and HENDERSON, J. V. (2017). Urban economics for the developing World: An introduction. *Journal of Urban Economics*, **98**, 1–5.
- GLAESER, E. L. (1999). Learning in cities. *Journal of urban Economics*, (46), 254–277.
- (2008). *Cities, Agglomeration and Spatial Equilibrium*. New York: Oxford University Press.
- and GOTTLIEB, J. D. (2008). The Economics of Place-Making Policies. *Brookings Papers on Economic Activity*, **2008** (1), 155–253.
- and — (2009a). The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States. *Journal of Economic Literature*, **47** (4), 983–1028.
- and — (2009b). The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States. *Journal of Economic Literature*, **47** (4), 983–1028.
- , — and ZIV, O. (2014a). Unhappy Cities.
- , GYOURKO, J. and SAKS, R. E. (2005). Urban growth and housing supply. *Journal of Economic Geography*, **6** (1), 71–89.
- and MARÉ, D. C. (2001). Cities and Skills. *Journal of Labor Economics*, **19** (2), 316–342.
- , PONZETTO, G. A. M. and TOBIO, K. (2014b). Cities, Skills and Regional Change. *Regional Studies*, **48** (1), 7–43.
- and RESSEGER, M. G. (2010). The Complementarity Between Cities and Skills. *Journal of Regional Science*, **50** (1), 221–244.

- , SCHEINKMAN, J. A. and SHLEIFER, A. (1995). Economic growth in a cross-section of cities. *Journal of Monetary Economics*, **36** (1), 117–143.
- and SHAPIRO, J. M. (2003). Urban growth in the 1990s: is city living back? *Journal of Regional Science*, **43** (1), 139–165.
- GLAESER, E. L. E. L. and SAIZ, A. (2004). The Rise of the Skilled City. *Brookings-Wharton Papers on Urban Affairs*, **2004** (1), 47–105.
- GOLDIN, C. (2006). The Quiet Revolution That Transformed Women’s Employment, Education, and Family. *The American Economic Review: Papers and Proceedings*, **96** (2), 1–21.
- (2014). A Pollution Theory of Discrimination. Male and Female Differences in Occupations and Earnings. In L. P. Boustan, C. Frydman and R. A. Margo (eds.), *Human Capital in History: The American Record*, October, University of Chicago Press, pp. 313–348.
- , KERR, S. P., OLIVETTI, C. and BARTH, E. (2017). The Expanding Gender Earnings Gap: Evidence from the LEHD-2000 Census. *American Economic Review*, **107** (May).
- GOLDSMITH-PINKHAM, P., SORKIN, I. and SWIFT, H. (2017). Bartik Instruments: What, When, Why and How.
- GOLLIN, D., KIRCHBERGER, M. and LAGAKOS, D. (2017). In Search of A Spatial Equilibrium in the Developing World.
- GONZAGA, G., MENEZES FILHO, N. and TERRA, C. (2006). Trade liberalization and the evolution of skill earnings differentials in Brazil. *Journal of International Economics*, **68** (2), 345–367.
- GORDON, N. and VEGAS, E. (2005). Educational finance, equalization, spending, teacher quality, and student outcomes: the case of Brazil’s FUNDEF. In E. Vegas (ed.), *Incentives to Improve Teaching: Lessons from Latin America*, 5, Washington D.C.: The World Bank.
- GREENWOOD, J., SESHADRI, A. and YORUKOGLU, M. (2005). Engines of Liberation. *Review of Economic Studies*, **72**, 109–133.
- GROOT, S. P. T., DE GROOT, H. L. F. and SMIT, M. J. (2014). Regional Wage Differences in the Netherlands: Micro Evidence on Agglomeration Externalities. *Journal of Regional Science*, **54** (3), 503–523.
- HAKOBYAN, S. and McLAREN, J. (2016). Looking for Local Labor Market Effects of NAFTA. *Review of Economics and Statistics*, **98** (4), 728–741.
- HANLON, W. W. (2015). Endogenous City Disamenities: Lessons from Industrial Pollution in 19th Century Britain.
- HANUSHEK, E. (1997). Assessing the effects of school resources on student performance: An update. *Educational evaluation and policy analysis*, **19** (2), 141.
- HANUSHEK, E. A. (2003). The failure of input-based schooling policies. *Economic Journal*, **113** (485), 64–98.

- HANUSHEK, E. A. and WOESSMANN, L. (2012). Schooling, educational achievement, and the Latin American growth puzzle. *Journal of Development Economics*, **99** (2), 497–512.
- HECKMAN, J. J., LAYNE-FARRAR, A. and TODD, P. E. (1996). Human Capital Pricing Equations with an Application to Estimating the Effect of Schooling Quality on Earnings. *The Review of Economics and Statistics*, **78** (4), 562–610.
- HERING, L. and PONCET, S. (2010). Market Access and Individual Wages: Evidence From China. *The Review of Economics and Statistics*, **92** (1), pp. 145–159.
- HEUERMANN, D. (2011). Human Capital Externalities in Western Germany. *Spatial Economic Analysis*, **6** (2), 139–165.
- HILLER, N. and LERBS, O. (2015). The capitalization of non-market attributes into regional housing rents and wages: evidence on German functional labor market areas. *Review of Regional Research*, **35** (1), 49–72.
- HOLMES, T. J. and LEE, S. (2010). Cities as Six-by-Six-Mile Squares Zipf’s Law? In E. L. Glaeser (ed.), *Agglomeration Economics*, February, The University of Chicago Press, pp. 105 – 131.
- HUNT, J. (2002). The Transition in East Germany: When Is a Ten-Point Fall in the Gender Wage Gap Bad News? *Journal of Labor Economics*, **20** (1), 148–169.
- HYMAN, B. J. (2017). Does Money Matter in the Long Run? Effects of School Spending on Educational Attainment. *American Economic Journal: Economic Policy*, **9** (4), 256–280.
- IBGE (2002). Censo Demografico 2000: Documentacao das Microdados da Amostra.
- (2016). *Arranjos Populacionais e Concentrações Urbanas do Brasil*. Rio de Janeiro: IBGE, 2nd edn.
- IPEA (2012). 2012: Desenvolvimento Inclusivo Sustentável?
- JACKSON, C. K., JOHNSON, R. C. and PERSICO, C. (2016). The Effects of School Spending on Educational and Economic Outcomes: Evidence from School Finance Reforms. *The Quarterly Journal of Economics*, (May), 157–218.
- JEDWAB, R. and VOLLRATH, D. (2016). The Urban Mortality Transition and the Rise of Poor Mega Cities.
- JOHNSON, M. and KEANE, M. P. (2013). A Dynamic Equilibrium Model of the US Wage Structure, 1968-1996. *Journal of Labor Economics*, **31** (1), 1–49.
- KE, S. (2010). Agglomeration, productivity, and spatial spillovers across Chinese cities. *Annals of Regional Science*, **45** (1), 157–179.
- and Yu, Y. (2014). The pathways from industrial agglomeration to TFP growth - the experience of Chinese cities for 2001-2010. *Journal of the Asia Pacific Economy*, **19** (January), 310–332.

- KLINE, P. (2010). Place Based Policies, Heterogeneity, and Agglomeration. *The American Economic Review: Papers and Proceedings*, **100** (2), 383–387.
- and MORETTI, E. (2014a). Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority. *The Quarterly Journal of Economics*, **129** (1), 275–331.
- and — (2014b). People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs. *Annual Review of Economics*, **6** (1), 629–662.
- KNIGHT, J. and GUNATILAKA, R. (2010). Great Expectations? The Subjective Well-being of Rural-Urban Migrants in China. *World Development*, **38** (1), 113–124.
- KOSEC, K. (2014). Relying on the Private Sector: The Income Distribution and Public Investments in the Poor. *Journal of Development Economics*, **107**, 320–342.
- KOVAK, B. K. (2013). Regional effects of trade reform: What is the correct measure of liberalization? *American Economic Review*, **103** (5), 1960–1976.
- KRUEGER, A. B. and LINDAHL, M. (2001). Education for Growth: Why and For Whom? *Journal of Economic Literature*, **XXXIX** (December), 1101–1136.
- LALL, S. V., SHALIZI, Z. and DEICHMANN, U. (2004). Agglomeration economies and productivity in Indian industry. *Journal of Development Economics*, **73** (2), 643–673.
- LAZEAR, E. P. and ROSEN, S. (1990). Male-Female Wage Differentials in Job Ladders. *Journal of Labor Economics*, **8** (1), 106–123.
- LEE, S. and LIN, J. (2017). Natural Amenities, Neighborhood Dynamics, and Persistence in the Spatial Distribution of Income. *Review of Economic Studies*, (13).
- LIN, H. L., LI, H. Y. and YANG, C. H. (2011). Agglomeration and productivity: Firm-level evidence from China’s textile industry. *China Economic Review*, **22** (3), 313–329.
- LIU, Y. and SHEN, J. (2014). Jobs or Amenities? Location Choices of Interprovincial Skilled Migrants in China, 2000–2005. *Population, Space and Place*, **20** (7), 592–605.
- LIU, Z. Q. (2007). The external returns to education: Evidence from Chinese cities. *Journal of Urban Economics*, **61** (3), 542–564.
- LLERAS-MUNEY, A. (2005). The Relationship Between Education and Adult Mortality in the United States. *The Review of Economic Studies*, (72), 189–221.
- LOCHNER, L. and MORETTI, E. (2004). The Effect of Education on Crime: Evidence from Prison Inmates, Arrest, and Self-Reports. *American Economic Review*, **94** (1), 155–189.
- LONG, F., GUO, M. and ZHENG, S. (2009). Estimating the Willingness-to-Pay for Urban Housing in Chinese Cities. *Tsinghua Science & Technology*, **14** (3), 360–366.
- MCKINSEY GLOBAL INSTITUTE (2010). *India’s Urban Awakening: Building Inclusive Cities, Sustaining Economic Growth*. Tech. Rep. April.

- MCLANAHAN, S. and ADAMS, J. (1987). Parenthood and Psychological Well-Being. *Annual review of sociology*, **13** (1987), 237–257.
- MELO, P. C., GRAHAM, D. J. and NOLAND, R. B. (2009). A meta-analysis of estimates of urban agglomeration economies. *Regional Science and Urban Economics*, **39** (3), 332–342.
- MENEZES-FILHO, N. and PAZELLO, E. (2006). Evaluating the Effects of FUNDEF on Wages and Test Scores in Brazil.
- and — (2007). Do teachers' wages matter for proficiency? Evidence from a funding reform in Brazil. *Economics of Education Review*, **26** (6), 660–672.
- MENEZES-FILHO, N. A. (2001). A Evolução da Educação no Brasil e seu Impacto no Mercado de Trabalho.
- MINCER, J. (1974). *Schooling, experience, and earnings*. New York: Columbia University Press.
- MONTIEL OLEA, J. L. and PFLUEGER, C. (2013). A Robust Test for Weak Instruments. *Journal of Business & Economic Statistics*, **31** (3), 358–369.
- MORETTI, E. (2004). Estimating the social return to higher education: evidence from longitudinal and repeated cross-sectional data. *Journal of Econometrics*, **121** (1-2), 175–212.
- (2011). Chapter 14: Local Labor Markets. In O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*, vol. 4B, vol. 4, Elsevier Inc., pp. 1237–1313.
- MORTEN, M. (2013). Temporary Migration and Endogenous Risk Sharing in Village India.
- and OLIVEIRA, J. (2016). The Effects of Roads on Trade and Migration: Evidence from a Planned Capital City.
- MUELLER, V. (2005). *Valuating Climate Amenities in Brazil Using a Hedonic Pricing Framework*. Doctoral dissertation, University of Maryland at College Park.
- MUNSHI, K. and ROSENZWEIG, M. (2009). Why is Mobility in India so Low ? Social Insurance, Inequality, and Growth.
- MUTH, R. (1969). *Cities and Housing*. Chicago and London: University of Chicago Press.
- NEUMARK, D. and SIMPSON, H. (2015). Place-Based Policies. *Handbook of Regional and Urban Economics*, **5**, 1197–1287.
- NOTOWIDIGDO, M. J. (2013). The Incidence of Local Labor Demand Shocks.
- OECD (2011). Brazil: Encouraging Lessons from a Large Federal System. In *Lessons from PISA for the United States*, 8, Paris: OECD Publishing, pp. 177–200.
- (2012). *Redefining "Urban" - A New Way to Measure Metropolitan Areas*. Paris: OECD Publishing.
- (2015). *OECD Urban Policy Reviews: China 2015*. Paris: OECD Publishing.

- OLIVEIRA, J. and PEREDA, P. (2015). The Value of Brazilian Climate: New Evidence from a Spatial Equilibrium Model with Two Sectors.
- OLIVETTI, C. and PETRONGOLO, B. (2008). Unequal pay or unequal employment?: a cross-country analysis of gender gaps. *Journal of Labor Economics*, **26** (4), 621–654.
- and — (2014). Gender gaps across countries and skills: Demand, supply and the industry structure. *Review of Economic Dynamics*, **17** (4), 842–859.
- OREOPOULOS, P. and SALVANES, K. G. (2011). Priceless: The Nonpecuniary Benefits of Schooling. *Journal of Economic Perspectives*, **25** (1), 159–184.
- PAES DE BARROS, R., OLINTO, P., LUNDE, T. and CARVALHO, M. (2011). The impact of access to free childcare on women’s labor market outcomes: evidence from a randomized trial in low-income neighborhoods of Rio de Janeiro. *Paper prepared for the 2011 World Bank Economists’ Forum*.
- PAN, Z. and ZHANG, F. (2002). Urban Productivity in China. *Urban Studies*, **39** (12), 2267–2281.
- PARTRIDGE, M. D., RICKMAN, D. S. and LI, H. (2009). Who wins from local economic development? *Economic Development Quarterly*, **23** (1), 13–27.
- PITTAU, M. G., ZELLI, R. and GELMAN, A. (2010). Economic disparities and life satisfaction in European regions. *Social Indicators Research*, **96** (2), 339–361.
- PONTHIEUX, S. and MEURS, D. (2015). Gender inequality. In *Handbook of Income Distribution*, vol. 2, 1st edn., Elsevier B.V., pp. 981–1146.
- PRITCHETT, L. (2001). Where Has All the Education Gone? *The World Bank Economic Review*, **15** (3), 367–391.
- (2006). Does Learning to Add Add Up? The Returns to Schooling in Aggregate Data. In *Handbook of the Economics of Education*, vol. 1, 11, pp. 635–695.
- RAMOS, R., SURIÑACH, J. and ARTÍS, M. (2010). Human capital spillovers, productivity and regional convergence in Spain. *Papers in Regional Science*, **89** (2), 435–447.
- RAUCH, J. E. (1993). Productivity Gains from Geographic Concentration of Human Capital: Evidence from the Cities. *Journal of Urban Economics*, **34**, 380–400.
- REHDANZ, K. and MADDISON, D. (2009). The amenity value of climate to households in Germany. *Oxford Economic Papers*, **61** (1), 150–167.
- REIS, E. J., PIMENTAL, M. and ALVARENGA, A. I. (2007). Areas Mínimas Comparáveis para os Períodos Intercensitários de 1872 a 2000.
- REIS, M. C. (2006). Os impactos das mudanças na demanda por trabalho qualificado sobre o desemprego por nível de qualificação durante os anos noventa no Brasil. *Revista Brasileira de Economia*, **60** (3), 297–319.
- RESENDE, G. M. (2013). Regional Development Policy in Brazil: A Review of Evaluation Literature. *REDES*, **18** (3), 202–225.

- RESENDE, M. (2004). Gibrat's law and the growth of cities in Brazil: A panel data investigation. *Urban Studies*, **41** (8), 1537–1549.
- RICE, P., VENABLES, A. J. and PATACCHINI, E. (2006). Spatial determinants of productivity: Analysis for the regions of Great Britain. *Regional Science and Urban Economics*, **36**, 727–752.
- ROBACK, J. (1982). Wages, rents, and the quality of life. *The Journal of Political Economy*, **90** (6), 1257–1278.
- RODRÍGUEZ-POSE, A. and TSELIOS, V. (2012). Individual Earnings and Educational Externalities in the European Union. *Regional Studies*, **1** (46), 39–57.
- and VILALTA-BUFÍ, M. (2005). Education, migration, and job satisfaction: The regional returns of human capital in the EU. *Journal of Economic Geography*, **5** (5), 545–566.
- ROSE, A. K. (2006). Cities and Countries. *Journal of Money, Credit and Banking*, **38** (8), 2225–2245.
- ROSEN, K. T. and RESNICK, M. (1980). The size distribution of cities: An examination of the Pareto law and primacy. *Journal of Urban Economics*, **8** (2), 165–186.
- ROSEN, S. (1979). Wage-Based Indexes of Urban Quality of Life. In P. Mieszkowski and M. Stras-Zheim (eds.), *Current Issues in Urban Economics*, Baltimore and London: Johns Hopkins University Press, pp. 74–104.
- SCHAFFAR, A. and DIMOU, M. (2012). Rank-size City Dynamics in China and India, 1981-2004. *Regional Studies*, **46** (6), 707–721.
- SCHWARTZMAN, S. (2005). Education-Oriented Social Programs in Brazil: The Impact of Bolsa Escola. *Instituto de Estudos do Trabalho e Sociedade, Rio de Janeiro*, pp. 1–28.
- SCORZAFAVE, L. G. and MENEZES-FILHO, N. A. (2001). Participação feminina no mercado de trabalho brasileiro: evolução e determinantes. *Pesquisa e Planejamento Econômico*, **31** (3), 441–478.
- and — (2005). Impacto da participação das mulheres na evolução da distribuição de renda do trabalho no Brasil. *Pesquisa e Planejamento Econômico*, **35** (2), 245–266.
- SERRANO, L. (2003). Human Capital Externalities: A Sectoral-Regional Application for Spain. *Fundación Centro de Estudios Andaluces, Documento de Trabajo*, **E2003/06**.
- SHAPIRO, J. M. (2006). Smart Cities: Quality of Life, Productivity and The Growth Effects of Human Capital. *The Review of Economics and Statistics*, **88** (May), 324–335.
- SHEN, J. (2013). Increasing internal migration in China from 1985 to 2005: Institutional versus economic drivers. *Habitat International*, **39**, 1–7.
- SIMON, C. J. and NARDINELLI, C. (2002). Human capital and the rise of American cities, 1900-1990. *Journal of Urban Economics*, **32**, 59–96.
- SOARES, S. and IZAKI, R. S. (2002). A Participação Feminina no Mercado do Trabalho. *Textos para Discussão*, (923).

- Soo, K. T. (2005). Zipf's Law for cities: a cross-country investigation. *Regional Science and Urban Economics*, **35** (3), 239–263.
- (2014). Zipf, Gibrat and geography: Evidence from China, India and Brazil. *Papers in Regional Science*, **93** (1), 159–181.
- SORENSEN, O. and DAHL, M. S. (2016). Geography, Joint Choices and the Reproduction of Gender Inequality. *American Sociological Review*, **81** (5), 900–920.
- SÜDEKUM, J. (2010). Human capital externalities and growth of high- and low-skilled jobs. *Jahrbucher fur Nationalökonomie und Statistik*, **230** (1), 92–114.
- TABUCHI, T. and YOSHIDA, A. (2000). Separating Agglomeration Economies in Consumption and Production. *Journal of Urban Economics*, **48**, 70–84.
- TOPEL, R. H. (1986). Local Labor Markets. *Journal of Political Economy*, **94** (3), 111–143.
- WORLD BANK (2014). *Urban China: Toward Efficient, Inclusive, and Sustainable Urbanization*. Tech. rep.
- WOZNIAK, A. (2010). Are college graduates more responsive to distant labor market opportunities? *Journal of Human Resources*, **45** (4), 944–970.
- XU, Z. (2009). Productivity and agglomeration economies in Chinese cities. *Comparative economic studies*, **51** (3), 284–301.
- YOUNG, A. (2013). Inequality, the Urban-Rural Gap, and Migration. *The Quarterly Journal of Economics*, **128** (4), 1727–1785.

Appendix A

Appendix to Chapter 1

A.1 Supplementary Figures for Chapter 1

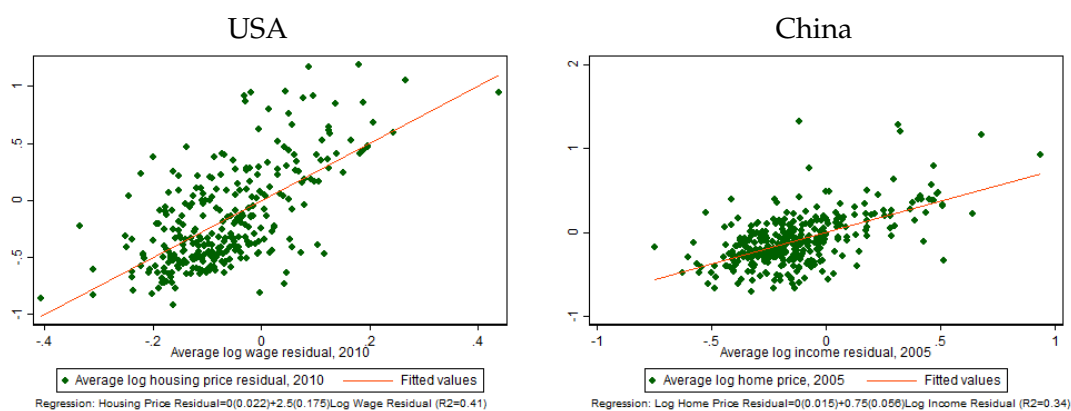


Figure A.1: *Income and housing prices, 2010*

Note: Samples restricted to areas with urban population of 100,000 or more. **Sources:** See data appendix.

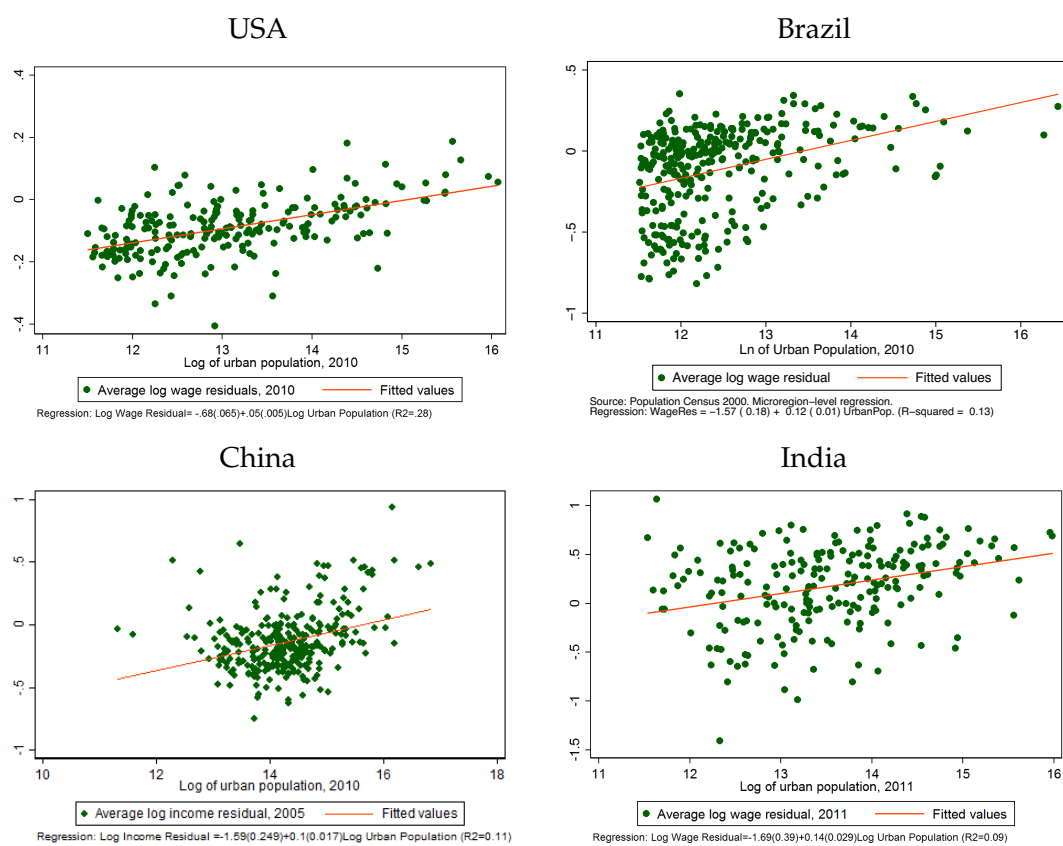


Figure A.2: *Urban population and income residuals, 2010*

Note: Samples restricted to areas with urban population of 100,000 or more. **Sources:** See data appendix.

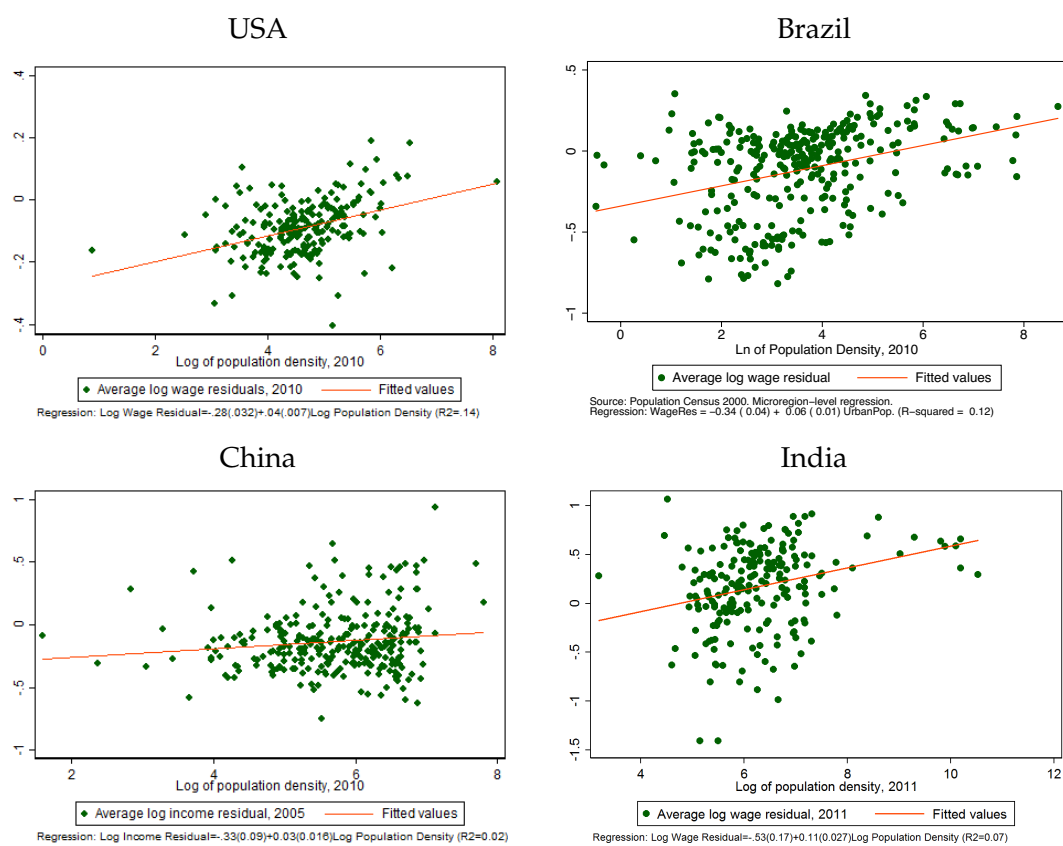


Figure A.3: *Population density and income residuals, 2010*

Note: Samples restricted to areas with urban population of 100,000 or more. **Sources:** See data appendix.

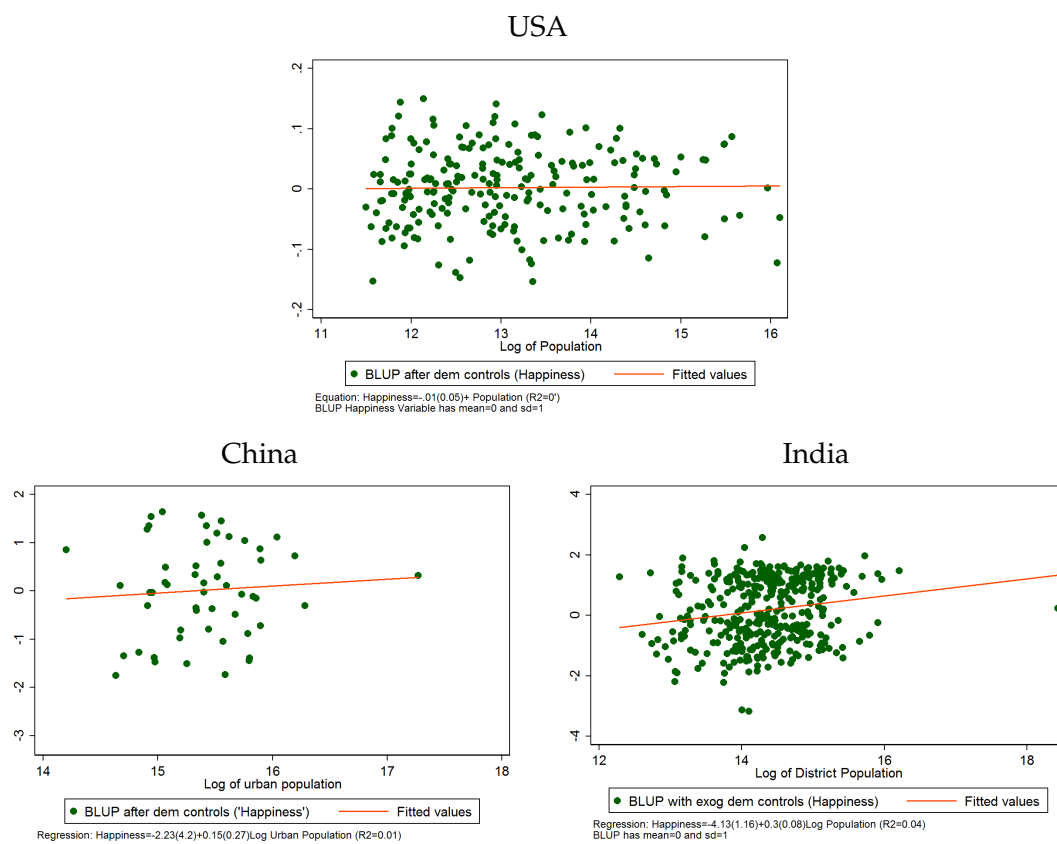


Figure A.4: *Happiness and population size*

Note: Samples restricted to areas with urban population of 100,000 or more. **Sources:** See data appendix.

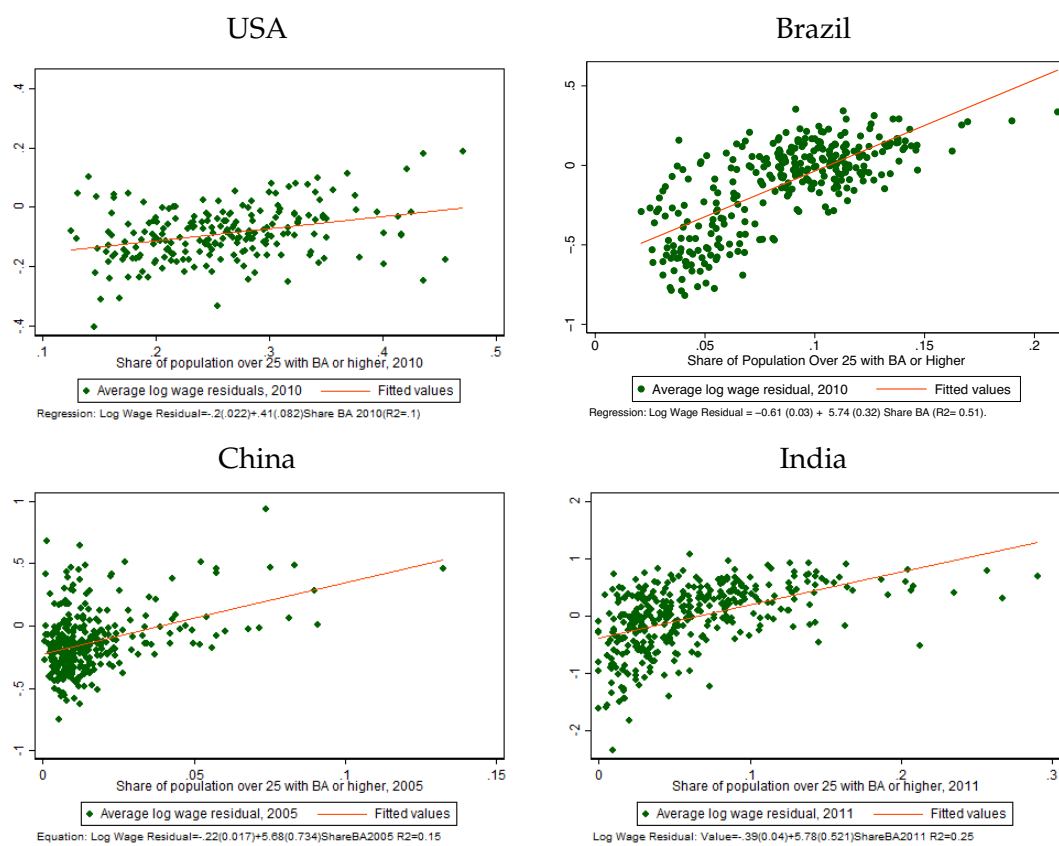


Figure A.5: *University graduates share and wage residuals, 2010*

Note: Samples restricted to areas with urban population of 100,000 or more. **Sources:** See data appendix.

A.2 Supplementary Tables for Chapter 1

Table A.1: *Happiness Regressions*

	USA (2010 MSAs)	China (2002 Cities)	India (2011 Districts)
	"Happiness"	"Happiness"	"Happiness"
Log Median Income (Disposable Income for China)	0.07*** (0.027)	0.54 (0.373)	0.24 (0.176)
Constant	-0.76*** (0.292)	-4.3 (2.972)	-2.44 (1.73)
R-Squared	0.03	0.04	0.01
Log of Population	0.00 (0.004)	0.15 (0.27)	0.3*** (0.08)
Constant	-0.01 (0.0522)	-2.23 (4.2)	-4.13*** (1.16)
R-Squared	0.00	0.01	0.04
Observations	267	54	369

Note: Regressions at the area level. Area "happiness" is measured as the best linear unbiased predictions (BLUPs) of the random effects after controlling for the exogenous demographic variables of age and race as in Glaeser *et al.* (2014a). Robust standard errors in parentheses.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

Table A.2: *Local prices and agglomeration, 2010*

	USA (MSAs)	Brazil (Microregions)	China (Cities)	India (Districts)	USA (MSAs)	China (Cities)
	Log rent				Log price	
OLS regressions						
Log of urban population	0.147*** (0.0136) R2= 0.432	0.128*** (0.021) R2=0.438	0.229*** (0.0762) R2=0.155	0.003 (0.005) R2=0.745	0.198*** (0.0427) R2=0.369	0.0982 (0.122) R2=0.409
Log of density	0.176*** (0.0141) R2=0.453	0.072*** (0.018) R2=0.420	0.373*** (0.119) R2=0.11	0.002 (0.004) R2=0.761	0.255*** (0.0526) R2=0.371	0.232*** (0.0535) R2=0.452
Observations	24.4M	818 K	6.7K	3,281	44M	25K
IV1 regressions						
Log of urban population	0.152*** (0.0131) R2=0.432	0.125*** (0.023) R2 = 0.438	0.365*** (0.130) R2=0.162	-0.004 (0.009) R2=0.760	0.197*** (0.0446) R=0.372	0.0599 (0.131) R2=0.405
Log of density	0.168*** (0.0156) R2=0.453	0.073*** (0.019) R2 = 0.420	0.588*** (0.221) R2=0.092	0.002 (0.004) R2=0.767	0.224*** (0.0574) R2=0.370	0.214*** (0.0651) R2=0.449
Observations	24.4M	818 K	6.4K	2,595	44M	24K
IV2 regressions						
Log of urban population	0.143*** (0.0229) R2=0.433	0.078** (0.033) R2 = 0.423	-0.007 (0.118) R2=0.087	-0.018* (0.010) R2=0.730	0.0920 (0.0623) R2=0.353	0.387 (0.227) R2=0.331
Log of density	0.141*** (0.0267) R2=0.453	0.057*** (0.021) R2 = 0.413	0.221** (0.105) R2=0.161	0.002 (0.004) R2=0.755	0.0376 (0.101) R2=0.326	0.451*** (0.109) R2=0.473
Observations	24.4M	744 K	4.4K	1,792	44M	19K
Dwelling characteristics	Yes	Yes	Yes	Yes	Yes	Yes

Note: Regressions at the urban household level, restricted to areas with urban population of 100,000 or more. All regressions include a constant. Standard errors clustered at the area level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

Appendix B

Appendix to Chapter 2

B.1 Supplementary Figures for Chapter 2

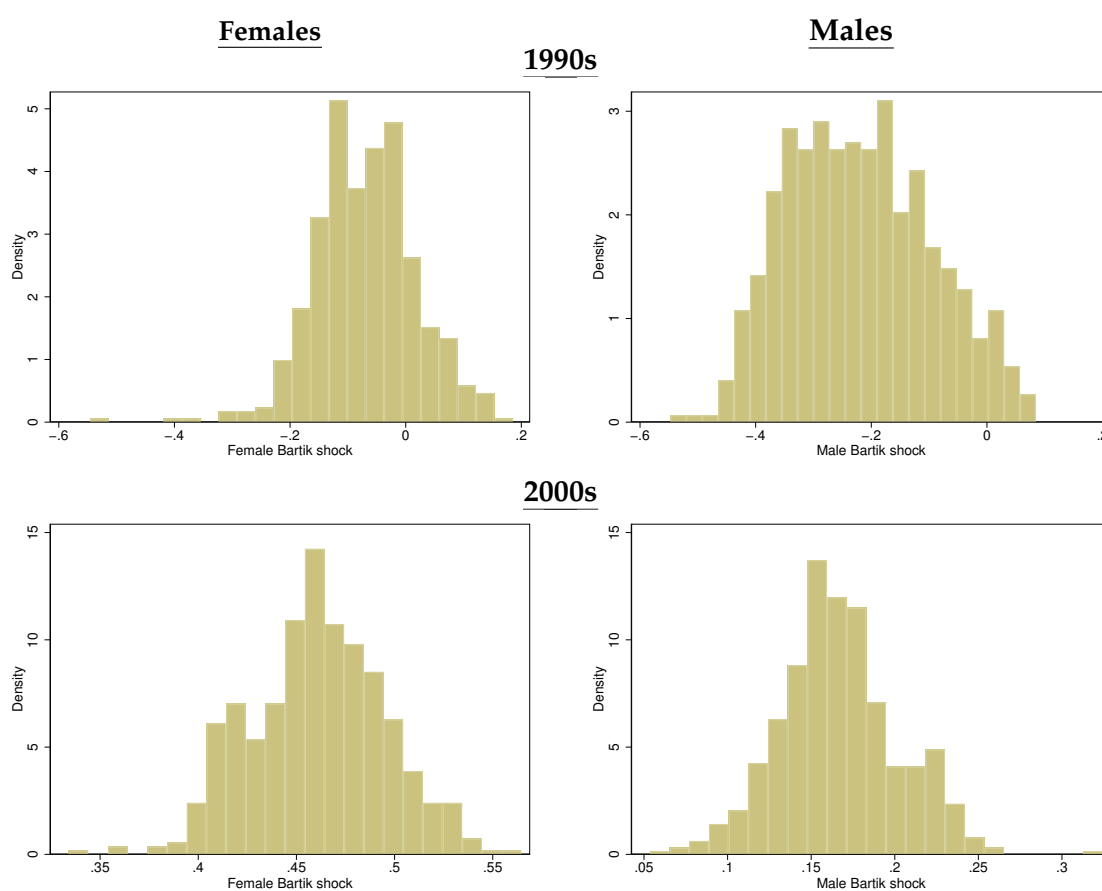


Figure B.1: *Distributions of Gender-specific Bartik shocks*

Note: Own calculations using census data. **Sources:** See data appendix.

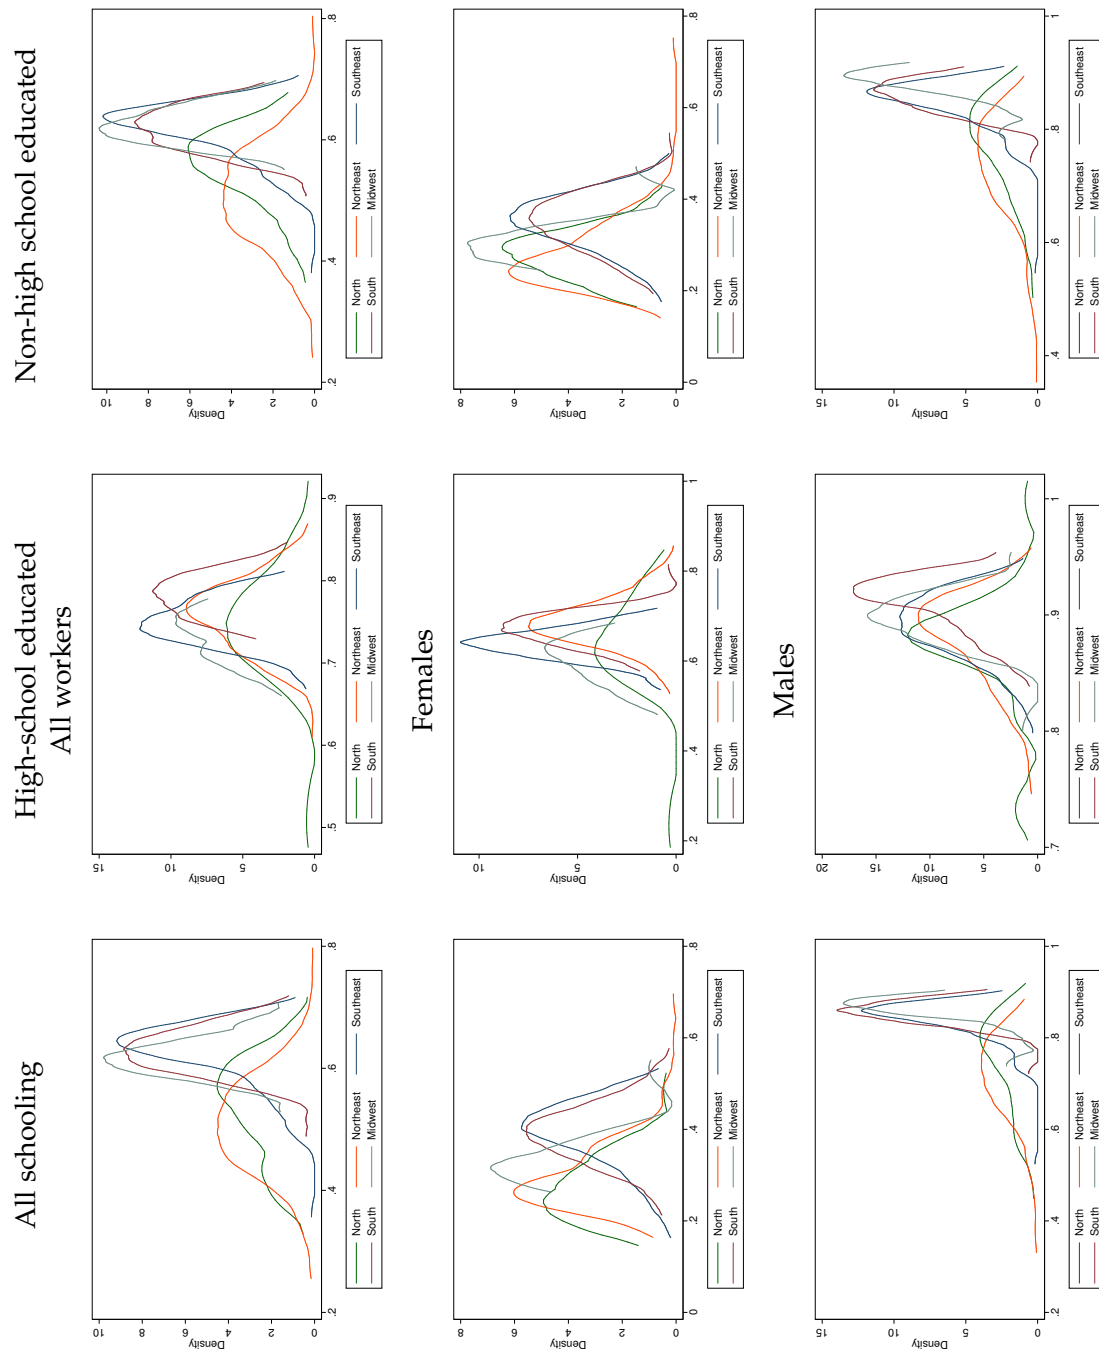
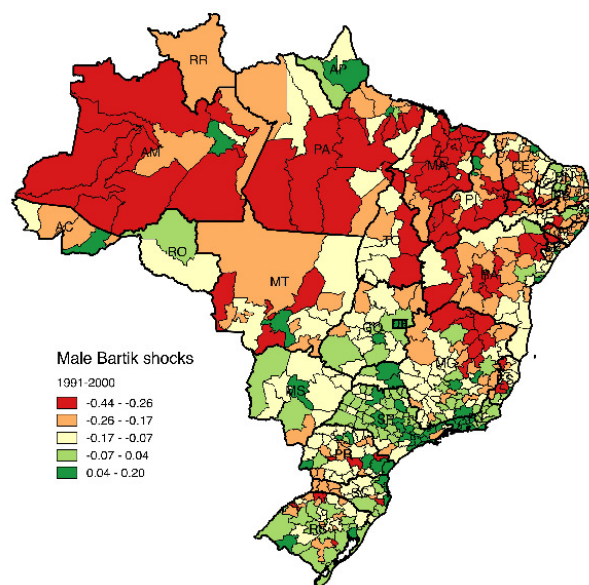


Figure B.2: *Labor force participation across microregions in Brazilian macro-regions, 2000*
Note: Own calculations using census data. **Sources:** See data appendix.

Males



Females

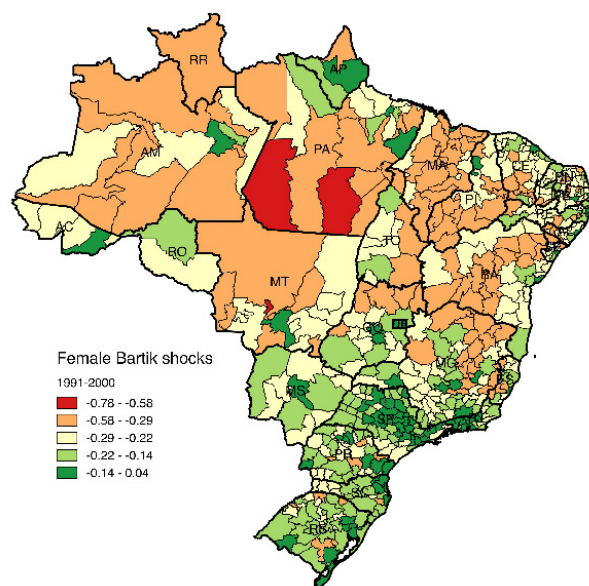


Figure B.3: *Geographic distribution of gender-specific Bartik shocks, Brazil 1991-2000*

Note: Own calculations using census data. **Sources:** See data appendix.

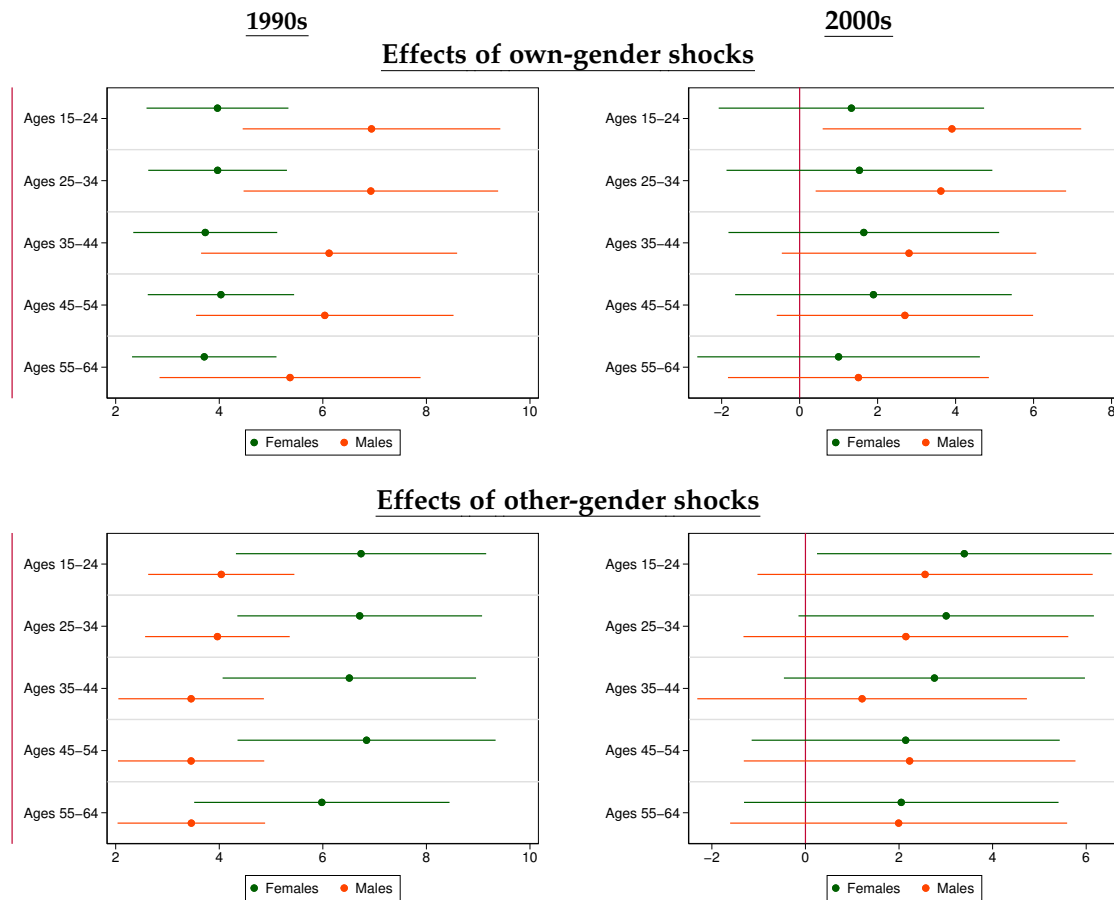


Figure B.4: *Effects of gender-specific shocks on migrant population by age*

Note: Own calculations using census data. **Sources:** See data appendix.

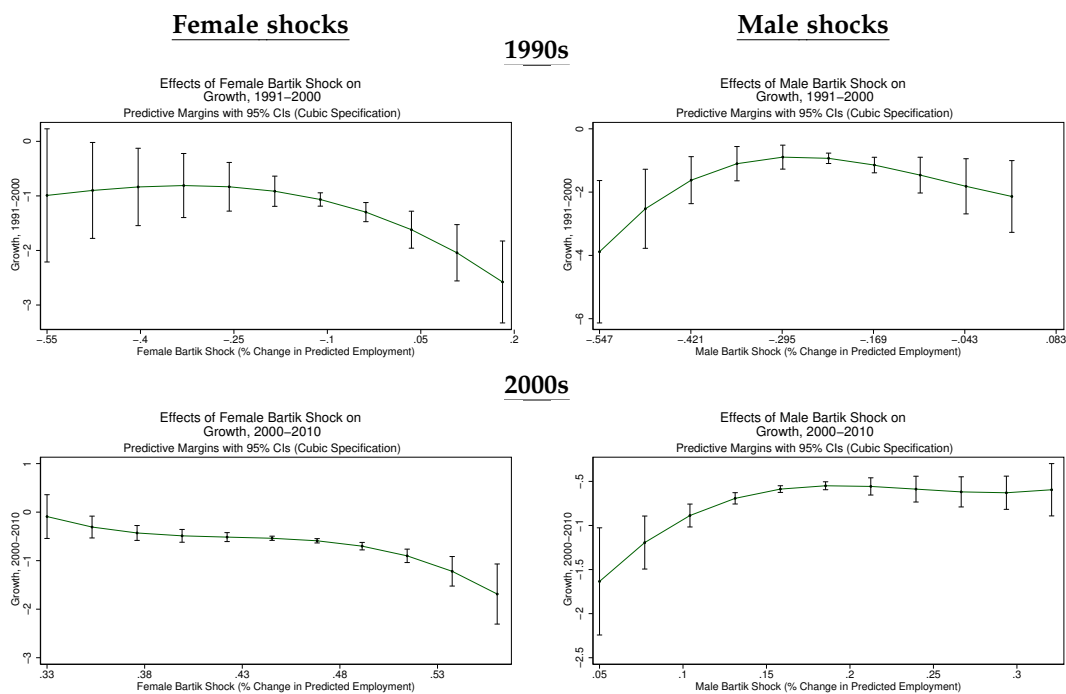


Figure B.5: *Effects of gender-specific shocks on the employment gap, predictive margins*

Note: Own calculations using census data. **Sources:** See data appendix.

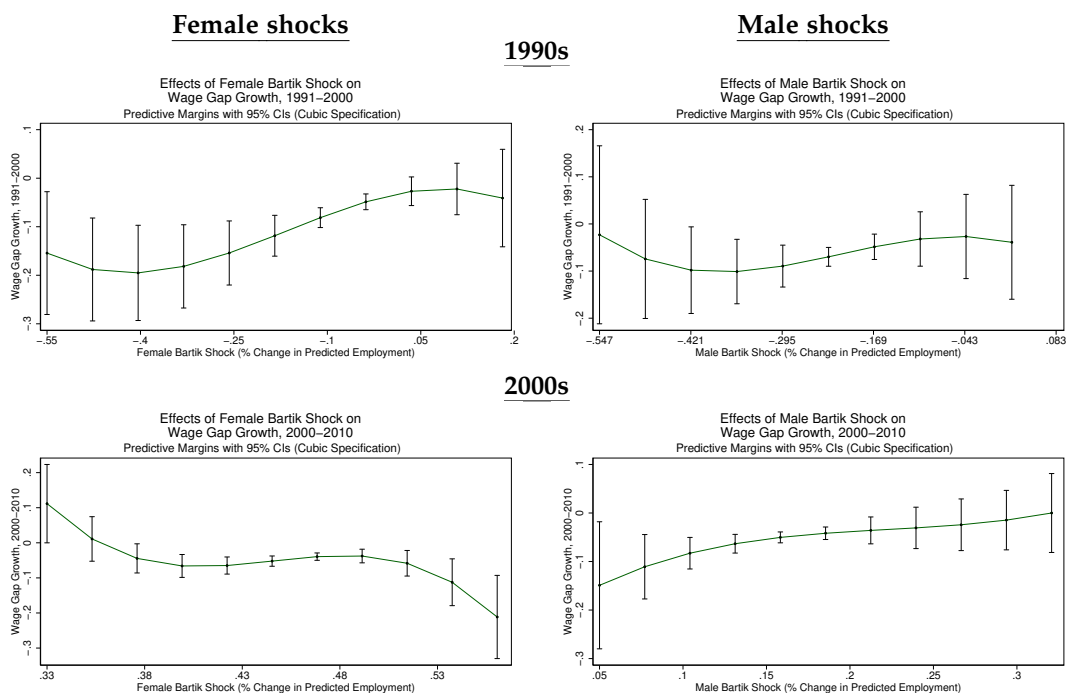


Figure B.6: *Effects of gender-specific shocks on the wage gap, predictive margins*

Note: Own calculations using census data. **Sources:** See data appendix.

B.2 Supplementary Tables for Chapter 2

Table B.1: *Summary statistics, 1990s*

	Mean	Std. Dev.	Min	Max
Shocks				
Bartik shocks, males	-0.07	0.09	-0.55	0.19
Bartik shocks, females	-0.22	0.12	-0.55	0.08
Main outcomes				
Δ_{91-00} Population	0.11	0.13	-0.46	0.93
Δ_{91-00} Female employment	0.31	0.21	-0.49	1.28
Δ_{91-00} Male employment	-0.06	0.2	-1.01	0.82
Δ_{91-00} Female average wage residual	0.01	0.17	-0.65	0.74
Δ_{91-00} Male average wage residual	-0.03	0.15	-0.66	0.37
Base year (1991) controls				
Log of population density	3.17	1.46	-1.65	8.51
Average log wage residual	-0.23	0.31	-1.1	0.74
Average temperature in the winter (C°)	20.86	4.16	11.83	27.25
Share of High-school educated	0.09	0.05	0	0.3
Formally-employed share in adult population	0.19	0.12	0.01	0.5
Informally-employed share in adult population	0.36	0.08	0.16	0.58
Unemployment rate	0.04	0.02	0	0.16
Share of population aged 0-14	0.37	0.06	0.27	0.53
Share of population aged 15-24	0.19	0.01	0.16	0.23
Share of population aged 25-34	0.15	0.02	0.1	0.2
Share of population aged 35-44	0.11	0.02	0.07	0.14
Share of population aged 45-44	0.07	0.01	0.05	0.11
Share of population aged 55-64	0.05	0.01	0.01	0.08
Urbanization rate	0.6	0.2	0.14	1
Share of employment in agriculture	0.45	0.21	0.01	0.92
Share of employment in manufacturing	0.1	0.08	0.01	0.52
Share of employment in government	0.03	0.01	0	0.15
Lagged changes controls				
Δ_{80-91} Population	0.23	0.22	-0.17	2.99
Δ_{80-91} Wage residual	-0.03	0.14	-0.56	0.46
Δ_{80-91} Formal employment	0.03	0.04	-0.1	0.21
Δ_{80-91} Informal employment	0	0.05	-0.21	0.17
Δ_{80-91} Unemployment rate	0.02	0.02	-0.19	0.14
Δ_{80-91} Urbanization rate	0.11	0.06	-0.2	0.48

Note: Own calculations with population censuses of 1980, 1991 and 2000. Outcomes calculated for individuals aged 15-64. N=539. **Sources:** See data appendix.

Table B.2: *Summary statistics, 2000s*

	Mean	Std. Dev.	Min	Max
Shocks				
Bartik shocks, males	0.46	0.03	0.33	0.56
Bartik shocks, females	0.17	0.04	0.05	0.32
Main outcomes				
Δ_{00-10} Population	0.18	0.11	-0.18	0.77
Δ_{00-10} Female employment	0.48	0.17	-0.07	1.35
Δ_{00-10} Male employment	0.2	0.13	-0.33	0.82
Δ_{00-10} Female average wage residual	0.03	0.11	-0.36	0.3
Δ_{00-10} Male average wage residual	0.03	0.12	-0.46	0.31
Base year (2000) controls				
Log of population density	3.29	1.46	-1.5	8.6
Average log wage residual	-0.26	0.26	-1.05	0.38
Average temperature in the winter (C°)	20.86	4.16	11.83	27.25
Share of High-school educated	0.15	0.07	0.02	0.38
Formally-employed share in adult population	0.17	0.1	0.01	0.48
Informally-employed share in adult population	0.34	0.06	0.15	0.51
Unemployment rate	0.13	0.04	0.03	0.26
Share of population aged 0-14	0.32	0.05	0.21	0.49
Share of population aged 15-24	0.2	0.02	0.16	0.24
Share of population aged 25-34	0.15	0.02	0.11	0.19
Share of population aged 35-44	0.12	0.02	0.08	0.17
Share of population aged 45-44	0.09	0.02	0.04	0.13
Share of population aged 55-64	0.06	0.01	0.03	0.1
Urbanization rate	0.67	0.18	0.19	1
Share of employment in agriculture	0.37	0.19	0	0.84
Share of employment in manufacturing	0.11	0.07	0.01	0.49
Share of employment in government	0.03	0.02	0.01	0.12
Lagged changes controls				
Δ_{91-00} Population	0.11	0.13	-0.46	0.93
Δ_{91-00} Wage residual	-0.03	0.14	-0.61	0.34
Δ_{91-00} Formal employment	-0.01	0.04	-0.13	0.1
Δ_{91-00} Informal employment	-0.02	0.06	-0.24	0.1
Δ_{91-00} Unemployment rate	0.08	0.03	-0.02	0.21
Δ_{91-00} Urbanization rate	0.07	0.05	-0.05	0.49

Note: Own calculations with population censuses of 1980, 1991 and 2000. Outcomes calculated for individuals aged 15-64. N=539. **Sources:** See data appendix.

Table B.3: Correlations, 1990s

Variables	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
1 Bartik shocks, males	1.00																							
2 Bartik shocks, females	0.87	1.00																						
3 Δy_{1-00} Population	0.26	0.38	1.00																					
4 Δy_{1-00} Female employment	0.07	0.03	0.56	1.00																				
5 Δy_{1-00} Male employment	0.49	0.59	0.73	0.55	1.00																			
6 Δy_{1-00} Female average wage residual	-0.11	-0.13	-0.25	-0.33	-0.35	1.00																		
7 Δy_{1-00} Male average wage residual	0.20	0.16	-0.16	-0.05	0.01	0.52	1.00																	
8 Log of population density	0.50	0.58	0.04	-0.28	0.15	0.23	0.15	1.00																
9 Average log wage residual	0.33	0.47	0.48	0.36	0.57	-0.74	-0.40	-0.02	1.00															
10 Average temperature in the winter (C°)	-0.42	-0.45	0.16	0.07	-0.25	0.02	-0.42	-0.34	-0.18	1.00														
11 Share of High-school educated	0.70	0.86	0.28	-0.02	0.49	-0.18	0.14	0.56	0.45	-0.39	1.00													
12 Formally-employed share in adult population	0.64	0.84	0.27	-0.05	0.50	-0.20	0.11	0.54	0.55	-0.58	0.83	1.00												
13 Informally-employed share in adult population	-0.55	-0.75	-0.21	0.01	-0.38	0.16	-0.10	-0.53	-0.46	0.53	-0.74	-0.91	1.00											
14 Unemployment rate	0.02	0.16	0.23	0.03	0.02	-0.14	-0.28	0.08	0.14	0.28	0.14	0.04	-0.25	1.00										
15 Share of population aged 0-14	-0.59	-0.67	0.06	0.11	-0.34	-0.05	-0.48	-0.49	-0.18	0.73	-0.68	-0.72	0.59	0.31	1.00									
16 Share of population aged 15-24	-0.05	0.03	0.12	0.11	0.13	-0.17	-0.09	-0.01	0.15	0.24	0.03	-0.10	0.11	0.18	0.16	1.00								
17 Share of population aged 25-34	0.59	0.74	0.20	0.13	0.56	-0.28	0.29	0.31	0.56	-0.63	0.75	0.80	-0.66	-0.15	-0.82	0.02	1.00							
18 Share of population aged 35-44	0.65	0.72	0.06	0.01	0.45	-0.13	0.38	0.40	0.39	-0.73	0.72	0.78	-0.64	-0.28	-0.92	-0.25	0.90	1.00						
19 Share of population aged 45-44	0.46	0.44	-0.23	-0.19	0.11	0.16	0.47	0.37	-0.03	-0.66	0.44	0.50	-0.41	-0.41	-0.89	-0.41	0.55	0.78	1.00					
20 Share of population aged 55-64	0.36	0.33	-0.27	-0.25	0.04	0.32	0.47	0.41	-0.21	-0.61	0.32	-0.10	-0.32	-0.40	-0.80	-0.49	0.37	0.62	0.89	1.00				
21 Urbanization rate	0.69	0.88	0.34	-0.01	0.52	-0.21	0.10	0.45	0.53	-0.35	0.82	0.84	-0.70	0.14	-0.65	0.02	0.75	0.71	0.43	0.31	1.00			
22 Share of employment in agriculture	-0.75	-0.96	-0.40	0.00	-0.56	0.13	-0.10	-0.58	-0.50	0.39	-0.87	-0.88	0.78	-0.21	0.66	-0.05	-0.74	-0.69	-0.40	-0.31	-0.92	1.00		
23 Share of employment in manufacturing	0.45	0.64	0.26	-0.01	0.41	-0.16	-0.01	0.47	-0.47	-0.42	0.52	0.78	-0.71	0.03	-0.49	-0.11	0.59	0.56	0.30	0.24	0.59	-0.70	1.00	
24 Share of employment in government	0.06	0.35	0.25	-0.03	0.29	-0.04	0.05	0.12	0.19	-0.07	0.47	0.39	-0.40	0.27	-0.24	0.15	0.29	0.19	0.10	0.07	0.42	-0.45	0.08	1.00

Note: Own-calculations based on population censuses of 1991 and 2000 Sources: See data appendix.

Table B.4: Correlations, 2000s

Variables	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
1 Bartik shocks, males	1.00																							
2 Bartik shocks, females	-0.74	1.00																						
3 Δ_{00-10} Population	-0.17	0.18	1.00																					
4 Δ_{00-10} Female employment	0.20	-0.22	0.56	1.00																				
5 Δ_{00-10} Male employment	-0.39	0.38	0.76	0.53	1.00																			
6 Δ_{00-10} Female average wage residual	-0.06	0.08	-0.08	-0.14	-0.03	1.00																		
7 Δ_{00-10} Male average wage residual	-0.07	0.13	-0.18	-0.22	-0.09	0.79	1.00																	
8 Log of population density	-0.52	0.56	-0.04	-0.36	0.14	0.10	0.05	1.00																
9 Average log wage residual	-0.33	0.37	0.14	0.24	0.27	-0.07	0.02	0.11	1.00															
10 Average temperature in the winter (C°)	0.14	-0.15	0.46	0.17	0.18	-0.26	-0.36	-0.33	-0.40	1.00														
11 Share of High-school educated	-0.61	0.71	-0.02	-0.21	0.20	0.23	0.32	0.54	0.62	-0.45	1.00													
12 Formally-employed share in adult population	-0.57	0.56	-0.08	-0.21	0.14	0.26	0.37	0.50	0.70	-0.65	0.86	1.00												
13 Informally-employed share in adult population	0.54	-0.44	0.01	0.07	-0.33	0.14	0.04	-0.37	-0.23	0.21	-0.40	-0.46	1.00											
14 Unemployment rate	-0.38	0.43	0.16	-0.01	0.36	-0.29	0.20	0.17	0.09	0.21	0.19	0.02	-0.57	1.00										
15 Share of population aged 0-14	0.41	-0.48	0.37	0.38	0.07	-0.44	-0.50	-0.49	-0.37	0.67	-0.71	-0.75	0.18	0.17	1.00									
16 Share of population aged 15-24	0.09	-0.08	0.43	0.18	0.20	-0.39	-0.39	-0.14	-0.17	0.69	-0.32	-0.46	0.06	0.39	0.67	1.00								
17 Share of population aged 25-34	-0.47	0.48	0.05	-0.01	0.25	0.26	0.30	0.35	0.74	-0.43	0.73	0.77	-0.17	0.02	-0.66	-0.27	1.00							
18 Share of population aged 35-44	-0.38	0.46	-0.27	-0.20	0.00	0.44	0.51	0.39	0.59	-0.73	0.76	0.84	-0.19	-0.18	-0.92	-0.70	0.77	1.00						
19 Share of population aged 45-44	-0.29	0.37	-0.44	-0.35	-0.15	0.46	0.51	0.42	0.33	-0.75	0.63	0.70	-0.15	-0.26	-0.94	-0.81	0.51	0.90	1.00					
20 Share of population aged 55-64	-0.08	0.13	-0.58	-0.44	-0.32	0.39	0.42	0.26	-0.15	-0.52	0.24	0.28	-0.03	-0.31	-0.75	-0.76	0.08	0.55	0.80	1.00				
21 Urbanization rate	-0.66	0.73	0.07	-0.23	0.26	0.26	0.36	0.44	0.63	-0.33	0.86	0.82	-0.36	0.25	-0.64	-0.25	0.73	0.70	0.54	0.17	1.00			
22 Share of employment in agriculture	0.79	-0.80	-0.15	0.19	-0.32	-0.18	-0.23	-0.56	-0.63	0.36	-0.86	-0.83	0.41	-0.30	0.61	0.20	-0.75	-0.67	-0.50	-0.10	-0.90	1.00		
23 Share of employment in manufacturing	-0.53	0.29	0.09	-0.01	0.17	0.14	0.12	0.38	0.52	-0.40	0.50	0.70	-0.31	-0.09	-0.42	-0.28	0.53	0.52	0.40	0.08	0.52	-0.62	1.00	
24 Share of employment in government	-0.35	0.22	0.11	-0.13	0.15	-0.11	-0.05	0.02	-0.08	0.21	0.08	-0.01	-0.27	0.37	0.07	0.14	-0.02	-0.14	-0.16	-0.07	0.08	-0.14	-0.18	1.00

Note: Own-calculations based on population censuses of 2000 and 2010 Sources: See data appendix.

Table B.5: Population and work in Brazil between 1991 and 2000

	All education levels			Less than high-school			High-school or higher		
	All	Male	Female	All	Male	Female	All	Male	Female
Panel A: Levels in 1991									
Wages (in 2010 Reais)									
Working-age population	88,770,975	43,466,944	45,304,030	64,779,259	32,309,153	32,470,104	13,433,827	6,234,115	7,199,712
Non-participation rate	39.40%	16.39%	61.48%	40.11%	12.73%	67.36%	19.89%	6.64%	31.37%
Employment rate	57.54%	80.08%	35.92%	56.81%	83.51%	30.24%	76.83%	90.29%	65.17%
Labor force	53,794,919	36,341,883	17,453,036	38,793,230	28,196,008	10,597,221	10,761,475	5,820,313	4,941,162
Formality rate	51.60%	49.67%	55.62%	43.91%	43.73%	44.37%	76.76%	74.64%	79.25%
Informality rate	43.36%	46.11%	37.63%	50.95%	51.96%	48.27%	19.15%	22.08%	15.71%
Unemployment rate	5.04%	4.22%	6.75%	5.14%	4.31%	7.36%	4.09%	3.29%	5.04%
Panel B: Changes 1991-2000									
Working-age population	21.04%	21.07%	21.01%	-10.56%	-9.21%	-11.92%	51.56%	48.69%	53.98%
Non-participation rate	-3.41%	5.21%	-11.67%	-4.24%	4.20%	-12.13%	-1.53%	2.14%	-5.14%
Employment rate	-3.80%	-11.62%	3.70%	-2.19%	-9.85%	4.91%	-4.38%	-6.37%	-2.15%
Labor force	26.52%	14.64%	47.48%	-3.71%	-14.15%	19.69%	53.46%	46.38%	61.20%
Formality rate	-10.01%	-7.92%	-14.26%	-7.85%	-5.87%	-11.76%	-16.30%	-13.58%	-19.41%
Informality rate	-0.98%	-0.54%	-0.05%	-1.83%	-1.14%	-2.37%	9.13%	8.86%	9.87%
Unemployment rate	10.99%	8.46%	14.31%	9.69%	7.01%	14.13%	7.17%	4.71%	9.53%

Note: Own calculations from data of the 1991 and 2000 population censuses. For working age population and labor force the changes are measures as log-differences. For all other variables the changes are simple differences of the rates across census years. **Sources:** See data appendix.

Table B.6: Geographic Mobility working-age population in Brazil 1991-2000

	All education levels			Less than high-school			High-school or higher		
	All	Male	Female	All	Male	Female	All	Male	Female
Panel A: Totals									
Population aged 15-65	109,561,279	53,664,398	55,896,881	58,288,370	29,467,303	28,821,067	22,497,118	10,144,649	12,352,469
Married share	56.96%	56.93%	56.99%	68.56%	67.49%	69.65%	59.11%	63.27%	55.69%
Singles share	43.04%	43.07%	43.01%	31.44%	32.51%	30.35%	40.89%	36.73%	44.31%
Migrant population aged 15-65	24,501,659	11,970,129	12,531,530	13,739,906	6,887,839	6,852,067	4,901,727	2,265,163	2,636,564
Share of migrants in total population	22.36%	22.31%	22.42%	23.57%	23.37%	23.77%	21.79%	22.33%	21.34%
Married share in migrant population	61.64%	61.32%	61.96%	71.22%	69.52%	72.92%	65.26%	68.56%	62.43%
Singles share in migrant population	38.36%	38.68%	38.04%	28.78%	30.48%	27.08%	34.74%	31.44%	37.57%
Panel B: By age groups									
Population aged 15-24	34,089,000	17,075,229	17,013,771	12,437,473	6,527,441	5,910,032	4,718,610	1,989,689	2,728,921
Migrant share	23.43%	21.57%	25.31%	29.06%	25.74%	32.73%	21.37%	20.05%	22.33%
Population aged 25-34	26,857,658	13,171,168	13,686,490	15,719,162	8,102,705	7,616,458	7,108,166	3,103,490	4,004,675
Migrant share	28.20%	27.88%	28.50%	29.62%	28.99%	30.29%	26.11%	25.98%	26.20%
Population aged 35-49	31,501,553	15,269,366	16,232,187	19,407,577	9,648,094	9,759,483	7,971,970	3,701,254	4,270,716
Migrant share	20.60%	21.84%	19.44%	20.59%	21.54%	19.66%	20.45%	22.36%	18.80%
Population aged 50-64	17,113,068	8,148,636	8,964,433	10,724,157	5,189,064	5,535,094	2,698,372	1,350,215	1,348,157
Migrant share	14.32%	15.71%	13.04%	13.74%	15.04%	12.51%	15.08%	17.20%	12.97%
Panel B: By Marital status									
Married population	62,405,772	30,552,713	31,853,059	39,960,041	19,887,395	20,072,647	13,297,914	6,418,211	6,879,704
Migrant share	24.20%	24.02%	24.37%	24.49%	24.08%	24.89%	24.06%	24.20%	23.92%
Single population	47,155,507	23,111,685	24,043,822	18,328,329	9,579,908	8,748,420	9,199,203	3,726,438	5,472,765
Migrant share	19.93%	20.03%	19.83%	21.58%	21.92%	21.21%	18.51%	19.11%	18.10%

Note: Own calculations from data of the 2000 population census. A person is considered a migrant if it they moved to their current municipality of residence over the previous 10 years. Age groups, marital status and schooling attainment correspond to the year 2000 (this information is not available for the pre-migration period). **Sources:** See data appendix.

Table B.7: Migrant population and work in Brazil, 2000

	All education levels			Less than high-school			High-school or higher		
	All	Male	Female	All	Male	Female	All	Male	Female
Panel A: All individuals									
Working-age population	109,561,280	53,664,399	55,896,880	58,288,370	29,467,304	28,821,068	22,497,118	10,144,648	12,352,469
Non-participation rate	35.99%	21.60%	49.80%	35.87%	16.94%	55.23%	18.36%	8.77%	26.23%
Employment rate	53.75%	68.45%	39.63%	54.62%	73.66%	35.15%	72.45%	83.93%	63.02%
Labor force	70,128,808	42,071,072	28,057,735	37,379,901	24,476,897	12,903,005	18,366,590	9,254,582	9,112,007
Formality rate	41.59%	41.74%	41.36%	36.05%	37.86%	32.61%	60.46%	61.06%	59.84%
Informality rate	42.38%	45.57%	37.58%	49.12%	50.82%	45.90%	28.28%	30.94%	25.59%
Unemployment rate	16.03%	12.68%	21.06%	14.83%	11.32%	21.49%	11.26%	8.00%	14.57%
Panel B: Migrants									
Working-age population	24,501,659	11,970,129	12,531,531	13,739,907	6,887,839	6,852,067	4,901,727	2,265,162	2,636,564
Non-participation rate	32.75%	16.68%	48.10%	32.63%	12.75%	52.61%	19.28%	7.85%	29.11%
Employment rate	56.46%	73.61%	40.09%	56.90%	77.68%	36.00%	71.62%	85.89%	59.35%
Labor force	16,477,140	9,973,510	6,503,631	9,256,945	6,009,896	3,247,048	3,956,450	2,087,323	1,869,126
Formality rate	41.87%	44.57%	37.73%	37.41%	41.13%	30.51%	58.86%	62.00%	55.36%
Informality rate	42.09%	43.78%	39.51%	47.05%	47.90%	45.46%	29.87%	31.21%	28.37%
Unemployment rate	16.04%	11.66%	22.76%	15.55%	10.97%	24.03%	11.27%	6.79%	16.28%

Note: Own calculations from data of the 2000 population census. A person is considered a migrant if it they moved to their current municipality of residence over the previous 10 years. **Sources:** See data appendix.

Table B.8: *Gender-specific Bartik shocks and start-year characteristics*

	1991-2000 shocks		2000-2010 shocks	
	Females (1)	Males (2)	Females (3)	Males (4)
Log population density	0.01*** (0.00)	0.01*** (0.00)	-0.00*** (0.00)	0.01*** (0.00)
Log wage residuals	0.01 (0.02)	0.00 (0.02)	0.01 (0.01)	0.02* (0.01)
Average winter temperature	-0.00* (0.00)	-0.00** (0.00)	-0.00** (0.00)	0.00* (0.00)
Share of high-school educated	0.59*** (0.12)	0.60*** (0.11)	-0.01 (0.04)	0.14*** (0.04)
Formality rate	-0.18 (0.13)	0.08 (0.12)	0.00 (0.04)	-0.16*** (0.04)
Informality rate	0.00 (0.13)	0.01 (0.11)	0.18*** (0.03)	-0.07* (0.03)
Unemployment rate	0.01 (0.23)	0.43** (0.19)	-0.00 (0.05)	0.12** (0.05)
Population share aged 0-14	1.36*** (0.48)	0.02 (0.37)	0.71*** (0.21)	-0.50*** (0.18)
Population share aged 15-24	1.44** (0.60)	0.44 (0.46)	0.93*** (0.23)	-0.31 (0.20)
Population share aged 25-34	-0.40 (0.57)	-0.29 (0.45)	0.11 (0.24)	-0.83*** (0.24)
Population share aged 35-44	3.88*** (1.01)	1.53** (0.62)	1.04*** (0.32)	0.30 (0.32)
Population share aged 45-54	1.59 (1.06)	-0.78 (0.73)	0.69* (0.37)	-0.74** (0.37)
Population share aged 55-65	1.97* (1.07)	-0.09 (0.85)	1.52*** (0.44)	-0.88* (0.45)
Urbanization rate	0.17*** (0.05)	0.27*** (0.04)	-0.09*** (0.02)	0.10*** (0.01)
Constant	-1.54*** (0.48)	-0.61 (0.38)	-0.20 (0.19)	0.50*** (0.18)
Observations	539	539	539	539
R-squared	0.61	0.86	0.63	0.72

Note: Robust standard errors clustered at the mesoregion level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

Table B.9: *Pre-trends tests*

	1991-2000 shocks		2000-2010 shocks	
	Females (1)	Males (2)	Females (3)	Males (4)
<u>Panel A: Employment growth residuals</u>				
Residuals of 1980-1991 female shocks	0.17 (0.28)			
Residuals of 1980-1991 male shocks		-0.15 (0.13)		
Residuals of 1991-2000 female shocks			1.93*** (0.45)	
Residuals of 1991-2000 male shocks				-0.71*** (0.23)
<u>Panel B: Wage growth residuals</u>				
Residuals of 1980-1991 female shocks	0.06 (0.12)			
Residuals of 1980-1991 male shocks		0.10 (0.08)		
Residuals of 1991-2000 female shocks			-1.53*** (0.43)	
Residuals of 1991-2000 male shocks				-0.26 (0.26)

Note: Robust standard errors clustered at the mesoregion level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

B.3 Model solutions appendix

This appendix describes the solutions of the model in greater detail.

B.3.1 Equilibrium under autarky

In the solution under autarky, regional population N_{jt} is assumed exogenous, and equilibrium is characterized by male labor, female labor, and housing markets clearing.

First I solve for the gender employment and wage gaps. Note that Equations 2.6 and 2.7 together yield a supply-side gender gap expression:

$$\frac{W_{Mjt}}{W_{Wjt}} = \frac{1}{1 + T_{jt}} \left(\frac{N_{Mjt}}{N_{Wjt}} \right)^{\frac{1}{\iota}} \quad (\text{B.1})$$

Combining equations B.1 and 2.3, I obtain:

$$\frac{N_{Mj}}{N_{Wj}} = \left[(1 + T_{jt}) \left(\frac{\psi_{Mjt}}{\psi_{Wjt}} \right)^{\sigma} \right]^{\frac{\iota}{1 - \iota(\beta\sigma - 1)}} \quad (\text{B.2})$$

$$\frac{W_{Mj}}{W_{Wj}} = \left(\frac{\psi_{Mjt}}{\psi_{Wjt}} \right)^{\frac{\sigma}{1 - \iota(\beta\sigma - 1)}} (1 + T_{jt})^{\frac{\iota(\beta\sigma - 1)}{1 - \iota(\beta\sigma - 1)}} \quad (\text{B.3})$$

These expressions in turn allow me to write the gender-specific inverse labor demand in terms of own-gender employment and exogenous parameters. To do this, I take the aggregate effective labor used by firms in region j as:

$$L_{jt} = \left[\left(\psi_{W_{jt}} N_{Wjt}^{\beta} \right)^{\sigma} + \left(\psi_{M_{jt}} N_{Mjt}^{\beta} \right)^{\sigma} \right]^{\frac{1}{\sigma}} \quad (\text{B.4})$$

which is a component of the production function (equation 2.1). Using B.3 I can re-write B.4 as

$$L_{jt} = N_{Wjt}^{\beta} \left[\psi_{W_{jt}}^{\sigma} + \psi_{M_{jt}}^{\sigma} \left[(1 + T_{jt}) \left(\frac{\psi_{Mjt}}{\psi_{Wjt}} \right)^{\sigma} \right]^{\frac{\beta\iota\sigma}{1 - \iota(\beta\sigma - 1)}} \right]^{\frac{1}{\sigma}}, \text{ or} \quad (\text{B.5})$$

$$L_{jt} = N_{Mj}^{\beta} \left[\psi_{W_{jt}}^{\sigma} \left[(1 + T_{jt}) \left(\frac{\psi_{Mjt}}{\psi_{Wjt}} \right)^{\sigma} \right]^{\frac{\beta\iota\sigma}{\iota(\beta\sigma - 1) - 1}} + \psi_{M_{jt}}^{\sigma} \right]^{\frac{1}{\sigma}} \quad (\text{B.6})$$

Labor market for females

Using equation B.5, female labor demand can be expressed as:

$$W_{Wjt} = \lambda_1 \psi_{Wjt}^\sigma N_{Wjt}^{\xi_1} \Psi_{Wjt}^{\frac{\gamma(1-\sigma)}{(1-\gamma)}} \quad (\text{B.7})$$

where $\Psi_{Wjt} := \left[\psi_{Wjt}^\sigma + \psi_{Mjt}^\sigma \left[(1 + T_t) \left(\frac{\psi_{Mjt}}{\psi_{Wjt}} \right)^\sigma \right]^{\frac{\beta\iota\sigma}{1-\iota(\beta\sigma-1)}} \right]^{\frac{1}{\sigma}}$, $\lambda_1 := \beta\gamma^{\frac{\gamma}{1-\gamma}} \bar{Z}^{\frac{1-\beta-\gamma}{1-\gamma}}$, and $\xi_1 := \frac{\beta\gamma(1-\sigma) + (1-\gamma)(\beta\sigma-1)}{(1-\gamma)}$.

Equating female labor demand in B.7 and labor supply in 2.6 yields equilibrium employment and wages:

$$\begin{aligned} N_{Wjt}^{*aut} &= N_{jt}^{\frac{1}{1-\iota\xi_1}} \left(\frac{\lambda_1}{1 + T_t} \right)^{\frac{\iota}{1-\iota\xi_1}} \psi_{Wjt}^{\frac{\iota\sigma}{1-\iota\xi_1}} \Psi_{Wjt}^{\frac{\gamma(1-\sigma)}{(1-\gamma)} \frac{\iota}{1-\iota\xi_1}} \\ W_{Wjt}^{*aut} &= N_{jt}^{\frac{\xi_1}{(1-\iota\xi_1)}} \left(\frac{\lambda_1}{1 + T_t} \right)^{\frac{1}{1-\iota\xi_1}} \psi_{Wjt}^{\frac{\sigma}{1-\iota\xi_1}} \Psi_{Wjt}^{\frac{\gamma(1-\sigma)}{(1-\gamma)} \frac{1}{1-\iota\xi_1}} \end{aligned}$$

Labor market for males

Using equation B.6, male labor demand can be written as:

$$W_{Mjt} = \lambda_1 \psi_{Mjt}^\sigma N_{Mjt}^{\xi_1} \Psi_{Mjt}^{\frac{\gamma(1-\sigma)}{(1-\gamma)}} \quad (\text{B.8})$$

where $\Psi_{Mjt} := \left[\psi_{Wjt}^\sigma \left[(1 + T_t) \left(\frac{\psi_{Mjt}}{\psi_{Wjt}} \right)^\sigma \right]^{\frac{\beta\iota\sigma}{\iota(\beta\sigma-1)-1}} + \psi_{Mjt}^\sigma \right]^{\frac{1}{\sigma}}$.

Equilibrium employment and wages for men follow from equating labor demand in B.8 and labor supply in 2.7:

$$\begin{aligned} N_{Mjt}^{*aut} &= N_{jt}^{\frac{1}{1-\iota\xi_1}} \lambda_1^{\frac{\iota}{1-\iota\xi_1}} \psi_{Mjt}^{\frac{\iota\sigma}{1-\iota\xi_1}} \Psi_{Mjt}^{\frac{\gamma(1-\sigma)}{(1-\gamma)} \frac{\iota}{1-\iota\xi_1}} \\ W_{Mjt}^{*aut} &= N_{jt}^{\frac{\xi_1}{(1-\iota\xi_1)}} \lambda_1^{\frac{1}{1-\iota\xi_1}} \psi_{Mjt}^{\frac{\sigma}{1-\iota\xi_1}} \Psi_{Mjt}^{\frac{\gamma(1-\sigma)}{(1-\gamma)} \frac{1}{1-\iota\xi_1}} \end{aligned}$$

Housing rents

Equation 2.9 can be re-written as:

$$R_{jt}^{*aut} = \left(\zeta \bar{W}_{jt}^{*aut} N_{jt} \right)^{\frac{1}{1+\rho}} \quad (\text{B.9})$$

with $\zeta := \frac{\alpha}{\bar{H}\left(\frac{1+r_t}{r_t}\right)^\rho}$. The net wage under autarky is in turn defined by the wage and employment equilibria in equations 2.13, 2.15, 2.14 and 2.16, and the average labor force participation costs of male and women in the workforce. Specifically:

$$\bar{W}_{jt}^{*aut} = \left(\frac{N_{Wjt}^{*aut}}{N_{jt}} W_{Wjt}^{*aut} - \bar{\varphi}_{Wjt} \right) + \left(\frac{N_{Mjt}^{*aut}}{N_{jt}} W_{Mjt}^{*aut} - \bar{\varphi}_{Mjt} \right)$$

The average participation costs correspond to the expected value for the population of each gender for whom their wages are weakly larger than the costs. Given the functional form assumption on $F(\varphi_i)$, these are give by:

$$\bar{\varphi}_{Wjt} = \frac{\iota}{\iota+1} (1 + T_{ij}) \left(\left(\frac{W_{Wjt}^{aut*}}{1 + T_{ij}} \right)^{\iota+1} - 1 \right) \quad (\text{B.10})$$

$$\bar{\varphi}_{Mjt} = \frac{\iota}{\iota+1} \left(\left(W_{Mjt}^{aut*} \right)^{\iota+1} - 1 \right) \quad (\text{B.11})$$

B.3.2 Equilibrium in the open region

When the region is open to labor migration, population becomes an endogenous variable. Under the spatial equilibrium assumption, migration arbitrages away household-level welfare differences across regions, such that household indirect utility equals the utility in the reservation region \underline{U} .

Spatial indifference curves and local population

Given the equilibrium rent equation in 2.9, the spatial indifference curve can be written as:

$$V_{jt}(\theta_j, \bar{W}_{jt}^{net}, N_{jt}) = \underline{U} = \zeta_t \theta_j \left(E(\bar{W}_{jt}^{net}) \right)^{\frac{1+\rho-\alpha}{1+\rho}} N_{jt}^{-\frac{\alpha}{1+\rho}}$$

where the net household wage enters the utility function as an expectation because, before migration, there is uncertainty about the individuals' participation costs. It is defined as the sum of the expected wage of men and women, namely $E(\bar{W}_{jt}^{net}) = E(\bar{W}_{Mjt}^{net}) + E(\bar{W}_{Wjt}^{net})$.

Households observe the distribution of labor force participation costs, and therefore know each of their members' probability of participating in city j given local wages, namely $\left(\frac{W_{Wjt}}{1+T_t} \right)^\iota$ for women and W_{Mjt}^ι for men, as well as the average costs of the people who

participate from equations B.10 and B.11. Combining these equations yields an expression for the expected net labor income for men and women in city j :

$$\begin{aligned} E(W_{Wj}^{net}) &= \left(\frac{W_{Wjt}}{1+T_t} \right)^\iota \left(W_{Wjt} - \frac{\iota(1+T_t)}{\iota+1} \left(\left(\frac{W_{Wjt}}{1+T_t} \right)^{\iota+1} - 1 \right) \right) \\ E(W_{Mj}^{net}) &= W_{Mjt}^\iota \left(W_{Mjt} - \frac{\iota}{\iota+1} (W_{Mjt}^{\iota+1} - 1) \right) \end{aligned}$$

Equilibrium outcomes

The spatial indifference curve can also be written as an expression for the local population in terms of expected household wages (equation 2.10). Using this and the solutions for the equilibrium under autarky, one can obtain equations that implicitly define the endogenous variables of the model in terms of the exogenous parameters. This in turn can be used to perform comparative static analysis of the effects of shocks to male and female local labor demand.

Because I can express male wages as a function of female wages and viceversa using equation B.3, I can write equations 2.11 and 2.12, and ultimately the population equation in 2.10 in terms of male wages (rather than in terms of expected net household wages):¹

$$N_{jt} = \left(\frac{\zeta \theta_j}{\underline{U}} \right)^{\frac{1+\rho}{\alpha}} \left[W_{Mjt}^\iota \left(W_{Mjt} \frac{\iota(1 - W_{Mjt}^{1+\iota})}{1+\iota} \right) + \left(\frac{\iota T_{Wt}(1 - \Phi_{jt}^{1+\iota})}{1+\iota} + T_{Wt}^\iota W_{Mjt} \left(\frac{\psi_{Mjt}}{\psi_{Wjt}} \right)^{\sigma_M} \right) \Phi_{jt}^\iota \right]^{\frac{1-\alpha+\rho}{\alpha}} \quad (\text{B.12})$$

where $T_{Wt} := 1 + T_t$, $\Phi_{jt} := T_{Wt}^{\iota_M - 1} W_{Mjt} \left(\frac{\psi_{Mjt}}{\psi_{Wjt}} \right)^{\sigma_M}$, $\iota_M := \frac{\iota(\beta\sigma - 1)}{\iota(\beta\sigma - 1) - 1}$, and $\sigma_M := \frac{\sigma}{\iota(\beta\sigma - 1) - 1}$.

Recall that the labor market solution under autarky expresses the male wage in terms of the population and exogenous parameters (equation 2.16). Replacing it into equation B.12 yields an expression that implicitly defines the population in terms of the exogenous parameters of the model.

This equation in turn can be used to obtain the equilibrium housing rents in the open region. To see this, notice that in equation 2.9 the product of household net wages and housing rents can be written as $\bar{W}_{jt}^{*aut} N_{jt} = N_{Wjt}^{*aut} W_{Wjt}^{*aut} + N_{Mjt}^{*aut} W_{Mjt}^{*aut} - N_{jt} (\bar{\varphi}_{Wjt} + \bar{\varphi}_{Mjt})$.

¹I can alternatively express the population in terms of female wages and exogenous parameters, obtaining equivalent solutions. I use the male wage expression because it yields a more succinct expression.

This in turn allows me to express population in terms of rents and the autarky solutions for employment and rents:

$$N_{jt} = \frac{N_{Wjt}^{*aut} W_{Wjt}^{*aut} + N_{Mjt}^{*aut} W_{Mjt}^{*aut} - \frac{R_{jt}^{1+\rho}}{\zeta}}{\bar{\varphi}_{Wjt} + \bar{\varphi}_{Mjt}} \quad (\text{B.13})$$

Replacing equation B.13 and the autarky employment and wage solutions in equations 2.13 through 2.16 into the open-city population equation yields an expression that implicitly defines local housing rents in terms of the exogenous parameters of the model.

Using similar processes I obtain equations that implicitly define gender-specific wages and employment. I take the autarky equilibrium male wage in equation 2.16 and replace N_{jt} with expression B.12, obtaining the equilibrium male wage in the open region. The employment solution involves two additional steps. First, I use the employment gap equilibrium under autarky in equation B.2 and combine it with the labor demand equation in 2.2 to express male wages in terms of male employment and exogenous parameters. Second, I plug in the resulting equation in the expression for the open-region equilibrium male wage. An equivalent procedure allows me to obtain equilibrium wages and employment for females. I omit reporting the full expressions for the sake of space, but they are available upon request.

Appendix C

Appendix to Chapter 3

C.1 Supplementary Figures for Chapter 3

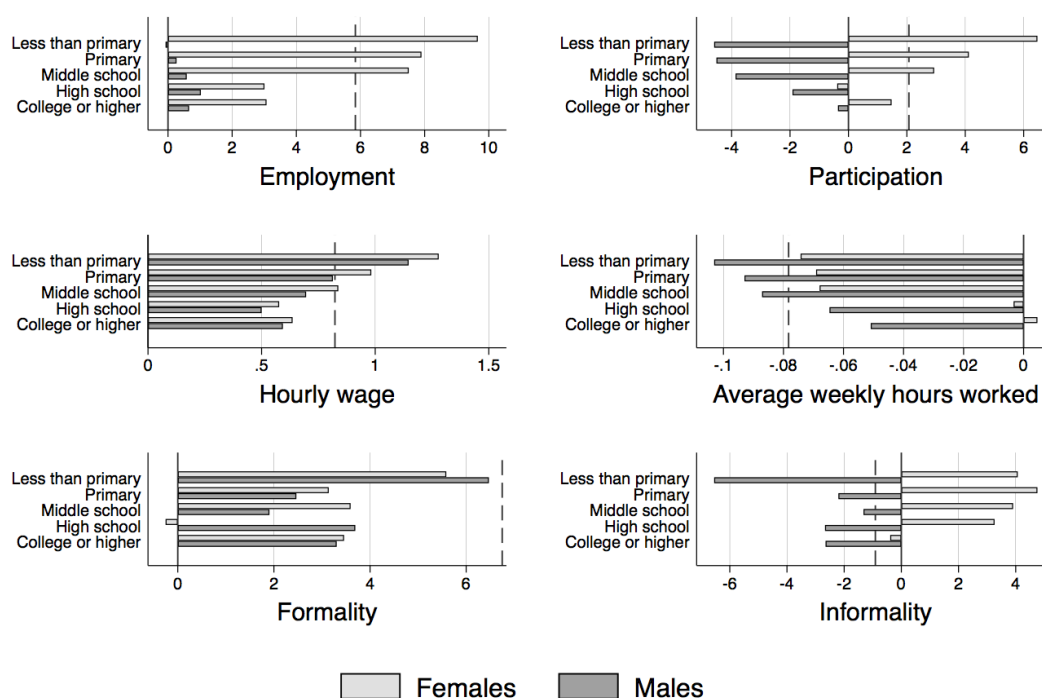


Figure C.1: *Changes in labor market outcomes 2000-2010 by educational attainment category and gender*

Note: Restricted to wage-earning population aged 23 through 64. Dashed lines denote population averages. All estimates are own calculations from microdata using sample weights. See the data appendix D for details on the measurement of each variable. **Source:** Population censuses of 2000 and 2010.

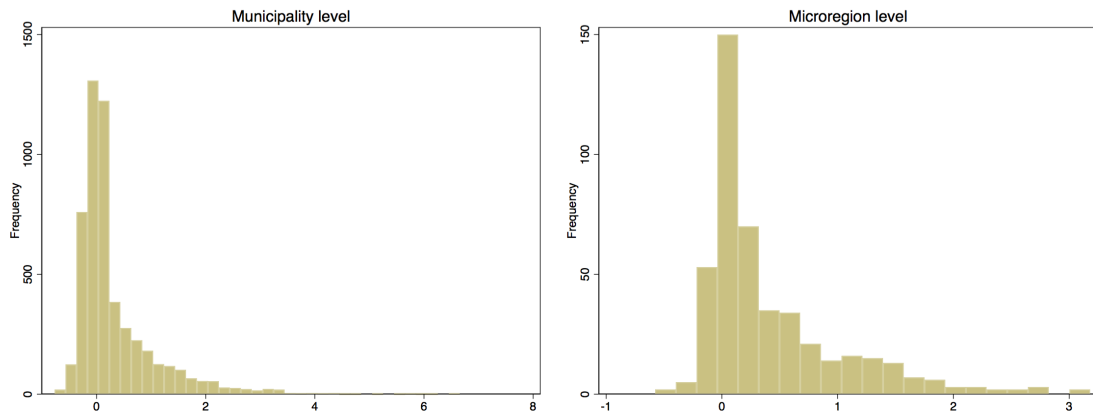


Figure C.2: *Distribution of FUNDEF shock across localities*

Sources: See data appendix.

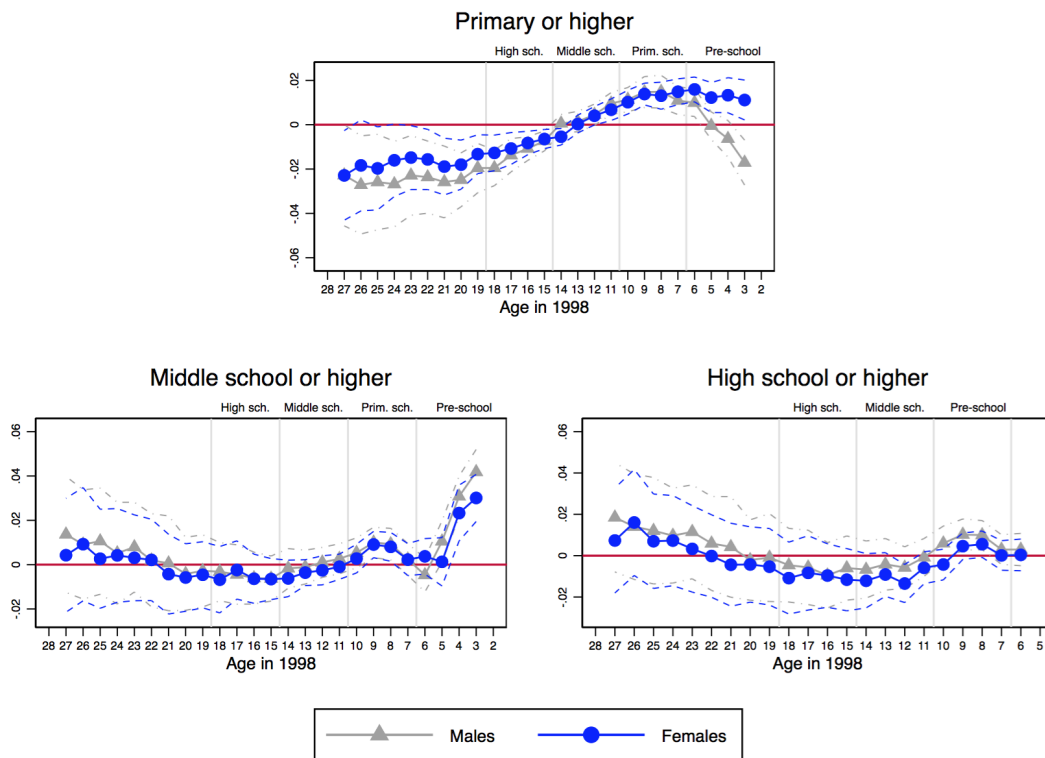


Figure C.3: *Effects of FUNDEF on individual's probability of reaching a specific educational attainment in 2010 by cohort and by gender*

Note: The markers represent the coefficient on the interaction of the FUNDEF treatment variable and each cohort dummy in equation 3.3. Dashed lines are 95% confidence intervals, with standard errors clustered at the municipality of education level. **Sources:** See data appendix.



Figure C.4: *Effects of FUNDEF on wages using alternative measures*

Note: The markers represent the coefficient on the interaction of the FUNDEF treatment variable and each cohort dummy in equation 3.3. Dashed lines are 95% confidence intervals, with standard errors clustered at the municipality of education level. **Sources:** See data appendix.

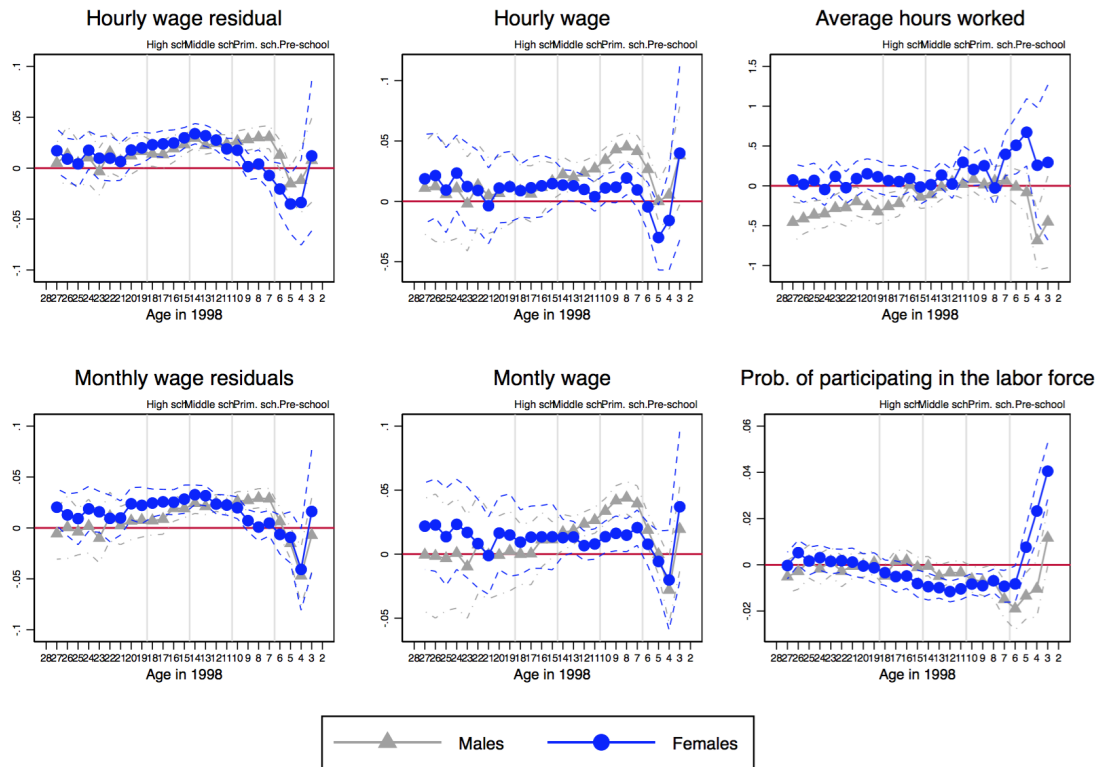


Figure C.5: *Effects of FUNDEF on wages and labor force participation by gender*

Note: The markers represent the coefficient on the interaction of the FUNDEF treatment variable and each cohort dummy in equation 3.3. Dashed lines are 95% confidence intervals, with standard errors clustered at the municipality of education level. **Sources:** See data appendix.



Figure C.6: *Effects of FUNDEF on wages by migrant status using alternative measures*

Note: The markers represent the coefficient on the interaction of the FUNDEF treatment variable and each cohort dummy in equation 3.3. Dashed lines are 95% confidence intervals, with standard errors clustered at the municipality of education level. **Sources:** See data appendix.

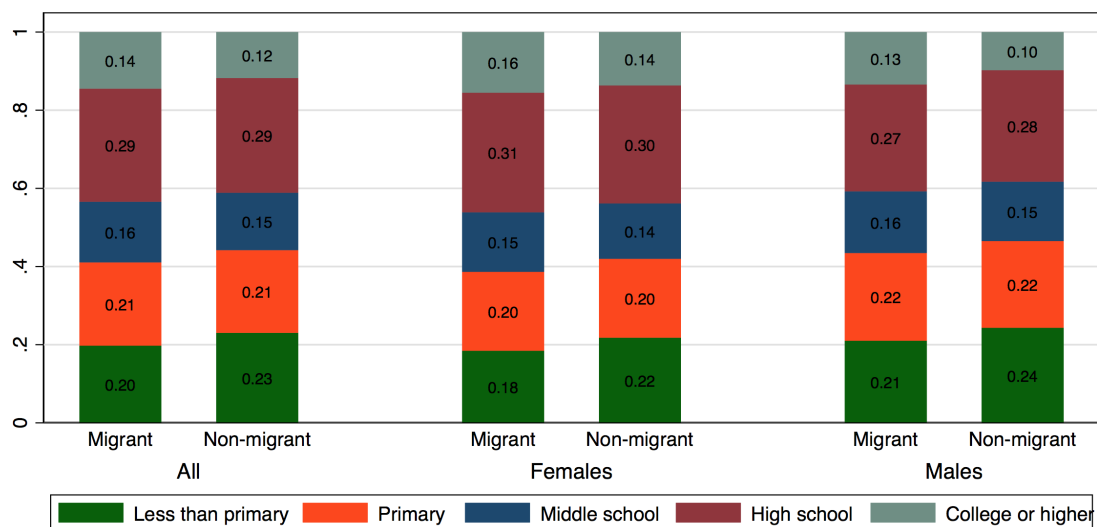


Figure C.7: *Percentage of Brazilian population age 23 or older in each educational attainment category by migrant status, 2010*

Note: Own calculations using census data. **Source:** Population censuses of 2000 and 2010.

C.2 Supplementary Tables for Chapter 3

Table C.1: *Individual summary statistics, 2010*

	Mean	Std. Dev.	Min	Max
Less than primary	0.15	0.36	0	1
Primary	0.20	0.40	0	1
Middle school	0.18	0.38	0	1
High school	0.37	0.48	0	1
College or higher	0.10	0.30	0	1
Hourly wage	7.389	21.924	0.002	8,660.505
Less than primary	4.966	17.704	0.003	3,464.204
Primary	4.570	11.347	0.002	3,464.204
Middle school	5.472	13.830	0.002	2,771.363
High school	6.675	16.389	0.002	4,618.936
College or higher	18.146	44.516	0.003	8,660.505
Monthly wage	1,082	2,208	1	800,000
Less than primary	662	920	1	102,010
Primary	674	1,267	1	600,000
Middle school	808	1,853	1	800,000
High school	996	1,570	1	400,000
College or higher	2,673	4,346	1	750,000
Probability of being a migrant	0.12	0.32	0	1
Less than primary	0.13	0.33	0	1
Primary	0.14	0.35	0	1
Middle school	0.14	0.35	0	1
High school	0.12	0.33	0	1
College or higher	0.17	0.38	0	1
Employment rate	0.66	0.47	0	1
Less than primary	0.56	0.50	0	1
Primary	0.59	0.49	0	1
Middle school	0.65	0.48	0	1
High school	0.72	0.45	0	1
College or higher	0.87	0.33	0	1

Continued on next page

Table C.1: *(continued)*

	Mean	Std. Dev.	Min	Max
Number of weekly hours worked	41.3	14.0	1	140
Less than primary	40.8	15.2	1	140
Primary	42.0	15.0	1	140
Middle school	42.2	14.5	1	140
High school	41.7	13.2	1	140
College or higher	39.1	12.3	1	140
Probability of participating in the labor force	0.73	0.44	0	1
Less than primary	0.62	0.48	0	1
Primary	0.66	0.47	0	1
Middle school	0.72	0.45	0	1
High school	0.80	0.40	0	1
College or higher	0.92	0.28	0	1
Probability of being formally employed *	0.54	0.50	0	1
Less than primary	0.39	0.49	0	1
Primary	0.40	0.49	0	1
Middle school	0.51	0.50	0	1
High school	0.62	0.49	0	1
College or higher	0.73	0.44	0	1
Probability of being informally employed*	0.36	0.48	0	1
Less than primary	0.51	0.50	0	1
Primary	0.50	0.50	0	1
Middle school	0.39	0.49	0	1
High school	0.27	0.45	0	1
College or higher	0.23	0.42	0	1
Probability of being unemployed*	0.10	0.29	0	1
Less than primary	0.10	0.29	0	1
Primary	0.11	0.31	0	1
Middle school	0.10	0.31	0	1
High school	0.10	0.31	0	1
College or higher	0.04	0.21	0	1

Source: Own calculations from 2010 population census using sampling weights.

* Probability conditional on participating in the labor force.

Table C.2: *Regional summary statistics, 2010*

	Mean	Std. Dev.	Min	Max
<u>Shocks (1997)</u>				
<i>Municipality level</i>				
Fundef shock	0.33	0.75	-0.77	6.66
Predicted FUNDEF shock	0.43	0.99	-0.73	10.61
<i>Regional level</i>				
Fundef shock	0.43	0.61	-0.58	3.18
Predicted FUNDEF shock	0.56	0.82	-0.23	4.61
<u>Main variables</u>				
Total population (1,000s)	291.76	769.38	4.95	13757.32
Working-age population (1,000s)	147.34	409.72	2.26	7443.07
Males	71.73	191.99	1.27	3528.61
Females	75.62	217.84	0.99	3914.46
Migrant share	0.25	0.12	0.03	0.71
Males	0.24	0.12	0.03	0.69
Females	0.25	0.12	0.03	0.72
Average hourly wage	2.70	1.19	0.70	7.40
Males	2.87	1.33	0.72	8.04
Females	2.40	1.04	0.36	6.69
Average montly wage	437.77	200.43	66.67	1217.12
Males	495.33	238.87	127.71	1417.59
Females	327.77	165.27	23.98	1002.02
Employment rate	0.54	0.08	0.22	0.81
Males	0.77	0.12	0.26	2.04
Females	0.30	0.13	0.04	0.72
Average weekly hours worked	42.73	3.26	32.36	53.96
Males	44.76	3.30	34.10	55.88
Females	38.33	3.35	25.47	50.79
Participation rate	0.58	0.09	0.27	0.82
Males	0.81	0.09	0.33	0.96
Females	0.34	0.14	0.04	0.74
Formality rate	0.31	0.17	0.01	0.82
Males	0.31	0.18	0.01	0.80
Females	0.35	0.15	0.02	0.87
Informality rate	0.62	0.17	0.16	0.97
Males	0.64	0.19	0.19	0.98
Females	0.55	0.14	0.11	0.95
Unemployment rate	0.07	0.05	0.00	0.26
Males	0.05	0.04	0.00	0.22
Females	0.10	0.08	0.00	0.77

Source: Own calculations with population censuses. Outcomes calculated for individuals aged 15-64. N=456.

Table C.3: *Regional summary statistics, 2000*

	Mean	Std. Dev.	Min	Max
<u>Main variables</u>				
Total population (1,000s)	317.72	815.62	12.95	12790.27
Working-age population (1,000s)	157.65	429.59	5.77	6859.49
Males	76.85	200.68	3.16	3216.30
Females	80.80	229.00	2.61	3643.20
Migrant share	0.19	0.08	0.03	0.62
Males	0.19	0.08	0.03	0.62
Females	0.19	0.08	0.03	0.62
Average hourly wage	2.56	0.96	0.92	6.57
Males	2.71	1.09	0.94	7.21
Females	2.29	0.78	0.78	5.77
Average montly wage	911.35	363.05	258.32	2375.99
Males	529.16	227.27	159.84	1417.59
Females	359.74	137.20	92.74	980.08
Employment rate	0.51	0.09	0.22	0.74
Males	0.71	0.11	0.26	0.99
Females	0.32	0.09	0.09	0.61
Average weekly hours worked	44.09	2.95	34.76	53.96
Males	46.48	3.04	38.84	55.88
Females	39.10	3.20	25.47	50.79
Participation rate	0.59	0.10	0.27	0.78
Males	0.78	0.10	0.33	0.91
Females	0.40	0.10	0.16	0.66
Formality rate	0.28	0.13	0.04	0.62
Males	0.28	0.14	0.03	0.64
Females	0.28	0.12	0.02	0.65
Informality rate	0.59	0.14	0.30	0.88
Males	0.63	0.15	0.31	0.93
Females	0.52	0.11	0.28	0.85
Unemployment rate	0.13	0.04	0.03	0.26
Males	0.09	0.04	0.01	0.22
Females	0.20	0.05	0.06	0.51
<u>1980s trends controls</u>				
Δ_{80-91} Formality rate	0.03	0.07	-0.18	0.29
Δ_{80-91} Average montly wage	0.05	0.16	-0.50	0.70

Source: Own calculations with population censuses. Outcomes calculated for individuals aged 15-64. N=456.

Table C.4: *Correlations among regional-level variables, 2010*

Variables	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1 Regional FUNDEF shock	1														
2 Regional predicted FUNDEF shock	0.99	1													
3 Total population	0.01	-0.04	1												
4 Working-age population	0.00	-0.05	1.00	1											
5 Share of migrants in working-age pop.	-0.19	-0.19	0.10	0.11	1										
6 Average log hourly wage	-0.38	-0.41	0.33	0.34	0.66	1									
7 Average log wage	-0.42	-0.45	0.31	0.32	0.64	0.98	1								
8 Employment rate	-0.40	-0.41	0.17	0.18	0.55	0.65	0.68	1							
9 Participation rate	-0.40	-0.41	0.22	0.23	0.53	0.73	0.77	0.94	1						
10 Average monthly hours worked	-0.43	-0.43	0.01	0.01	0.13	0.08	0.25	0.40	0.42	1					
11 Formality rate	-0.47	-0.50	0.25	0.26	0.51	0.74	0.73	0.67	0.68	0.31	1				
12 Informality rate	0.44	0.48	-0.27	-0.28	-0.44	-0.76	-0.76	-0.57	-0.67	-0.29	-0.96	1			
13 Unemployment rate	0.05	0.05	0.11	0.10	-0.33	0.06	0.19	-0.28	0.07	-0.09	-0.03	-0.26	1		
14 Change 1980-1991 in formality rate	-0.08	-0.10	-0.25	-0.25	0.02	-0.15	-0.13	-0.01	-0.04	0.07	-0.12	0.13	-0.05	1	
15 Change 1980-1991 in average log wage	0.00	-0.02	0.12	0.12	0.16	0.26	0.21	0.16	0.17	0.01	0.19	-0.19	0.02	0.03	1

Note: Own-calculations with population census data. **Sources:** See data appendix.

Table C.5: *Correlations among regional-level variables, 2000*

Variables	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1 Regional FUNDEF shock	1														
2 Regional predicted FUNDEF shock	0.99	1													
3 Total population	0.01	-0.04	1												
4 Working-age population	0.00	-0.05	1.00	1											
5 Share of migrants in working-age pop.	-0.13	-0.13	0.05	0.05	1										
6 Average log hourly wage	-0.48	-0.53	0.39	0.39	0.42	1									
7 Average log wage	-0.51	-0.55	0.36	0.37	0.46	0.99	1								
8 Employment rate	-0.51	-0.53	0.15	0.16	0.36	0.69	0.73	1							
9 Participation rate	-0.55	-0.57	0.23	0.23	0.40	0.76	0.78	0.96	1						
10 Average monthly hours worked	-0.46	-0.45	0.05	0.05	0.45	0.47	0.56	0.61	0.63	1					
11 Formality rate	-0.53	-0.57	0.23	0.24	0.25	0.82	0.82	0.74	0.80	0.50	1				
12 Informality rate	0.52	0.56	-0.27	-0.27	-0.25	-0.78	-0.78	-0.59	-0.72	-0.45	-0.96	1			
13 Unemployment rate	-0.04	-0.06	0.19	0.18	0.04	0.06	0.03	-0.36	-0.10	-0.08	0.07	-0.36	1		
14 Change 1980-1991 in formality rate	-0.08	-0.10	-0.26	-0.26	0.03	-0.21	-0.17	-0.02	-0.05	0.08	0.00	0.02	-0.07	1	
15 Change 1980-1991 in average log wage	0.00	-0.02	0.13	0.12	0.19	0.30	0.29	0.12	0.16	0.03	0.19	-0.20	0.09	0.03	1

Note: Own-calculations with population census data. Sources: See data appendix.

Table C.6: *Effects of FUNDEF on migrant population by educational attainment group*

	All (1)	Males (2)	Females (3)
Less than primary	-0.206*** (0.042)	-0.214*** (0.044)	-0.196*** (0.041)
Primary or higher	-0.045* (0.026)	-0.025 (0.030)	-0.064*** (0.024)
Middle school or higher	0.044** (0.022)	0.027 (0.027)	0.062*** (0.022)
High school or higher	0.047* (0.026)	0.014 (0.030)	0.076*** (0.025)
College or higher	0.201*** (0.043)	0.096* (0.050)	0.287*** (0.047)

Note: The table reports the coefficients on the treatment variable in equation 3.4. Regressions are at the microregion level (N=456). Robust standard errors clustered at the mesoregion level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

Table C.7: *Effects of FUNDEF on local education attainment controlling for 1990s' migration composition*

	Change in share of educated in adult population			
	2000-2010		1991-2000 (placebo test)	
	Mid-school (1)	High-school (2)	Mid-school (3)	High-school (4)
All individuals	0.041*** (0.007)	0.020*** (0.005)	-0.001 (0.005)	-0.003 (0.003)
Males only	0.034*** (0.007)	0.013*** (0.005)	-0.002 (0.005)	-0.003 (0.003)
Females only	0.049*** (0.008)	0.027*** (0.006)	-0.001 (0.005)	-0.004 (0.003)
Changes in formality rates and wages in the 1980s	Yes	Yes	Yes	Yes
1990s migration as a share of local population	Yes	Yes	Yes	Yes
Share of high-school educated in 1990s migrants	Yes	Yes	Yes	Yes

Note: The table reports the coefficients on the treatment variable in equation 3.4. Regressions are at the microregion level (N=456). Robust standard errors clustered at the mesoregion level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1. **Sources:** See data appendix.

Appendix D

Data Appendix

Table D.1: *Databases Used*

Acronym	Database	Years	Source
<i>International</i>			
WDI	World Development Indicators, The World Bank	1960-2014	http://databank.worldbank.org
<i>U.S.A.</i>			
ACS	American Community Survey 5-Year sample (5-in-100 national random sample of the population)	2006-2010	https://usa.ipums.org/usa
USPC	Tabulated data from the Population Census, which mostly come from the National Historical Geographic Information System, which in turn compiles data from the U.S. Census.	2010, 2000, 1990, 1980	https://www.nhgis.org/
BEA	Income data rural/urban from Bureau of Economic Analysis	2010	bea.gov/iTable/index_nipa.cfm
CCDB	City and County Data book (amenities variables)	1990	From County and City Data book 1994 (hard copy only)
BRFSS	Behavioral Risk Factor Surveillance System (BRFSS) conducted by the Centers for Disease Control		From Glaeser <i>et al.</i> (2014a)

Continued on next page

Table D.1: (continued)

Acronym	Database	Years	Source
<i>Brazil</i>			
BPC	Population Census micro-data sample	2010, 2000, 1991, 1980	fflch.usp.br/centrodametropole
IPD	IPEA Municipal areas and climate data	2010	www.ipeadata.gov.br
ETDB	Evolution of territorial division of Brazil - municipalities microdata.	1980-2010	http://downloads.ibge.gov.br/
IPUMS	Integrated Public Use Microdata series, international (IPUMS). Brazilian census	2010, 2000, 1991, 1980	https://international.ipums.org This sample is smaller than the official census sample, and geographic identifiers are less disaggregated. However, it includes homogenized variables that are comparable across time and across countries.
IPEA1	IPEA - Municipality areas	2010	www.ipeadata.gov.br
IPEA2	IPEA - Climate data	2002	www.ipeadata.gov.br
IBGE1	IBGE - Municipality Borders GIS files	2010	https://mapas.ibge.gov.br/bases-e-referenciais/bases-cartograficas/malhas-digitais.html
IBGE2	IBGE - Evolution of municipality borders over census years	1872-2010	www.ibge.gov.br/home/geociencias/geografia/default_evolucao.shtm
IBGE3	IBGE - National consumer price index	1980-2010 (monthly)	ww2.ibge.gov.br/home/estatistica/indicadores/precos/inpc_ipca/default_seriesHist.shtm
SC	INEP - Brazilian school census	1997,1998	portal.inep.gov.br/microdados
TRES	STN - Brazilian National and State Treasuries	1997,1998	tesouro.fazenda.gov.br
<i>China</i>			
CPSS	1% Population Sample Survey	2005	
CPC	China County Population Census Data with GIS Maps - All China Marketing Research Co., Ltd.	2010, 2000, 1990, 1982	All China Marketing Research Co., Ltd., an agent for non-confidential data collected by the National Bureau of Statistics of China.
CHIPS	China Household Income Project (CHIP) Survey Urban Dataset	2007	http://www.ciadbnu.org/chip
CCSY	China City Statistical Yearbook	2005, 2002	

Continued on next page

Table D.1: *(continued)*

Acronym	Database	Years	Source
PSY	Provincial Statistical Yearbook		
CMA	China Meteorological Administration		
<i>India</i>			
IHDS	The India Human Development Survey IHDS-II	2012	www.ihds.umd.edu
IPC	Census tables, Census of India (Office of the Registrar General & Census Commissioner).	2011, 2001, 1991	www.censusindia.gov.in
IPC2	India: District Boundaries. ML InfoMap Pvt. Ltd., New Delhi.	2011, 2001, 1991, 1981	ML InfoMap sources the data from the Census of India (Office of the Registrar General & Census Commissioner).
INSTAT	IndiaStat	various	http://www.indiastat.com

Table D.2: *Country-level variables (Chapter 1)*

Variable	Country	Sources	Samples	Comments
Urban population (% of total)	All	WDI	USA, Brazil, China and India; all years available.	Refers to people living in urban areas as defined by each national statistics office.
Income share held by quantile	All	WDI	USA (1991, 2000, 2010), Brazil (1990, 2001, 2011), China (1990, 1999, 2010), and India (1993, 2004, 2009).	Selected years are those for which figures are based on current data as opposed to projections.
Income levels in 2014 \$ PPP	USA	BEA	1990, 2000, 2010	Income is expressed in 2014 dollars using WDI's consumer price index when appropriate, and then transformed to PPP dollars using WDI's conversion factor.
	BRA	BPC, WDI	1991, 2000, 2010	Income aggregates estimated from the census micro data in local currency, expressed in 2014 reals using cruzeiro-real current exchange and WDI's consumer price index when appropriate, and then transformed to PPP dollars using WDI's conversion factor.

Continued on next page

Table D.2: *(continued)*

Variable		Sources	Samples	Comments
	CHN	PSY	1990, 1999, 2010	Income expressed in 2014 reals using yuan-real current exchange and WDI's consumer price index when appropriate, and then transformed to PPP dollars using WDI's conversion factor.
	IND	IHDS	2005, 2011	Income aggregates estimated from the IHDS micro data in local currency, expressed in 2014 reals using rupee-real current exchange and WDI's consumer price index when appropriate, and then transformed to PPP dollars using WDI's conversion factor.
Migrants in the last 5 years (% of population)	USA	ACS, USPC	1990, 2000, 2010	
	BRA	BPC-IPUMS	National aggregates 1991, 2000, 2010	National aggregates of the homogenized variable "migrate5" from IPUMS, which refers to the person's place of residence 5 years ago.
	CHN	CPC	2000, 2010	
	IND	INSTAT	1993, 2001, 2011	This variable refers to migration in the past four years, the only data available.

Table D.3: *Individual-level variables (Chapter 1)*

Variable	Country	Sources	Samples	Comments
Wages (income)	USA	ACS	All males ages 25 to 55 micro-data sample, 2010.	Annual wage income.
	BRA	PBC	All males ages 25 to 55 micro-data sample, 2010.	Annualized monthly labor income from main occupation.
	CHN	CPSS	All males ages 25 to 55 micro-data sample, 2005.	Annualized monthly income.

Continued on next page

Table D.3: *(continued)*

Variable		Sources	Samples	Comments
	IND	IHDS	All males ages 25 to 55 micro-data sample, 2011.	Income per capita.
Housing rents	USA	ACS	All renter households with rent data.	Annualized monthly gross rent.
	BRA	BPC	All renter households with rent data.	Annualized monthly housing rent.
	CHN	CHIPS	All renter households with rent data.	Annualized estimated rental value of the property.
	IND	IHDS	All renter households with rent data.	Annualized monthly housing rent.
Housing value	USA	ACS	All owner-occupied households with value data.	Housing value.
	CHN	CPSS	All owner-occupied households with value data	Housing value.
Log real wage	USA	ACS	All males ages 25 to 55 micro-data sample, 2010.	Ln of wages - 0.33 x Ln of housing rents
	BRA	BPC	All males ages 25 to 55 micro-data sample, 2010	Ln of wages - 0.33 x Ln of housing rents
	CHN	CPSS	All males ages 25 to 55 micro-data sample, 2005.	Ln of wages - 0.33 x Ln of housing rents
	IND	IHDS	All males ages 25 to 55 micro-data sample, 2011.	Ln of wages - 0.33 x Ln of housing rents
Log wage residual	USA	ACS	Employed males aged 25-55 who work full-time and earn more than half the federal minimum wage for a full-time worker.	We run a regression of the log of wage and salary income on age, race, and education controls, calculate the residuals, and take the average at the MSA level.

Continued on next page

Table D.3: *(continued)*

Variable		Sources	Samples	Comments
	BRA	BPC	All workers with wage data.	We run a regression of the log of wage as the dependent variable, on education and age group controls, calculate the residuals, and take the average at the microregion level.
	CHN	CPSS	Males aged 25-55 with urban Hukou.	We run a regression the log of income on age and education controls, calculate the residuals, and take the average at the city level.
Log wage residual	IND	IHDS	Males aged 25-55.	We run a regression of the log of earnings on age and education controls, calculate the residuals, and take the average at the district level.
Log rent residual	USA	ACS	All renter households.	We run a regression of the log of monthly contract rent as the dependent variable on dwelling characteristics controls. We then calculate the residuals and take the average at the MSA level.
	BRA	BPC	All renter households with rent data.	We run a regression of the log of annual rent on dwelling characteristics controls. We then calculate the residuals and take the average at the microregion level.
	CHN	CHIPS	All renter households that report rent data.	We run a regression of the log of rent as the dependent variable on dwelling characteristics controls. We then calculate the residuals and take a weighted average at the city level.
	IND	IHDS	All renter households that report rent data.	We run a regression of the log of rent on dwelling characteristics controls. We then calculate the residuals and take an average at the district level.
Log housing price residual	USA	ACS	All owner-occupied households with price data.	We run a regression of the log of housing price as the dependent variable on dwelling characteristics controls. We then calculate the residuals and take the average at the MSA level.

Continued on next page

Table D.3: *(continued)*

Variable		Sources	Samples	Comments
	CHN	CPSS	All owner-occupied households with price data.	We run a regression of the log of housing price as the dependent variable on dwelling characteristics controls. We then calculate the residuals and take the average at the MSA level.
Education group controls	USA	ACS	All workers with wage and schooling data.	Individuals are classified in four educational categories: None or some grammar, grammar, high school, or college. Dummies for these schooling groups are used as controls.
	BRA	BPC	All workers with wage and schooling data.	Individuals are classified in five educational categories: less than primary school, primary school, junior high, high school, and college and higher. Dummies for these schooling groups are used as controls.
	CHN	CPSS	All workers with wage and schooling data.	Individuals are classified in four educational categories: Elementary, junior high, senior high, and higher education. Dummies for these schooling groups are used as controls.
	IND	IHDS	All workers with wage and schooling data.	Individuals are classified in five educational categories: None, elementary, secondary, high secondary, or higher education. Dummies for these schooling groups are used as controls.
Age (demographic) group controls	USA	ACS	All workers with wage data.	Individuals are classified in five-years age groups. Dummies for these age groups are used as controls in calculations.
	BRA	BPC	All workers with wage data.	Individuals are classified in five-years age groups, except for two larger groups including, respectively, individuals younger than 15 and individuals older than 65. Dummies for these age groups are used as controls in calculations.
	CHN	CPSS	All workers with wage data.	Individuals are classified in five-years age groups. Dummies for these age groups are used as controls in calculations.

Continued on next page

Table D.3: *(continued)*

Variable		Sources	Samples	Comments
	IND	IHDS	All workers with wage data.	Individuals are classified in five-years age groups. Dummies for these age groups are used as controls in calculations.
Dwelling characteristics controls	USA	ACS	All renter and owner-occupied households that report these variables.	Number of bedrooms, and number of rooms, age of the structure, units in the structure.
	BRA	BPC	All renter households that report these variables.	Number of rooms; number of bedrooms; type of dwelling (house, apartment, etc); predominant material in dwelling's external wall; and water provision mechanism (general network, well, etc.).
	CHN	CPSS, CHIPS	All renter and owner-occupied households that report these variables	Number of rooms; source of the house (self-built, commercial housing, affordable housing and public owned housing purchased after housing reforms, commercial renting, and low renting); existence of plumbing; existence of a bathroom; structure (includes reinforced concrete, brick-wood, wood/bamboo/grass), and number of stories.
	IND	IHDS	All renter households that report these variables.	Number of rooms; building house type (house with no shared walls; house with shared walls, flat, chawl, slum housing, or other); housing surrounded by sewage (yes/no); housing surrounded by stagnant water (yes/no); animals (no animals, animals inside living area, animals in an attached room, animals outside); predominant wall type; predominant roof type; and predominant floor type.

Table D.4: *Individual-level variables (Chapters 2 and 3)*

Variable	Samples	Description and comments
Montly Wage	BPC 1980, 1991, 2000 and 2010; IBGE3.	Monthly labor income in main occupation in the reference period, in 2010 reais.* **
Log montly wage residual	BPC 1980, 1991, 2000 and 2010.	Residuals of an individual-level regression of the log of wage on individual characteristics including age categories, schooling categories, sex and race. Regressions are restricted to the correspondent subpopulation (e.g. female wage residuals are estimated using only female workers observations). All regressions use sample weights provided in the IBGE microdata samples.
Weekly hours worked	BPC 2000 and 2010.	Usual number of hours worked at main job during the reference week (variables V0653 in the 2010 census and V0453 in the 2000 census.)* **
Hourly wage	BPC 2000 and 2010.	Montly wage divided by 4.33, and then by the weekly hours worked.* **
Log hourly wage residual	BPC 2000 and 2010.	Residuals of an individual-level regression of the log of the hourly wage on individual characteristics (same procedure as in the monthly wage residuals calculations)* **
Participant	BPC 2000 and 2010.	Individual that is either formally employed, informally employed or unemployed.** ***
Formally employed	BPC 1980.	Individual that worked over the period of reference as employee and contributed to social security, or was an employer.***
	BPC 1980, 1991, 2000 and 2010.	Individual that worked over the period of reference with a signed work card or as civil-service employede, or was an employer.** ***
Informally employed	BPC 1980.	Individual that worked over the period of reference as employee and did not contribute to social security, or was self-employed.
	BPC 1980, 1991, 2000 and 2010.	Individual that worked over the period of reference as a private sector or domestic employee without a signed work card, or was self-employed.**
Employed	BPC 1980, 1991, 2000 and 2010.	Individual either formally or informally employed.
Unemployed	BPC 1980, 1991, 2000 and 2010.	Individual that declared that they looked for employment but were not employed over the period of reference.**
Migrant (chapter 2)	BPC 2000, 2010.	Individual that declares that its time of residence in their current municipality is less or equal to 10 years (numerical response in variable V0416 in 2000 and V0624 in 2010).****

Continued on next page

Table D.4: *(continued)*

Variable	Samples	Description and comments
Migrant (chapter 3)	BPC 2000, 2010.	Individual that declares that its time of residence in their current municipality is less or equal to the year they finished schooling (numerical response in variable V0416 in 2000 and V0624 in 2010).
High-school educated	BPC 1980, 1991.	Individuals that completed at least high-school-equivalent education (2do grau, colegial o medio 2do ciclo) based on variables V523 and V524 in 1980, V0328 and V3241 in 1991.
	BPC 2000.	Individuals that completed at least high-school-equivalent education (2do grau, antigo classico, cientifico, etc. completed) based on variables V0432 and V4300 in 2000.
	BPC 2010.	Individuals that completed at least high-school-equivalent education education (regular or supletivo de ensino medio, antigo classico, cientifico, etc. completed) based on variables V0633 and V0634.
Rent	BPC 1991, BPC 2010, IBGE3.	Monthly value of housing rent.*
Rent residual	BPC 1991, BPC 2010.	Residuals of an household-level regression of the log of rent on individual housing unit characteristic including number of rooms, number of bedrooms, dwelling type, walls' material, and water source. Regressions are restricted to households that pay positive rents All regressions use sample weights provided in the IBGE microdata samples.
Industry of employment	BPC 1980, 1991, 2000 and 2010.	Industry code for employed workers (from Dix-Carneiro and Kovak 2017.)
Major industry of employment	BPC 1980, 1991, 2000 and 2010.	Four major industries based on CNAE - Domiciliar definition (Agriculture, Manufacturing, Services and Government.)

* All monetary values are expressed in 2010 reais. Variables are converted from prior currencies to reais and deflated using the national consumer price index (INCP) provided by the IBGE. The original INCP deflators are adjusted to account for inconsistencies derived from a dual-currency period in 1994, following the method proposed by Corseuil and Foguel (2002).

** The reference period changed between the censuses up to 1991 (when it was defined as the prior 12 months before the survey) and the censuses of 2000 and after (when it was defined as the prior week before the survey.)

*** Civil service employees and employers are excluded from the computations of the regional-level aggregate labor-market variables.

**** In all microregion-level aggregates the migrant definition is adjusted, to the extent the data allows, in order to include only those who lived in a different microregion before migrating (i.e. the definition excludes migrants from a different municipality within the same microregion). This correction is based on variables V4250 in 2000 (which only provides region of residence 5 years earlier) and V6254 in 2010.

Table D.5: *Area-level variables (Chapter 1)*

Variable	Country	Sources	Samples	Comments
Areas	USA	USPC	2010	
	BRA	BPC	2010	For cross-section regressions, tables, and graphs, we use 2010 Microregions as defined by the IBGE.
	CHN	CPC	2010	
	IND	IPC2	2010	
Time-consistent areas	USA	USPC	2010	For growth regressions we use time-consistent MSAs. Counties' current information is aggregated using 1980s' MSA identifiers for the 1980-2010 regressions.
	BRA	BPC, ETDB	All municipalities' 1980, boundaries for the 1991, 2000 and 2010 censuses.	For growth regressions we use time-consistent Microregions. First we construct Minimum Comparable Areas (MCAs) for the period 1980-2010 (as in Reis <i>et al.</i> , 2007), and then we group MCAs into Minimum Comparable Microregions (MCMs), following Kovak (2013).
	CHN	CPC	All counties' boundaries for the 2010, 2000, 1990 and 1982 censuses.	For growth regressions we use time-consistent cities. Current counties and districts' shape files are geo-matched to 1982 boundaries, and the information is aggregated using 1982 city boundaries.
	IND	IPC2	All district's boundaries for the 2011, 2001, 1991 and 1981 censuses.	For growth regressions we use time-consistent districts. Current districts' shape files are geo-matched to 1981 boundaries, and the information is aggregated using 1981 district boundaries.
Population / Urban population	USA	USPC	All MSAs	In the US census and official surveys any person living in a county belonging to a MSA is considered an urban dweller.
	BRA	BPC	All Microregions identified in the census	The Brazilian Population census counts as urban dweller every person living in an area that is a municipal seat ("city"), district seat ("town") or "isolated urban area".
	CHN	CPC	All Cities identified in the census	The Census of China defines as urban all individuals holders of urban hukou.

Continued on next page

Table D.5: *(continued)*

Variable		Sources	Samples	Comments
	IND	IPC	All districts identified in the census	The census of India defines as an urban are “all places with a Municipality, Corporation or Cantonment or Notified Town Area”; and all other places that have: “a) a minimum population of 5,000; b) at least 75% of the male working population was non-agricultural; c) adensity of population of at least 400 sq. Km. (i.e. 1000 per sq. Mile.)
Population Density	USA	ACS, USPC	All MSAs	Average population by square mile.
	BRA	BPC, IPD	All Microregions identified in the census	Average population by square kilometer. Microregion areas are aggregates of constituent municipalities’ areas.
	CHN	CPC	All Cities identified in the census	Average population by square kilometer.
	IND	IHDS, INSTAT	All districts identified in the survey	Average population by square kilometer.
Average log wage / income per capita	USA	USPC	All MSAs	Log of the median household income by MSA.
	BRA	BPC	All Microregions identified in the census, 2010 and 1980.	Weighted average (over all urban income-earning workers in the region) of the log of the annualized monthly labor income from main occupation.
	CHN	CHIPS	All cities identified in the CHIPs dataset	Log of the average disposable income of urban workers by city.
	IND	IHDS-II	All districts identified in the IHDS datasets	Log of the median income per capita of urban workers in the district.
Average log wage residual	USA	ACS	All MSAs identified in the ACS microdata.	Weighted average at the MSA level of the urban individual log wage residuals.
	BRA	BPC	All Microregions identified in the census, 2010.	Weighted average (over all urban income-earning workers in the region) of the individual log wage residuals.
	CHN	CPSS	All cities identified in the microdata, 2005.	Weighted average at the city level of the urban individual log wage residuals.

Continued on next page

Table D.5: *(continued)*

Variable		Sources	Samples	Comments
	IND	IHDS	All districts identified in the microdata, 2011	Weighted average of the district level of the urban individual log wage residuals
Average log rent residual	USA	ACS	All MSAs identified in the ACS microdata, 2010.	Weighted average (over urban households) of the individual log rent residuals.
	BRA	BPC	All Microregions identified in the census, 2010.	Weighted average (over urban households) of the individual log rent residuals.
	CHN	CHIPS	All cities identified in the microdata, 2005.	Weighted average (over urban households) of the individual log rent residuals.
	IND	IHDS	All districts identified in the microdata, 2011.	Weighted average (over urban households) of the individual log rent residuals.
Average log housing price residual	USA	ACS	All MSAs identified in the ACS microdata, 2010.	Weighted average (over all urban homeowner households) of the individual log value residuals.
	CHN	CPSS	All cities identified in the microdata, 2005.	Weighted average (over all urban homeowner households) of the individual log value residuals.
BA share	USA	USPC 1980, ACS 2010	All MSAs identified in the ACS microdata, 2010.	Fraction of urban population age 25 or higher that completed BA-equivalent university degree or higher.
	BRA	BPC	All Microregions identified in the census, 2010 and 1980.	Fraction of urban population age 25 or higher that completed BA-equivalent university degree or higher.
	CHN	CPC 1982, 2010	All cities identified in the census, 1982 and 2010.	Fraction of urban population age 25 or higher that completed BA-equivalent university degree or higher.
	IND	IPC 1990	All districts identified in the census, 1990.	Fraction of urban population age 25 or higher that completed BA-equivalent university degree or higher.

Continued on next page

Table D.5: *(continued)*

Variable		Sources	Samples	Comments
Absolute difference from ideal temperature by season	USA	CCBD	All MSAs identified in the Census, 1990	Summer temperature (July) and winter temperature (July) in Celsius, expressed as the absolute difference from the “ideal” temperature (assumed to be 21.11 Celsius or 70 Fahrenheit). When used as controls the “raw” variables (as opposed to the deviations from the ideal) are used.
	BRA	IPD	All Microregions identified in the census, 2010	Municipal-level figures are averaged at the Microregions level. Summer temperature (December-February) and average winter temperature (June-August) in Celsius, expressed as the absolute difference from the “ideal” temperature (assumed to be 21.11 Celsius or 70 Fahrenheit). When used as controls the “raw” variables (as opposed to the deviations from the ideal) are used.
	CHN	CMA	All cities, 2005	Summer temperature (July) and winter temperature (July) in Celsius, expressed as the absolute difference from the “ideal” temperature (assumed to be 21.11 Celsius or 70 Fahrenheit). When used as controls the “raw” variables (as opposed to the deviations from the ideal) are used.
	IND	INSTAT	All districts, 2011	Maximum and minimum yearly temperatures in Celsius, expressed as the absolute difference from the “ideal” temperature (assumed to be 21.11 Celsius or 70 Fahrenheit). When used as controls the “raw” variables (as opposed to the deviations from the ideal) are used.
Average annual rainfall	USA	CCBD	All MSAs identified in the Census, 1990	
	BRA	IPD	All Microregions identified in the census, 2010	Municipal-level figures are averaged at the Microregions level.
	CHN	CMA	All cities, 2005	
	IND	INSTAT	All districts, 2011	

Continued on next page

Table D.5: (continued)

Variable		Sources	Samples	Comments
Life satisfaction	USA	BRFSS	367 MSAs covered by BRFSS	We use the BLUP variable from Glaeser, Gottlieb, and Ziv (Glaeser <i>et al.</i> , 2014a)
	CHN	CHIPS	All cities identified in the CHIPS dataset	The source question asks: “Generally speaking, do you feel happy?” with possible responses being very happy, happy, so-so, not very happy, not happy at all, and don’t know. The CHIPS dataset gives “happy” a value of 1, and “not happy at all” a value of 5. We code these values so “happy” is 5 and drop the “don’t know” observations. After controlling for the exogenous demographic variables of age and race, we estimate a best linear unbiased predictions (BLUPs) of the random effects. We then take the urban average of the BLUPs by city to use as our happiness variable.
Life satisfaction	IND	IHDS-II	All districts identified in the IHDS datasets	The source question asked “What is the Level of satisfaction with economic situation (0 to 2)”, with “2” being the most satisfied. After controlling for the exogenous demographic variables of age and race, we estimate a best linear unbiased predictions (BLUPs) of the random effects. We then take the urban average of the BLUPs by city to use as our happiness variable.
Urban Population and Density IV1	USA	USPC	All time-consistent MSAs 1910-2010	1980 Log of urban population or density.
	BRA	BPC	All time-consistent Microregions 1980-2010	1980 Log of urban population or density.
	CHN	CPC	All time-consistent cities 1950-2010	1980 Log of urban population or density.
	IND	IPC	All time consistent districts 1951-2011.	1980 Log of urban population or density.
Urban Population and Density IV2	USA	USPC	All time-consistent MSAs 1900-2010	1900 Log of urban population or density.
	BRA	BPC	All time-consistent Microregions 1980-2010	1920 Log of urban population or density.

Continued on next page

Table D.5: *(continued)*

Variable		Sources	Samples	Comments
BA share IV1	CHN	CPC	All time-consistent cities 1950-2010	1950 Log of urban population or density.
	IND	IPC	All time consistent districts 1951-2011.	1951 Log of urban population or density.
	USA	USPC	All time-consistent MSAs 1940-2010	BA share in MSA in 1980
	BRA	BPC	All time-consistent Microregions 1980-2010	BA share in microregion in 1980
	CHN	CPC	All time-consistent cities, 1980-2010	BA share in city in 1980
BA share IV2	IND	IPC	All time-consistent districts, 1991-2011	BA share in city in 1991
	USA	USPC	All time consistent MSAs 1980-2010	MSA-specific weighted average of 2010 national BA shares by age groups, where the weights are predicted 2010 shares in total population of each age group (based on 1980 data). Age groups are the same used as controls in individual regressions.
	BRA	BPC	All time-consistent Microregions 1980-2010	Region-specific weighted average of 2010 national BA shares by age groups, where the weights are predicted 2010 shares in total population of each age group (based on 1980 data). Age groups are the same used as controls in individual regressions.
	CHN	CPC	All cities, 1948	Number of educational institutions in 1948.
	IND	IPC	All time-consistent districts, 1991-2011	District-specific weighted average of 2011 national BA shares by age groups, where the weights are predicted 2010 shares in total population of each age group (based on 1991 data). Age groups are the same used as controls in individual regressions.

Table D.6: *Area-level variables (Chapters 2 and 3)*

Variable	Samples	Description and comments
<i>Main variables of interest</i>		
FUNDEF shock (muni)	SC, TRES, BPC 2000	Change in the municipal-level fundamental education budget induced by FUNDEF, expressed as a fraction of the resources contributed the fund by local governments (equation 3.1.)
FUNDEF shock (region)		Weighted sum of the FUNDEF shock from the municipalities belonging to the microregion, using the share of each municipality in the region's school-age population as weights (equation 3.2.)
Migrant population	BPC 2000, 2010.	Total population of adult migrants.
Population	BPC 1980, 1991, 2000 and 2010.	Total population calculated over all observations (including population of all ages, not only adults).
Working-age pop.	BPC 2000, 2010.	Total population aged 15 through 64.
Average log rent residual	BPC 1991, 2010.	Average of the log rent residual at the region level, for households reporting positive monthly rent payments.
Average log wage residual	BPC 1980, 1991, 2000 and 2010.	Average of the log of the wage residual at the region level, for adult individuals reporting positive wage.
Employment	BPC 1980, 1991, 2000 and 2010.	Total employed adult population.
Employment rate	BPC 2000, 2010.	Employed individuals as a share of the working age population.
Participation rate	BPC 2000, 2010.	Individuals that participate in the labor force as a share of the working-age population.
Formality rate	BPC 2000, 2010.	Share of formally employed in participant population.
Informality rate	BPC 1980, 1991, 2000 and 2010.	Share of informally employed in participant population.
Unemployment rate	BPC 1980, 1991, 2000 and 2010.	Share of unemployed in participant population.
Non-participant population	BPC 1980, 1991, 2000 and 2010.	Total adult population that is not in the labor force.
Wage gap	BPC 1980, 1991, 2000 and 2010.	Average log wage for males minus average log wage for females at the microregion level.
Employment gap	BPC 1980, 1991, 2000 and 2010.	Ratio between share of employed in adult males and share of employed in adult females.
<i>Other variables</i>		
Log of population density	BPC 1980, 1991, 2000 and 2010; IPEA1.	Log of the ratio Population / Area.

Continued on next page

Table D.6: *(continued)*

Variable	Samples	Description and comments
Average winter temperature	IPEA2	Average winter temperature (June-August) in celsius. Microregion-level variable is an area-weighted average of municipal-level measures.
Share of high-school educated	BPC 1980, 1991, 2000 and 2010.	Share of high-school educated in adult population.
Formally-employed share in adults	BPC 1980, 1991, 2000 and 2010.	Share of formally employed in adult population.
Informally-employed share in adults	BPC 1980, 1991, 2000 and 2010.	Share of informally employed in adult population.
Unemployment rate	BPC 1980, 1991, 2000 and 2010.	Share of unemployed in the labor force.
Age group share (seven age groups)	BPC 1980, 1991, 2000 and 2010.	Share of each age group in region's population age groups are defined as: 1) 0-14; 2) 15-24; 3) 25-34; 4) 35-44; 5) 45-54; 6) 55-64; 7) 65 or older. Group 7 is the omitted group in all regressions.
Urbanization rate	BPC 1980.	Calculated from municipality aggregates published by the IBGE. In (IBGE2 source in subsection ??).
	BPC 1980, 1991, 2000 and 2010.	Share of total population living in locations classified as urban by the census (includes the urban districts of each municipality, as well as settlements that satisfy other geographic conditions (contiguity, infrastructure, and the availability of services.)
Major industry share in employment	BPC 1980, 1991, 2000 and 2010.	Shared of major industry in regional employment.
Microregion	BPC 1980, 1991, 2000, and 2010; IBGE2.	Time-consistent boundary of microregion. The definitions are constructed in two steps, following a procedure similar to that described in Kovak (2013). First, I construct time-consistent municipality boundaries (known in the literature as minimum-comparable areas - MCAs) by joining municipalities with common ancestors for the period 1980-2010, based on the official IBGE municipality family tree (see source IBGE2). IPEA provides a similar definition for the period 1872-2007 (Reis <i>et al.</i> 2007) but in this source MCAs are more aggregated than needed for accurate comparisons in recent decades. Second, I generate time-consistent microregions by aggregating MCAs that share common ancestors also for the period 1980-2010.

Continued on next page

Table D.6: *(continued)*

Variable	Samples	Description and comments
Microregion (robustness)	BPC 1980, 1991, 2000 and 2010.	Time-consistent microregion from Dix-Carneiro and Kovak (2017), which in turn takes as an input the original MCAs definitions provided by IPEA (Reis <i>et al.</i> 2007).
Arranjos populacio-nais	BPC 1980, 1991, 2000 and 2010.	Time-consistent Arranjos Populacionais (AP). Takes the original definition of AP (IBGE 2016) and joins arranjos that share a common MCA for the 1980-2010 period (using the same procedure as in the case of microregions).
Average log rent	BPC 1991, BPC 2010, IBGE3.	Monthly rent paid. The average of the log rent (i.e. the geometric average) is calculated over all renter households in the region.
Average log wage	BPC 1980, 1991, 2000 and 2010.	Average of the log wage (i.e. the geometric average) over employed adults with positive wage.
Non-employed share	BPC 1980, 1991, 2000 and 2010.	Share of non-employed (non-participant or unemployed) in adult population.
Labor force	BPC 1980, 1991, 2000 and 2010.	Adult population that is either formally employed, informally employed or unemployed.
Area (in square km)	IPEA1.	Geographic area in square kilometers, calculated aggregating the areas of the municipalities in each microregion.
Industry share in employment	BPC 1980, 1991, 2000 and 2010.	Shared of industry in regional employment (used to compute the Bartik shocks).