



Contents lists available at ScienceDirect

Journal of Urban Economics

journal homepage: www.elsevier.com/locate/jue

What is different about urbanization in rich and poor countries? Cities in Brazil, China, India and the United States[☆]

Juan Pablo Chauvin^{a,*}, Edward Glaeser^{a,b}, Yueran Ma^a, Kristina Tobio^a

^aHarvard University, United States

^bHarvard University and NBER, United States

ARTICLE INFO

Article history:

Received 5 February 2016

Revised 3 May 2016

Available online xxx

JEL classification:

O15

O18

R12

R23

Keywords:

Urbanization

Developing countries

Spatial equilibrium

Agglomeration economies

Human capital externalities

ABSTRACT

Are the well-known facts about urbanization in the United States also true for the developing world? We compare 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 huge, 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.

© 2016 Elsevier Inc. All rights reserved.

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 hypoth-

esis, which in its standard version assumes free mobility across metropolitan areas. Does that assumption make sense in a country like China, which historically imposed legal barriers to mobility such as the Hukou system (Au and Henderson, 2006b)?

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 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

[☆] We thank William Strange and an anonymous referee for excellent comments and suggestions. We acknowledge support from the Taubman Center for State and Local Government. Chauvin acknowledges support from the Center for International Development at Harvard.

* Corresponding author.

E-mail address: chauvin@g.harvard.edu, jpchauvin@gmail.com (J.P. Chauvin).

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 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 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 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, 2009](#)), 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 .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% 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.

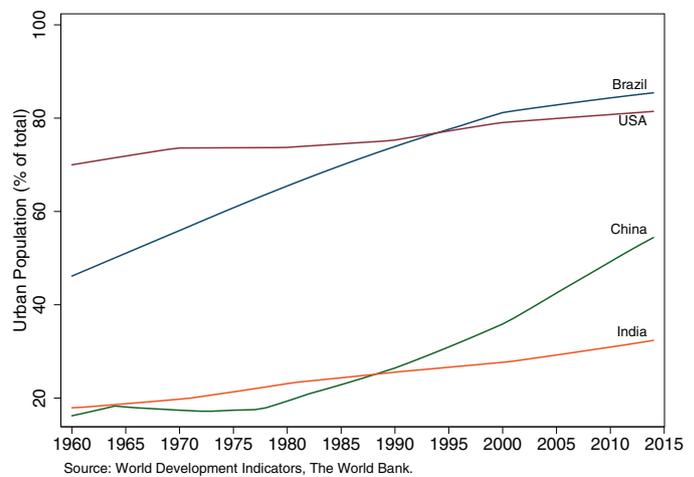


Fig. 1. Share of total population living in urban areas, 1960–2014.

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.

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. Fig. 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 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 \$7500 (in 2012 dollars) in the 1920s, when the U.S. became 50% urban, Brazilian per capita GDP only reached that level in 2011, when it was 80% 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% in 1960 to 31% 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% 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).

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 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.

2.2. The distribution of populations across area sizes

Table 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% of the population in such areas in 2000 and 41% in 2010. Collectively the other four population groupings contain only 26% of the U.S. population in 2010.

Brazil also has a large share of its population (32% 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% 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

Table 1
Share of people living in urban areas of different size.

	Areas of 100 K – 250 K (percent)	Areas of 250 K – 500 K (percent)	Areas of 500 K – 1 M (percent)	Areas of 1 M – 1.5 M (percent)	Areas of 1.5 M+ (percent)	Population in areas 100 K+ (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.

Table 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)						
	1990	Income per capita		1991	Income per capita	
Urban Areas	37,195	2000	2010	2000	2010	2010
Rural areas	26,816	44,071	45,124	3899	5830	7543
Difference	10,379	31,342	34,835	1141	1846	2731
		12,729	10,289	2758	3984	4812
China						
	1990	1999	2010	India		
				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)						
	1990	Income per capita		Earnings per capita		
Urban Areas	1451	1999	2010	2005	2011	
Rural areas	660	3064	6179	3123	5027	
Difference	792	1099	4265	1129	2265	
		1965	1914	1994	2762	

Sources: See data appendix.

increased by 18%, 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 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 high-

est share of its income going to the top ten percent (42%) 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% of national income goes to the top ten percent of the income distribution. In China and the U.S., 4.7% 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% 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% 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% higher than rural incomes, which is significant but not extreme.

By contrast, urban Indians in our sample earn 122% more than rural Indians. Urban Brazilians earn 176% 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.

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. Giesen and Sudekum, 2011; Rose, 2006; 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 $\frac{1}{N}$ times the population of the largest city. In large 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.

3.1. Zipf's Law across four countries

Fig. 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.

Table 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.

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 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.

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

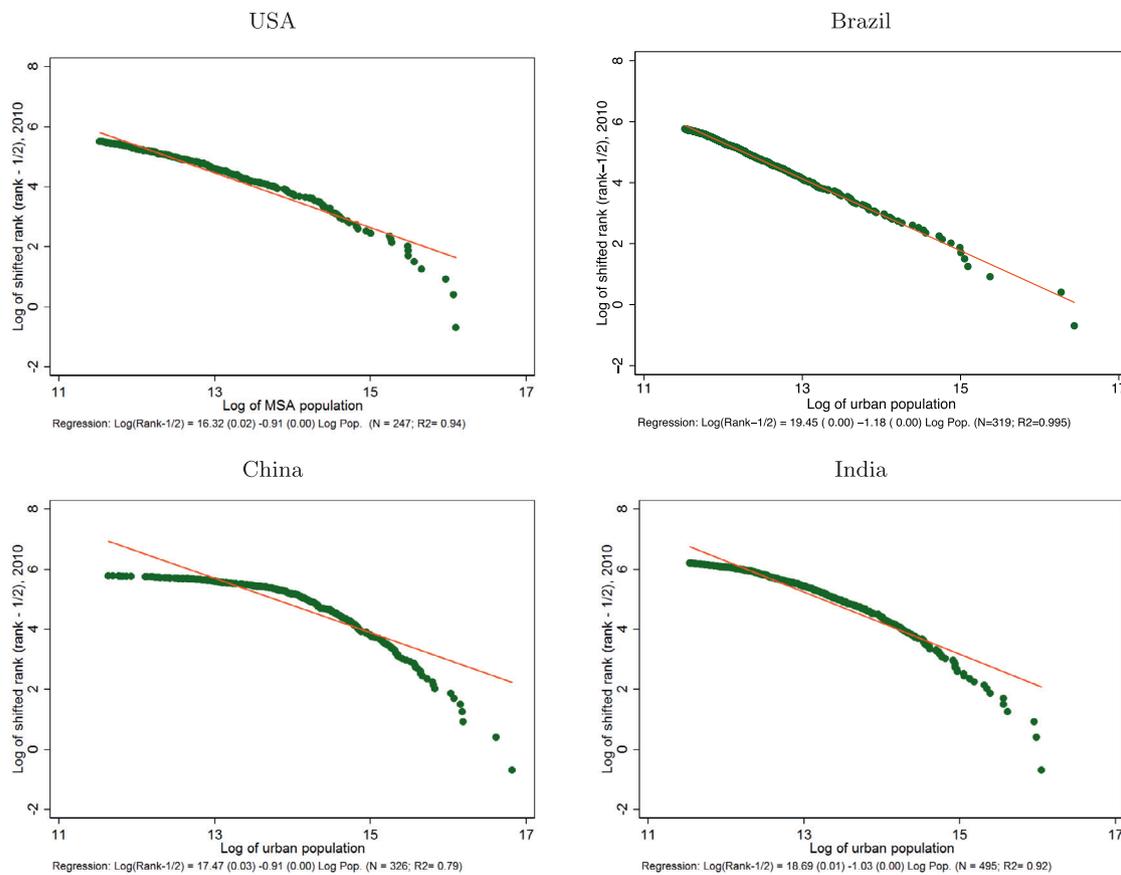


Fig. 2. Zipf's Law. Urban populations and urban population ranks, 2010. *Note:* Regression specifications and standard errors based on [Gabaix and Ibragimov \(2011\)](#). Samples restricted to areas with urban population of 100,000 or larger. *Sources:* See data appendix.

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.

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.

3.2. Gibrat's Law across four countries

Table 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

Table 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** (0.010) 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.

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 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.

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.

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 homogeneous 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 (2015) 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.¹

¹ Albouy (2015) provides a comprehensive discussion of the connection between deductibility and spatial equilibrium.

Table 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 = 29 M 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 = 1484 R2 = 0.304	1.922*** (0.172) N = 56 M R2 = 0.396	1.122 *** (0.073) N = 24.5 K R2 = 0.521
Average log wage residual	1.612*** (0.159) N = 29 M 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 = 1484 R2 = 0.304	2.887*** (0.256) N = 56 M R2 = 0.403	1.097*** (0.122) N = 24.8 K 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.

The log-separable formulation suggests the relationship:

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

The linear-separable formulation suggests:

$$\text{Price} = \text{After Tax Wage} + \text{Amenities} \quad (2)$$

If log price is regressed on log wages, then the first formulation implies that the coefficient will be $\frac{1}{\alpha} (1 + \frac{\text{Cov}(\text{Log}(\text{Wage}), \text{Log}(\text{Amenities}))}{\text{Var}(\text{Log}(\text{Wage}))})$. 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 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 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.

Fig. 3 shows the core relationship visually at the area level using housing rents, and Appendix Fig. 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, $\frac{1}{3}$ 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 mis-measured 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 5 and the second graph in Fig. 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% 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

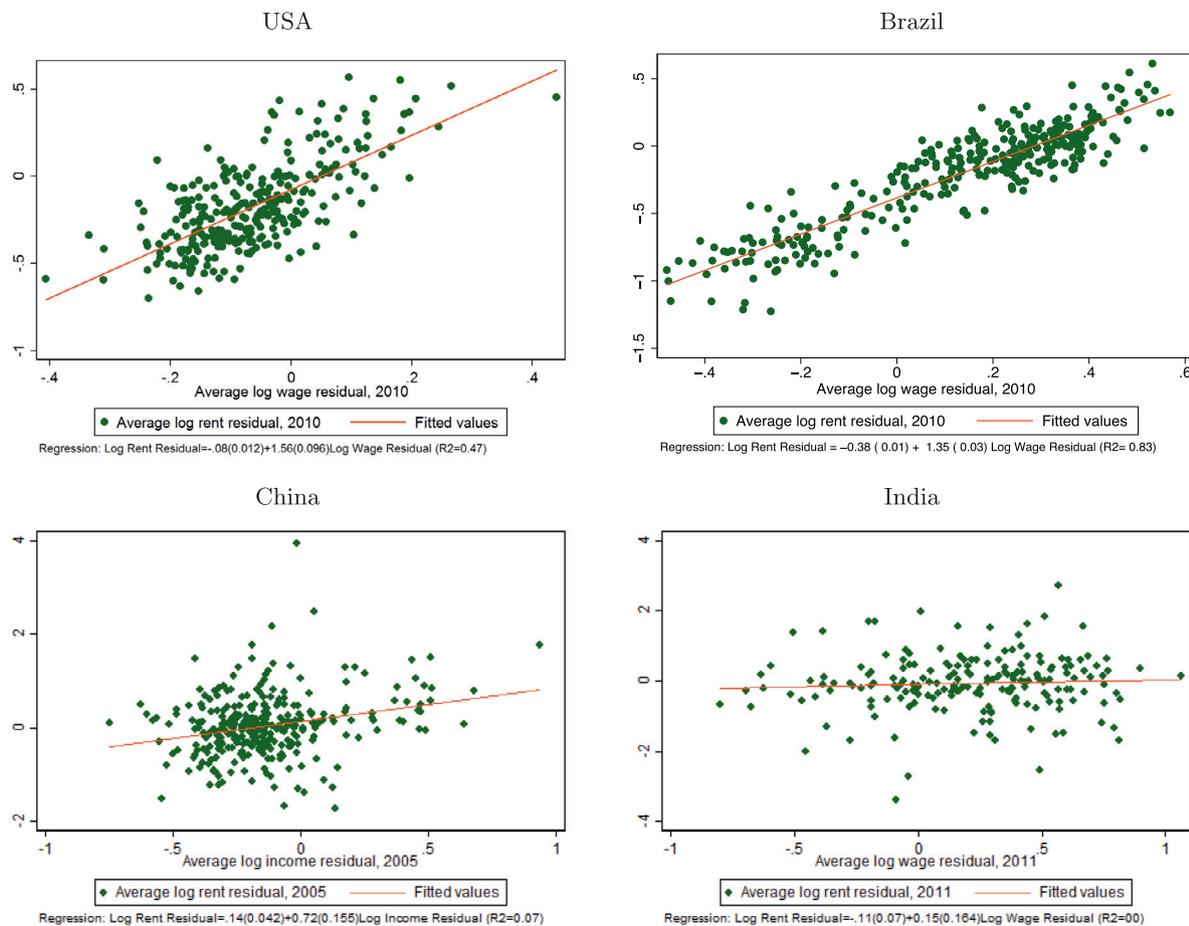


Fig. 3. Income and rents, 2010. *Note:* Samples restricted to areas with urban population of 100,000 or more. *Sources:* See data appendix.

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 5 and the third graph in Fig. 3 provide the results for China using housing rents as dependent variable, and the sixth column of Table 5 and second graph in Appendix Fig. 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 Fig. 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% of Chinese and 13.4% 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 Fig. 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 5. Possibly,

Table 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	8497	24,125	44,765	2157	2157	2157
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.

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.

4.2. Real wages and amenities

Eqs. (1) and (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, 2009),

France (Cavailhes et al., 2009) and Germany (Rehdanz and Madison, 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° Fahrenheit (21.11° Celsius). Consequently, the value can increase if the area is either particularly hot or particularly cold. The choice of 70° represents a middle ground within the 65° to 75° 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 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 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 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.

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.

Fig. 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. Gerdttham 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.

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

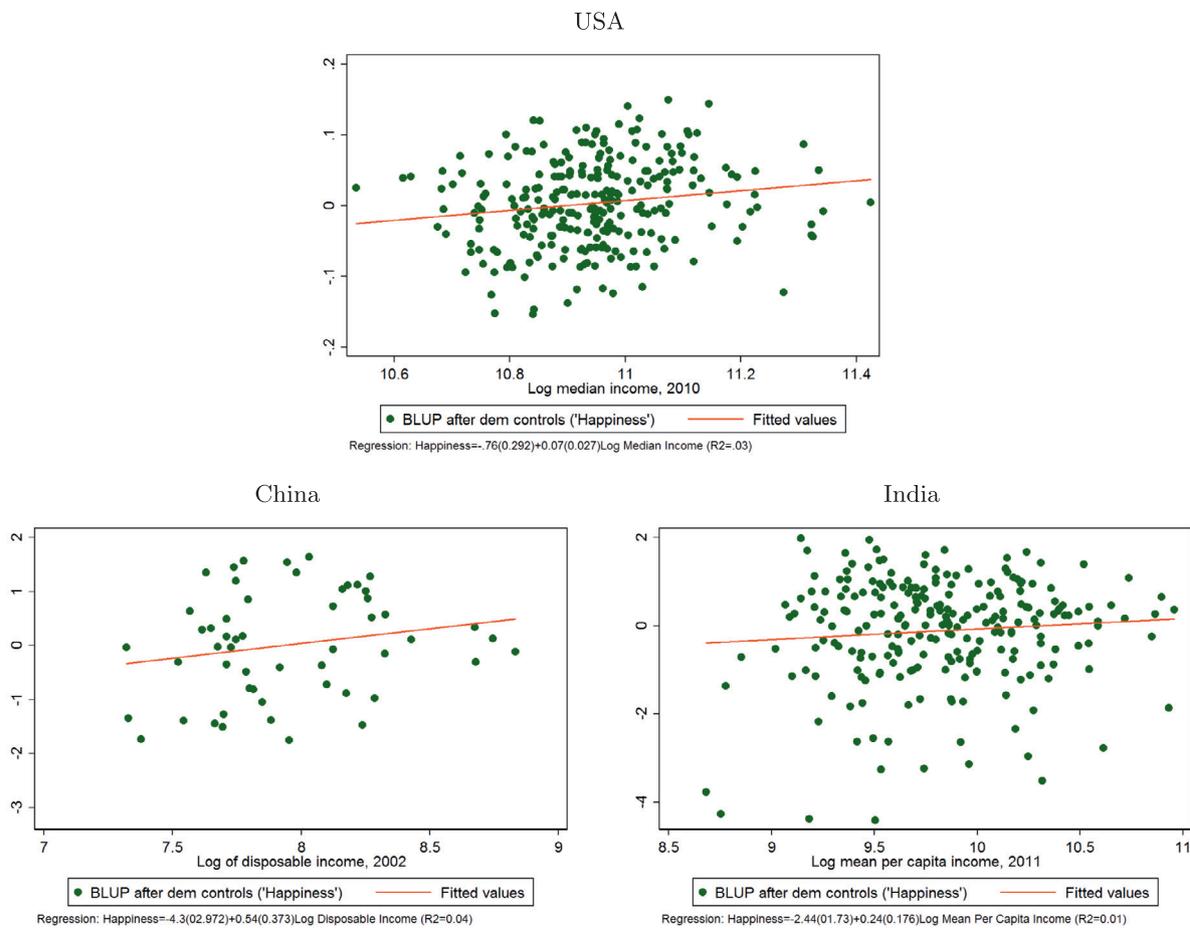


Fig. 4. Happiness and income levels. *Note:* Samples restricted to areas with urban population of 100,000 or more. *Sources:* See data appendix.

Table 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	1993	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.

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 7 shows, according to the 2000 Census, 21% 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 metropoli-

tan 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% of Americans had moved counties, states or countries during the previous five years. Only 7.1% 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 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 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 (Fieiss and Verner, 2003) but they are less mobile than Americans. Brazil's mobility rate has also declined over time. In 2000, 9.1% of the population had made a major or minor move over the previous five years. In 2010, 7.1% had made a major or minor move. Major moves are particularly rare. Only 2.4% 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% of the population had made a major or minor move over the previous five years. In 2010, 12.8% of the population had moved. Shen (2013) also documents this increase in mobility. Somewhat remarkably, China is now a more geographically mobile county 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 het-

erogeneity 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.

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.

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.²

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 (Combes et al., 2008; D'Costa and Overman, 2014). Cities may also attract more productive firms (Combes et al., 2012).

While there has been copious work on these issues for decades (Combes et al., 2010; Glaeser and Maré, 2001), 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.

² Theory 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-serviced neighborhoods in economically successful localities.

Table 8
Income and agglomeration, 2010.

	USA (MSAs) Log wage	Brazil (Microregions) Log wage	China (Cities) Log wage	India (Districts) 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	28.5 M	2172 K	147 K	9778
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	28.5 M	2172 K	143 K	7627
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	28.5 M	1998 K	112 K	5245
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.

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 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 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 pop-

ulation 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 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 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 coeffi-

cients. 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 (Bank, 2014). Pan and Zhang (2002) use firm production data and show that as city size doubles, firm productivity increases by 3.6%. 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% 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 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 Figs. 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 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, Au and Henderson (2006a) 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 Fig. 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,

Table 9
Real income and agglomeration, 2010.

	USA (MSAs) Log real wage	Brazil (Microregions) Log real wage	China (Cities) Log real wage	India (Districts) 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	28.5 M	2172 K	147 K	2102
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	28.5 M	2172 K	143 K	1649
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	28.5 M	1998 K	112 K	1141
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.

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 (De la Roca and Puga, 2014; Glaeser and Maré, 2001) and rising agglomeration effects over time (Brühlhart and Mathys, 2008).

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

Table 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)	34 M	27 M	2172 K	21,712 K	147 K	147 K	12 K	12 K
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	27 M	27 M	2172 K	2172 K	147 K		11 K	11 K
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)	17 M	16 M	2172 K	2172 K	147 K		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.

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 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 Fig. A.5 shows the relationship between area average log earnings residuals and area-level education.

The first two columns of Table 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).

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 rate 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% 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% 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% 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 $\frac{\text{Cov}(\text{Log}(\text{Wage}), \text{Unobservable Human Capital})}{\text{Var}(\text{Observed Human Capital})}$. 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.

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; Shapiro, 2006) 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 Germany (Sudekum, 2010), Spain (Ramos et al., 2010), and Europe more broadly (Badinger and Tondl, 2002; Rodríguez-Pose and Vilalta-Buffi, 2005). Using a sample of over 1500 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 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 11 uses total rather than urban populations.³

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% increase in population between 1980 and 2010. This effect is statistically and economically meaningful and it is shown in Fig. 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.

³ In 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 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.

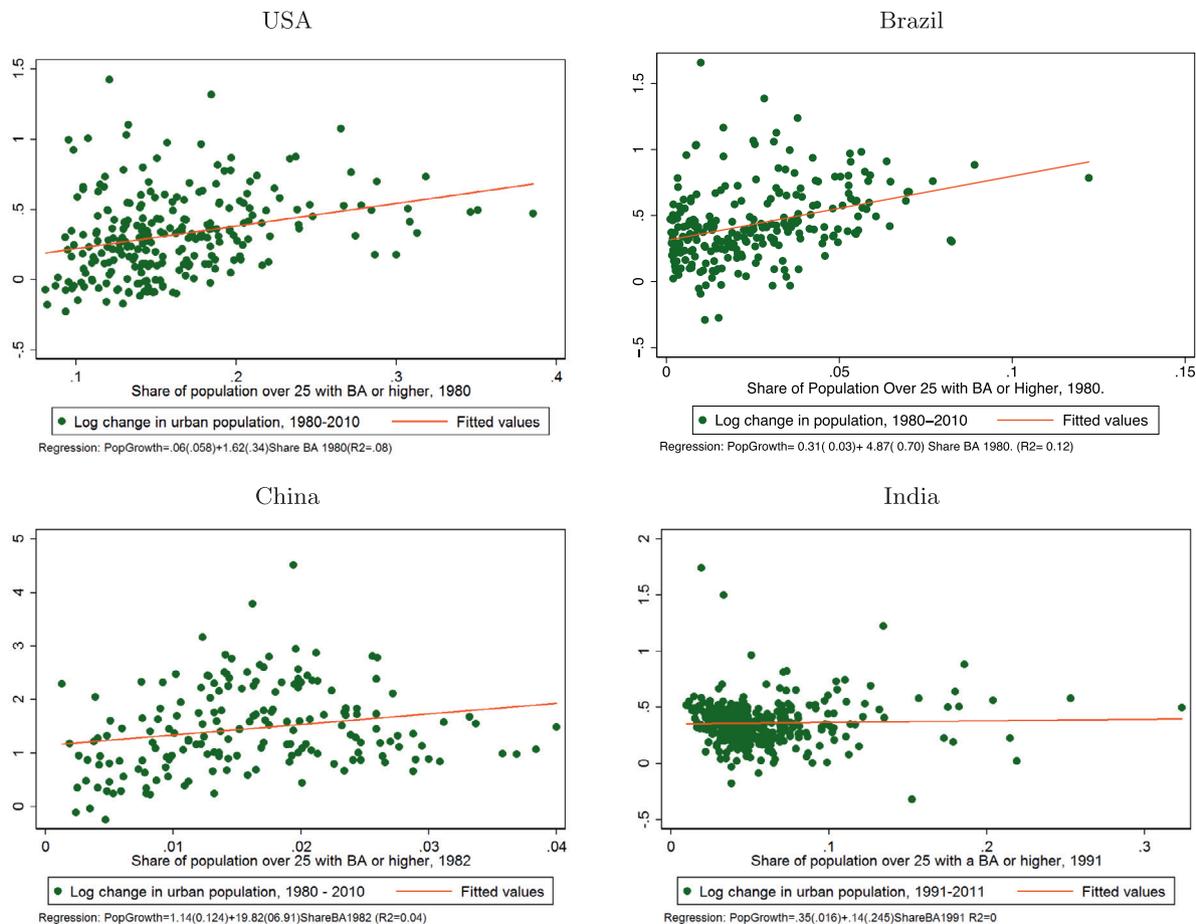


Fig. 5. University graduates share and population growth 1980–2010. Note: Samples restricted to areas with total population of 100,000 or more in 1980. Sources: See data appendix.

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% 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 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.

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% over thirty years. A one standard deviation increase in a Chinese area's education is predicted to increase population growth by around 52%. 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.

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.

Table A1
Happiness Regressions.

	USA (2010 MSAs) "Happiness"	China (2002 Cities) "Happiness"	India (2011 Districts) "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 A2
Local prices and agglomeration, 2010.

	USA (MSAs) Log rent	Brazil (Microregions) Log rent	China (Cities) Log rent	India (Districts) Log rent	USA (MSAs) Log price	China (Cities) 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	24.4 M	818 K	6.7 K	3281	44 M	25 K
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	24.4 M	818 K	6.4 K	2595	44 M	24 K
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	24.4 M	744 K	4.4 K	1792	44 M	19 K
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.

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% 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 homogeneous human capital, a spatial equilibrium may eventually appear in countries like India, where it still has not emerged.

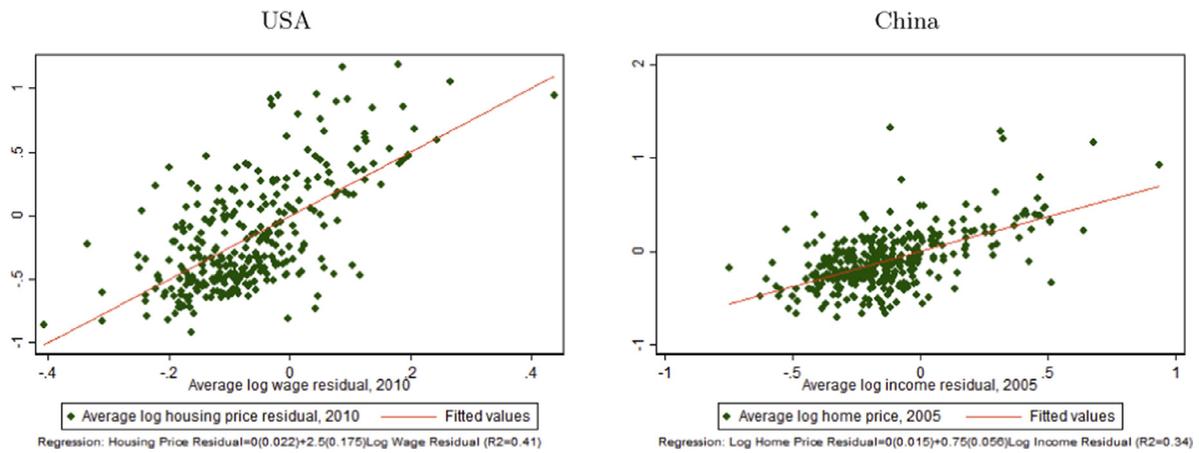


Fig. A1. Income and housing prices, 2010. *Note:* Samples restricted to areas with urban population of 100,000 or more. *Sources:* See data appendix.

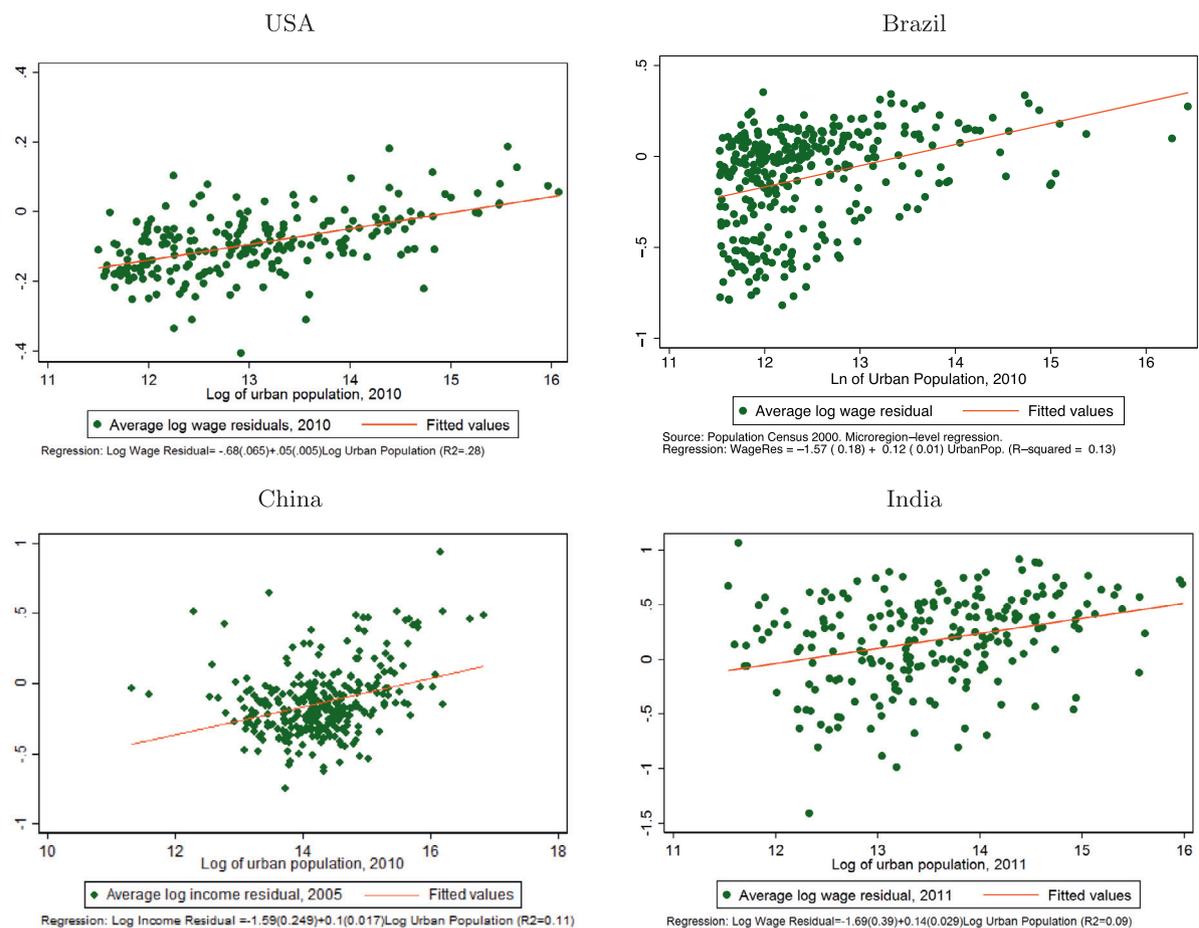


Fig. A2. Urban population and income residuals, 2010. *Note:* Samples restricted to areas with urban population of 100,000 or more. *Sources:* See data appendix.

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.

Appendix. Tables and Figures Appendix

1. Tables A.1 and A.2.
2. Fig. A.1 to A.5.

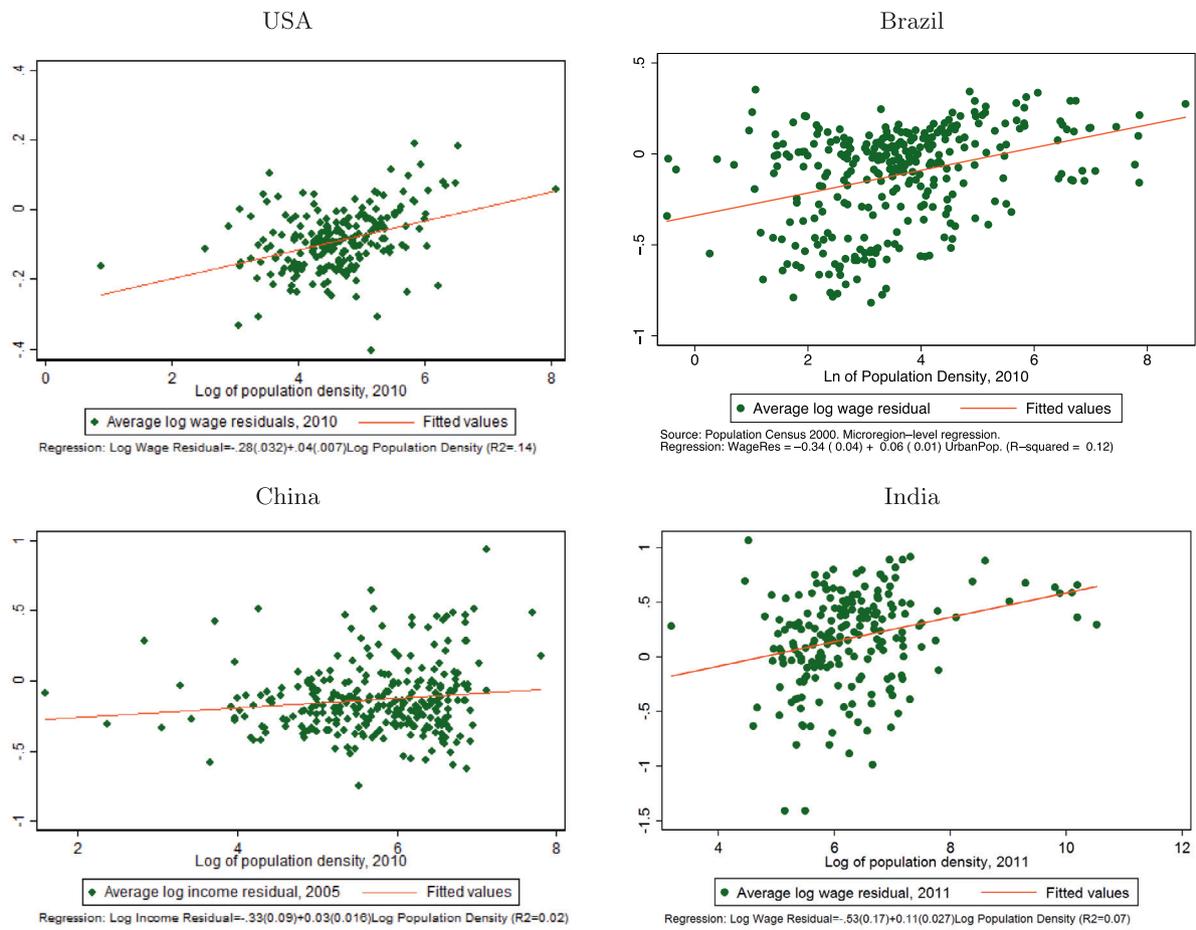


Fig. A3. Population density and income residuals, 2010 *Note:* Samples restricted to areas with urban population of 100,000 or more. *Sources:* See data appendix.

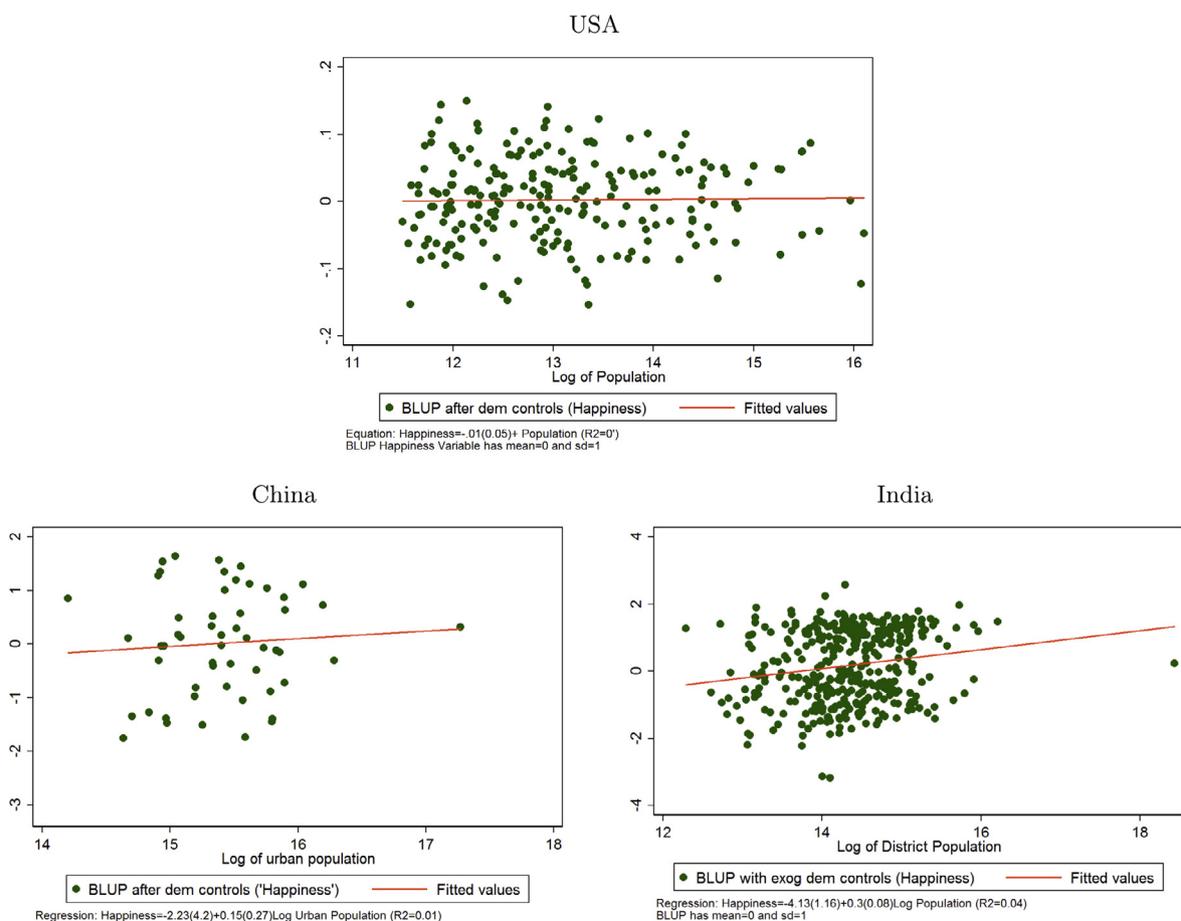


Fig. A4. Happiness and population size *Note:* Samples restricted to areas with urban population of 100,000 or more. *Sources:* See data appendix.

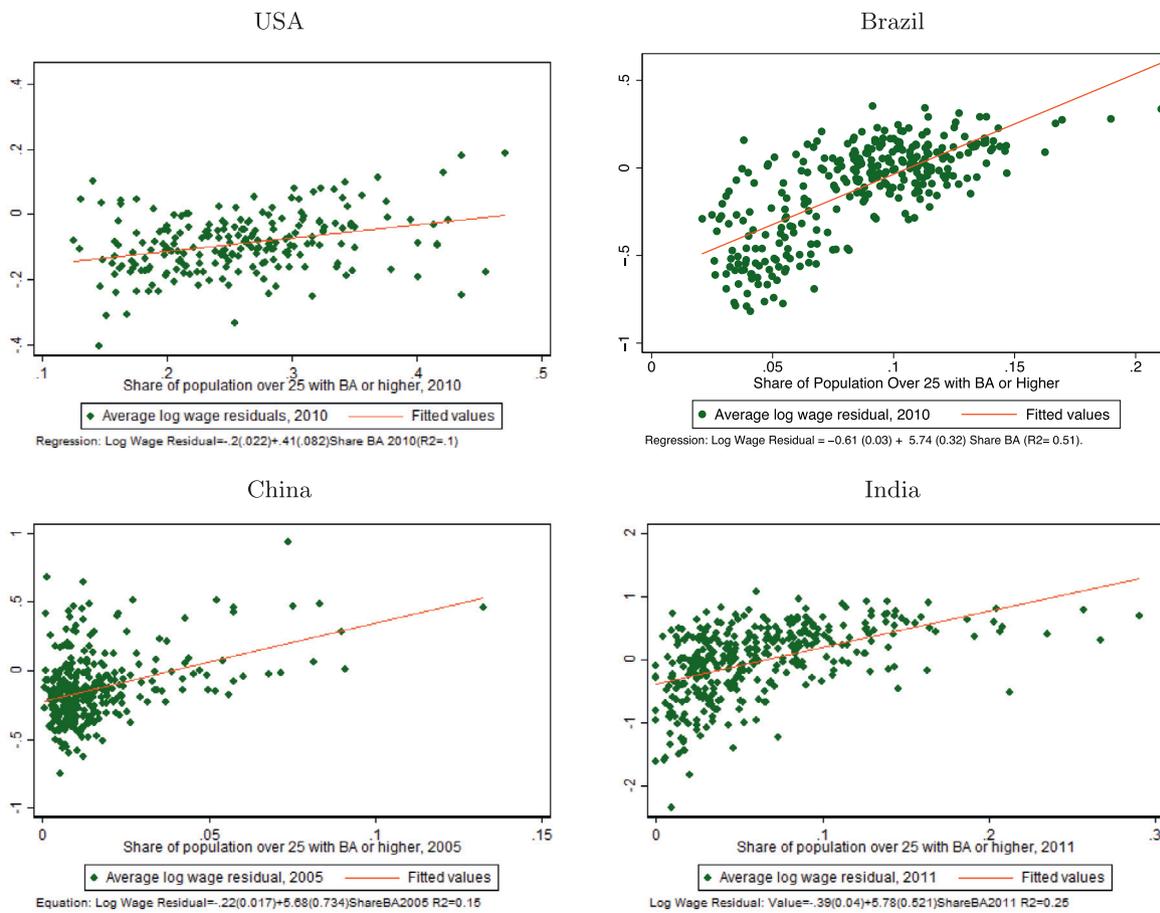


Fig. A5. University graduates share and wage residuals, 2010. Note: Samples restricted to areas with total population of 100,000 or more. Sources: See data appendix.

Appendix. Data Appendix

Databases Used			
Acronym	Database	Years	Source
International			
WDI	World Development Indicators, The World Bank	1960–2014	http://databank.worldbank.org
U.S.A.			
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	http://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 et al. (2014a)
Brazil			
BPC	Population Census micro-data sample	2010, 2000, 1991, 1980	www.flch.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.
China			
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.ciidbnu.org/chip
CCSY	China City Statistical Yearbook	2005, 2002	
PSY	Provincial Statistical Yearbook		
CMA	China Meteorological Administration		
India			
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

Country-level Variables Description				
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 reais 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)

(Continued from previous page)

Country-level Variables Description				
Variable	Country	Sources	Samples	Comments
Migrants in the last 5 years (% of population)	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	
	USA	ACS, USPC	1990, 2000, 2010	National aggregates of the homogenized variable "migrate5" from IPUMS, which refers to the person's place of residence 5 years ago.
	BRA	BPC-IPUMS	National aggregates 1991, 2000, 2010	
	CHN	CPC	2000, 2010	
IND	INSTAT	1993, 2001, 2011	This variable refers to migration in the past four years, the only data available.	
Area-level Variables Description				
Variable	Country	Sources	Samples	Comments
Areas	USA	USPC	2010	For cross-section regressions, tables, and graphs, we use 2010 Microregions as defined by the IBGE.
	BRA	BPC	2010	
	CHN	CPC	2010	
Time-consistent areas	IND	IPC2	2010	For growth regressions we use time-consistent MSAs. Counties' current information is aggregated using 1980s' MSA identifiers for the 1980–2010 regressions.
	USA	USPC	2010	
	BRA	BPC, ETDB	All municipalities' 1980, boundaries for the 1991, 2000 and 2010 censuses.	
Population / Urban population	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.
	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.
Population Density	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.
	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 5000; b) at least 75% of the male working population was non-agricultural; c) density of population of at least 400 sq. Km. (i.e. 1000 per sq. Mile.)
Average log wage/income per capita	USA	USPC	All MSAs	Average population by square mile.
	BRA	BPC	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.
Average log wage residual	IND	IHDS, INSTAT	All districts identified in the survey	Average population by square kilometer.
	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.
Average log wage residual	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.

(continued on next page)

(Continued from previous page)

Area-level Variables Description				
Variable	Country	Sources	Samples	Comments
Average log rent residual	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.
	IND	IHDS	All districts identified in the microdata, 2011	Weighted average of the district level of the urban individual log wage residuals
	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.
Average log housing price residual	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.
	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.
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	
Life satisfaction	IND	INSTAT	All districts, 2011	
	USA	BRFSS	367 MSAs covered by BRFSS	We use the BLUP variable from Glaeser, Gottlieb, and Ziv (Glaeser et al., 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.
	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.

(continued on next page)

(Continued from previous page)

Area-level Variables Description				
Variable	Country	Sources	Samples	Comments
Urban Population and Density IV2	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.
	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.
	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.
BA share IV1	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.

Individual-level Variables Description				
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.
	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.
Housing value	IND	IHDS	All renter households with rent data.	Annualized monthly housing rent.
	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)

(Continued from previous page)

Individual-level Variables Description				
Variable	Country	Sources	Samples	Comments
Log rent residual	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.
	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.
	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.
Log housing price residual	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.
	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.
Education group controls	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.
	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.
Age (demographic) group 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.
	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.
Dwelling characteristics controls	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.
	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.).

(continued on next page)

(Continued from previous page)

Individual-level Variables Description				
Variable	Country	Sources	Samples	Comments
	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 (grass/thatch, mud/unburnt bricks, plastic, wood, burned bricks, GI sheets or other metal, stone, cement/concrete, other); predominant roof type (grass/thatch/mud/wood, tile, slate, plastic, GI metal/asbestos, cement, brick, stone concrete, other); and predominant floor type (mud, wood/bamboo, brick, stone, cement, tiles/mosaic, others).

References

- Acemoglu, D., Angrist, J., 2001. How large are human-capital externalities? evidence from compulsory schooling laws. *NBER Macroeconomics Annual* 2000 15 (15), 9–74.
- Albouy, D., 2015. What are cities worth? Land rents, local productivity, and the total value of amenities.
- Alonso, W., 1964. *Location and Land Use*. Harvard University Press, Cambridge, Mass, Cambridge, MA.
- Anderson, G., Ge, Y., 2005. The size distribution of chinese cities. *Regional Science and Urban Economics* 35 (6), 756–776.
- Arriaga, E., 1968. Components of city growth in selected latin american countries. *The Milbank Memorial Fund Quarterly* 46 (2), 237–252.
- Au, C.C., Henderson, J.V., 2006a. Are chinese cities too small? *Review of Economic Studies* 73, 549–576.
- Au, C.C., Henderson, J.V., 2006b. How migration restrictions limit agglomeration and productivity in china. *Journal of Development Economics* 80, 350–388.
- Azzoni, C. R., Menezes-Filho, N., Silveira-Neto, R., 2000. Geography and income convergence among brazilian states. (May), 30.
- Azzoni, C.R., 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., 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.
- Bank, W., 2014. *Urban china: Toward efficient, inclusive, and sustainable urbanization*. Technical report.
- Bell, M., Muhidin, S., 2009. *Cross-National Comparisons of Internal Migration*.
- Bosker, M., Brakman, S., Garretsen, H., 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.
- Brühlhart, M., Mathys, N.A., 2008. Sectoral agglomeration economies in a panel of european regions. *Regional Science and Urban Economics* 38 (4), 348–362.
- Buettner, T., Ebertz, A., 2009. Quality of life in the regions: results from german counties. *Annals of Regional Science* 43, 89–112.
- Cavailles, J., Joly, D., Cardot, H., Hilal, M., Wavresky, P., 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.
- 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., 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.
- Chomitz, K.M., da Mata, D., de Carvalho, A.Y., Magalhães, J.C., 2005. Spatial dynamics of labor markets in brazil. *World Bank Policy Research Working Paper*, 3752.
- Ciccone, A., 2002. Agglomeration effects in europe. *European Economic Review* 46 (2), 213–227.
- Ciccone, A., Hall, R.E., 1996. Productivity and the density of economic activity. *The American Economic Review* 86 (1), 54–70.
- Combes, P.P., Démurger, S., Li, S., 2013. *Urbanisation and migration externalities in china*. Centre for Economic Policy Research Discussion Paper, 9352.
- Combes, P.P., Duranton, G., 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., Roux, S., 2012. The productivity advantages of large cities: distinguishing agglomeration from firm selection. *Econometrica* 80 (6), 2543–2594.
- Combes, P.P., Duranton, G., Gobillon, L., Roux, S., 2010. Estimating agglomeration economies with history, geology, and worker effects. In: *The economics of agglomeration*, I, pp. 15–65.
- Combes, P.P., Gobillon, L., 2015. The empirics of agglomeration economies. In: Duranton, G., Henderson, J., Strange, W.C. (Eds.), *Handbook of Regional and Urban Economics*, 5. North Holland, pp. 247–348.
- Corbi, R.B., Menezes-Filho, N.A., 2006. Os determinantes empíricos da felicidade no brasil. *Revista de Economia Política* 26 (4), 518–536.
- Cuñado, J., Pérez de Gracia, F., 2012. Does education affect happiness? evidence for spain. *Social Indicators Research* 108 (1), 185–196.
- D'Costa, S., Overman, H.G., 2014. The urban wage growth premium: sorting or learning? *Regional Science and Urban Economics* 48, 168–179.
- Dalmazzo, A., De Blasio, G., 2007. Production and consumption externalities of human capital: an empirical study for italy. *Journal of Population Economics* 20 (2), 359–382.
- Diamond, R., 2015. The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000.
- Durand, J.D., 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*. In: *The Milbank Memorial Fund Quarterly*, 43, pp. 166–196.
- Falck, O., Fritsch, M., Heblich, S., 2011. The phantom of the opera: cultural amenities, human capital, and regional economic growth. *Labour Economics* 18 (6), 755–766.
- Feler, L., Henderson, J.V., 2011. Exclusionary policies in urban development: under-servicing migrant households in brazilian cities. *Journal of Urban Economics* 69 (3), 253–272.
- Fiess, N. M., Verner, D., 2003. *Migration and human capital in brazil during the 1990s*. Technical Report 3093.
- Fleisher, B., Li, H., Zhao, M.Q., 2010. Human capital, economic growth, and regional inequality in china. *Journal of Development Economics* 92 (2), 215–231.
- Fontes, G.G., Simoes, R.F., De Oliveira, A.M., Hermeto, C., 2009. Urban attributes and wage disparities in brazil: a multilevel hierarchical model. *Regional Studies* 44 (5), 596–607.
- Foster, N., Stehrer, R., 2009. Sectoral productivity, density and agglomeration in the wider europe. *Spatial Economic Analysis* 4 (4), 427–446.
- Freguglia, R.S., 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., Gabriel, S.A., 2012. Labor migration, human capital agglomeration and regional development in china. *Regional Science and Urban Economics* 42 (3), 473–484.
- Gabaix, X., 1999. Zipf's law for cities: an explanation. *The Quarterly Journal of Economics* (August) 114, 739–767.

- Gabaix, X., 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.
- Gennaioli, N., La Porta, R., Lopez De Silanes, F., Shleifer, A., 2014. Growth in regions. *Journal of Economic Growth* 19, 259–309.
- Gerdtham, U.G., 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., Sudekum, J., 2011. Zipf's law for cities in the regions and the country. *Journal of Economic Geography* 11 (4), 667–686.
- Glaeser, E., Gottlieb, J.D., 2006. Urban resurgence and the consumer city. *Urban Studies* 43 (8), 1275–1299.
- Glaeser, E.L., 1999. Learning in cities. *Journal of Urban Economics* 46, 254–277.
- Glaeser, E.L., 2008. *Cities, Agglomeration and Spatial Equilibrium*. Oxford University Press, New York.
- Glaeser, E.L., Gottlieb, J.D., 2009. The wealth of cities: agglomeration economies and spatial equilibrium in the United States. *Journal of Economic Literature* 47 (4), 983–1028.
- Glaeser, E. L., Gottlieb, J. D., Ziv, O., 2014a. *Unhappy Cities*.
- Glaeser, E.L., Maré, D.C., 2001. Cities and skills. *Journal of Labor Economics* 19 (2), 316–342.
- Glaeser, E.L., Ponzetto, G.A.M., Tobio, K., 2014b. Cities, skills and regional change. *Regional Studies* 48 (1), 7–43.
- Glaeser, E.L., Scheinkman, J.A., Shleifer, A., 1995. Economic growth in a cross-section of cities. *Journal of Monetary Economics* 36 (1), 117–143.
- Glaeser, E.L., Saiz, A., 2004. The rise of the skilled city. *Brookings-Wharton Papers on Urban Affairs* 2004 (1), 47–105.
- Groot, S.P.T., De Groot, H.L.F., Smit, M.J., 2014. Regional wage differences in the Netherlands: micro evidence on agglomeration externalities. *Journal of Regional Science* 54 (3), 503–523.
- Hering, L., Poncet, S., 2010. Market access and individual wages: evidence from China. *The Review of Economics and Statistics* 92 (1), 145–159.
- Heurmann, D., 2011. Human capital externalities in western Germany. *Spatial Economic Analysis* 6 (2), 139–165.
- Hiller, N., 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., Lee, S., 2010. Cities as six-by-six-mile squares zipf's law?. In: Glaeser, E.L. (Ed.) *Agglomeration Economics*. No. February. The University of Chicago Press, pp. 105–131.
- IBGE, 2002. *Censo demográfico 2000: Documentação das microdadas da amostra*.
- IPEA, 2012. 2012: *Desenvolvimento inclusivo sustentável?*
- Jedwab, R., Vollrath, D., 2016. The Urban Mortality Transition and the Rise of Poor Mega Cities..
- Ke, S., 2010. Agglomeration, productivity, and spatial spillovers across Chinese cities. *Annals of Regional Science* 45 (1), 157–179.
- Ke, S., 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) 19 (2), 310–332.
- Knight, J., Gunatilaka, R., 2010. Great expectations? the subjective well-being of rural-urban migrants in China. *World Development* 38 (1), 113–124.
- Kovak, B.K., 2013. Regional effects of trade reform: what is the correct measure of liberalization? *American Economic Review* 103 (5), 1960–1976.
- Lall, S.V., Shalizi, Z., Deichmann, U., 2004. Agglomeration economies and productivity in Indian industry. *Journal of Development Economics* 73 (2), 643–673.
- Lin, H.L., Li, H.Y., Yang, C.H., 2011. Agglomeration and productivity: firm-level evidence from China's textile industry. *China Economic Review* 22 (3), 313–329.
- Liu, Y., 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.
- Long, F., Guo, M., Zheng, S., 2009. Estimating the willingness-to-pay for urban housing in Chinese cities. *Tsinghua Science & Technology* 14 (3), 360–366.
- McLanahan, S., Adams, J., 1987. Parenthood and psychological well-being. *Annual Review of Sociology* 13 (1987), 237–257.
- Melo, P.C., Graham, D.J., Noland, R.B., 2009. A meta-analysis of estimates of urban agglomeration economies. *Regional Science and Urban Economics* 39 (3), 332–342.
- 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.
- Morten, M., 2013. *Temporary Migration and Endogenous Risk Sharing in Village India*.
- Mueller, V., 2005. *Valuating climate amenities in Brazil using a hedonic pricing framework*. In: *Doctoral dissertation, University of Maryland at College Park*.
- Munshi, K., Rosenzweig, M., 2009. Why is mobility in India so low? social insurance, inequality, and growth.
- Muth, R., 1969. *Cities and Housing*. University of Chicago Press, Chicago and London.
- OECD, 2012. *Redefining "Urban" - A New Way to Measure Metropolitan Areas*. OECD Publishing, Paris..
- OECD, 2015. *OECD Urban Policy Reviews: China 2015*. OECD Publishing, Paris.
- Pan, Z., Zhang, F., 2002. Urban productivity in China. *Urban Studies* 39 (12), 2267–2281.
- Pittau, M.G., Zelli, R., Gelman, A., 2010. Economic disparities and life satisfaction in European regions. *Social Indicators Research* 96 (2), 339–361.
- Ramos, R., Suriñach, J., 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., Maddison, D., 2009. The amenity value of climate to households in Germany. *Oxford Economic Papers* 61 (1), 150–167.
- Reis, E. J., Pimental, M., Alvarenga, A. I., 2007. *Áreas mínimas comparáveis para os períodos intercensitários de 1872 a 2000*.
- 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., 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.
- De la Roca, J., Puga, D., 2014. *Learning by working in big cities*. CEMFI Working Paper.
- Rodríguez-Pose, A., Tselios, V., 2012. Individual earnings and educational externalities in the European Union. *Regional Studies* 46 (1), 39–57.
- Rodríguez-Pose, A., Vilalta-Bufi, 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., 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: Mieszkowski, P., Straszheim, M. (Eds.), *Current Issues in Urban Economics*. Johns Hopkins University Press, Baltimore and London, pp. 74–104.
- Schaffar, A., Dimou, M., 2012. Rank-size city dynamics in China and India, 1981–2004. *Regional Studies* 46 (6), 707–721.
- 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, 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., Nardinelli, C., 2002. Human capital and the rise of American cities, 1900–1990. *Journal of Urban Economics* 32, 59–96.
- Soo, K.T., 2005. Zipf's law for cities: a cross-country investigation. *Regional Science and Urban Economics* 35 (3), 239–263.
- Soo, K.T., 2014. Zipf, Gibrat and geography: evidence from China, India and Brazil. *Papers in Regional Science* 93 (1), 159–181.
- Sudekum, J., 2010. Human capital externalities and growth of high- and low-skilled jobs. *Jahrbuch für Nationalökonomie und Statistik* 230 (1), 92–114.
- Tabuchi, T., Yoshida, A., 2000. Separating agglomeration economies in consumption and production. *Journal of Urban Economics* 48, 70–84.
- 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.