

Urban productivity in the developing world

Edward L. Glaeser* and Wentao Xiong**

Abstract: Africa is urbanizing rapidly, and this creates both opportunities and challenges. Labour productivity appears to be much higher in developing-world cities than in rural areas, and historically urbanization is strongly correlated with economic growth. Education seems to be a strong complement to urbanization, and entrepreneurial human capital correlates strongly with urban success. Immigrants provide a natural source of entrepreneurship, both in the US and in Africa, which suggests that making African cities more livable can generate economic benefits by attracting talent. Reducing the negative externalities of urban life requires a combination of infrastructure, incentives, and institutions. Appropriate institutions can mean independent public authorities, public–private partnerships, and non-profit entities, depending on the setting.

Keywords: urban productivity, developing-world cities, agglomeration, education, entrepreneurship, immigration

JEL classification: O18, R00, L26

I. Introduction

Can cities help turn poor countries into rich countries? During the 1990s, much of Africa seemed to be experiencing urbanization without growth (Fay and Opal, 2000; Gollin *et al.*, 2016), and many feared that this would produce a permanent continent of slums. Those fears echoed the 1960s literature on urbanization without industrialization in Latin America (Durand *et al.*, 1965; Arriaga, 1968). Brazil became 50 per cent urban in 1964 when its *per capita* income was only \$2,000 in 2016 dollars.

Yet Latin America has subsequently grown, and today Brazil's *per capita* GDP is over \$11,000. Urbanizing Africa may also be experiencing its own 'growth miracle' (Young, 2012). For cities to provide pathways to prosperity, they must increase productivity either through improved technology, the spread of human capital, or more physical capital. The simplest 'compositional' growth model suggests that urbanization increases growth by enabling rural–urban migrants to experience the higher wages that are prevalent in cities. Consequently, the contribution that urbanization makes to growth is a simple function of the productivity gap between rural and urban areas.

To assess this gap, section II discusses the evidence on productivity and urbanization in the developing world. There are two standard means of assessing local productivity:

*Harvard University and NBER, e-mail: eglaeser@harvard.edu

**Harvard University, e-mail: wxiong@fas.harvard.edu

Glaeser thanks the Taubman Center at Harvard Kennedy School for financial support. Helpful comments were provided by Paul Collier and an anonymous referee.
doi:10.1093/oxrep/grx034

© The Authors 2017. Published by Oxford University Press.

For permissions please e-mail: journals.permissions@oup.com

firm-level output data and individual earnings data. Both typically show large differences across space within countries. For example, in 2014, labour productivity in São Paulo was 89 per cent higher than the rest of Brazil.¹ As agglomeration size doubles, wages rise by approximately 5 per cent in the US and Brazil, but the link is much larger in India and China (Chauvin *et al.*, 2016).

Jones *et al.* (2016) also find a substantial urban wage premium in several African countries. They perform two tests to determine whether this premium reflects unobserved individual heterogeneity that follow Glaeser and Mare (2001) and Combes, Duranton and Gobillon (2008). They conclude the urban wage premium is real in Africa.

Skills are a potent predictor of area-level earnings in the US (Rauch, 1993; Moretti, 2004), but the correlation is far stronger in Brazil, China, and India than in the US. Skills are also a particularly strong predictor of area population growth in Brazil and China (Chauvin *et al.*, 2016). This fact suggests that if urbanization is to generate large growth effects, it must be accompanied by investment in human capital.

We use firm-level data from China to examine productivity more directly. The link between prefecture-level² density and labour productivity across manufacturing firms is also stronger than in the US, but it is somewhat weaker than the link between density and earnings in China. We also find that export-oriented industries are particularly likely to locate in dense cities and to form industrial clusters. One interpretation of this fact is that agglomeration economies are particularly important when developing-world industries attempt to produce high-quality goods for the global market.

In section III, we turn to the relative importance of local entrepreneurship, immigration, and foreign direct investment (FDI). One hypothesis is that countries with relatively low levels of human capital can only grow by attracting multi-national corporations, which have much better management practices (Bloom and Van Reenen, 2010). An alternative view champions the role of local entrepreneurs. A third view emphasizes the complementarity between the local entrepreneurship and foreign investment.

While it may be hard to imagine that local entrepreneurship will lead to new export industries in today's poorest countries, initially small operators have often played a large role in the growth of East Asian and Latin American economies. Soichiro Honda began his remarkable career as a car mechanic. US-based research shows that measures of local entrepreneurship, such as an abundance of small establishments, strongly predict subsequent employment growth (Glaeser *et al.*, 2015). Dirubhai Ambani began as a small Mumbai spice exporter in 1966 and built the massive Reliance conglomerate.

John Sutton's enterprise maps of various African countries³ document that many successful businesses were built by immigrants or their children. Immigrants provide a natural source of entrepreneurial capital that can complement investments in native education. The economic advantage of attracting outside talent also implies that improving the quality of life in urban Africa is a form of economic development strategy.

¹ <http://www.euromonitor.com/sao-paulo-city-review/report>

² In China, a prefecture is a sub-provincial administrative unit, typically consisting of 2–4 urban districts and 4–8 surrounding suburban districts and counties. As of 2004, there were 337 prefectures in China, but this paper restricts empirical analysis to 287 'prefecture-level cities', and excludes 50 (often remote) 'prefecture-level regions' with little manufacturing activity and quite limited prefecture-level data.

³ 'The Enterprise Map Project' by John Sutton and co-authors; detailed description on John Sutton's homepage at LSE (<http://personal.lse.ac.uk/sutton/>)

Section IV focuses on making African cities more livable. The higher productivity of particularly skilled cities suggests significant benefits from making cities of the developing world more habitable. Reducing the downsides of density, such as contagious disease, congestion, and crime, could make it possible for successful cities to expand and allow more people to enjoy enhanced productivity. Yet fighting the downsides of density requires a combination of infrastructure, incentives, and institutional reform.

Section V provides a discussion of the potential larger benefits of urbanization. There is a strong track record of cities leading to democratic uprisings that topple dictators. It is less clear that cities enable the establishment of effective democracies. The relation between cities and political improvement remains another important topic for future research.

Section VI concludes with four policy implications. First, we stress the importance of reducing artificial barriers to urban growth, such as excessive land-use controls. Second, we discuss the policy approaches to the downsides of urban life. Third, we reiterate the need to better educate developing-world cities. Fourth, we emphasize the need to explore entrepreneurship-related interventions.

II. Productivity differences across cities

Within both rich and developing countries, there are large differences in income over space. In 2015, *per capita* gross domestic product (GDP) was over three times higher in Shenzhen than in the rest of China. Workers in urban Uganda earned 0.48 log points more than workers in the countryside (Jones *et al.*, 2016). Bangalore's *per capita* income was more than 2.5 times the Indian average in 2015.⁴ Similarly, *per capita* GDP in London and Paris is much higher than in the rest of Britain or France.

While these gaps exist in rich and poor countries alike, it seems particularly important to understand the gaps in regional productivity in the poorer world. After all, the poor regions of India or Africa are far more destitute than the poor regions of the United Kingdom or France. If poor countries are to have more widespread prosperity, then we must better understand why some parts of those countries have managed to become richer. Moreover, we must understand whether the growth of productive places is limited by local disamenities or artificial limits on housing. Perhaps housing markets and amenities can be improved in ways that enable more people to benefit from the productivity that appears in cities such as Shenzhen and Bangalore.

In this section, we examine the evidence on the correlation between urban density and productivity in Brazil, China, and India, drawing upon Chauvin *et al.* (2016). We first discuss productivity differences across space and the link between productivity and agglomeration size. We then discuss the correlation between productivity, human capital, and growth.

We have two ways of assessing local productivity: earnings and total factor productivity. In a neoclassical model, earnings equal the marginal product of labour. Consequently, differences in earnings should capture differences in productivity across

⁴ <http://www.thehindu.com/news/cities/bangalore/bengaluru-urban-tops-state-in-per-capita-income-kalaburagi-last/article8376124.ece>

space. Since earnings can differ for reasons other than productivity, including labour market regulations, we also supplement earnings-based data with firm-level data on labour productivity in China.

(i) Heterogeneity in productivity across space

We begin by discussing the heterogeneity in productivity across space and its link with city size and city density levels. Perhaps the simplest evidence of productivity disparity is the comparison between urban and rural earnings. In the US, for example, urbanites earn approximately 30 per cent more than rural residents. [Glaeser and Mare \(2001\)](#) find that this gap does not disappear with controls for individual human capital attributes, including test scores. Presumably, higher housing costs and other urban disamenities offset higher urban wages. Yet even though higher urban wages are not a free lunch, they still provide an indication of higher productivity in cities.

The urban–rural wage gap also exists in Brazil, India, and China. [Chauvin *et al.* \(2016\)](#) find that urban earnings are 45 per cent higher than rural earnings in China, 122 per cent higher in India, and an astounding 176 per cent higher in Brazil. Indeed, in that paper the urban residents of Brazil were the highest paid group, but the rural residents of Brazil earned almost as little as the rural residents of India, the lowest paid group in the sample.

[Chauvin *et al.* \(2016\)](#) use the regional average residual from an equation in which the logarithm of wages is regressed on individual controls as the primary measure of local productivity. In the US, the bulk of metropolitan areas are contained in an 0.4 log point band, meaning the most productive areas are about 50 per cent more productive than the least productive areas. In Brazil, China, and India, metropolitan area average log wage residuals spread over a full log point range, meaning that the most productive areas are about 170 per cent more productive than the least productive areas.

[Jones *et al.* \(2016\)](#) estimate the urban wage premium in Tanzania, Uganda, and Nigeria. They generally find significant urban–rural differences, especially in the primate city of each country. They use two tests to assess the importance of sorting on unobserved characteristics. They show that most of the urban wage premium is offset by higher costs of living, which would presumably not be true if urban workers were just more able ([Glaeser and Mare, 2001](#)). They also find little sorting on observable characteristics, which might also mean that there is little sorting on unobservable characteristics.

This huge dispersion in earnings seems to be matched with a huge dispersion in labour productivity, which we define as the log of value added per worker. Using Chinese firm-level data, discussed in the [Data appendix](#) of this paper, [Figure 1](#) shows the strong correlation between labour productivity and log of earnings per worker. This relationship treats each industry–prefecture–year as an observation,⁵ taking the average of firm-level variables within each observation and including industry–year fixed effects. This correlation supports the view that wage-based heterogeneity is likely to reflect underlying heterogeneity in firm-level productivity.

[Figure 2](#) illustrates the spatial productivity gap across industries. Within an industry, from high to low, we rank prefectures by the average labour productivity of local firms,

⁵ While the raw Chinese firm-level data cover mining, manufacturing, and utility industries, we restrict empirical analysis to approximately 420 manufacturing industries.

and take the difference between the most productive prefecture (Prefecture 1) and the runner-up (Prefecture 2). While the average industry is present in about 90 prefectures, an enormous productivity gap already exists between the top two prefectures: on average, Prefecture 1 is 0.48 log point, or about 60 per cent, more productive than Prefecture 2. It is noteworthy that this spatial gap differs significantly across industries.

One reasonable question is how such large productivity gaps persist in equilibrium. Why don't workers flood into high-wage locales? One explanation is that we have not controlled adequately for unobserved worker human capital. Another is that migration is limited either by strong place-specific tastes or by explicit government policies, such as the household registration (*hukou*) system. A third possibility is that high housing costs and disamenities offset the higher wages in urban areas. If this third explanation is correct, then there is scope for government policies that support affordable housing and reduce the disamenities of urban life.


(ii) Agglomeration economies

We now turn to the question of whether productivity is higher in large metropolitan areas or denser prefectures. Economists have long argued that big cities enjoy

Figure 1: Productivity and earnings across prefectures and industries in China



Figure 2: Spatial productivity gaps across industries in China



agglomeration economies because urban scale makes it easier to ship goods or hire well-matched workers or exchange ideas. Typically, these agglomeration economies are measured either with higher wages (Glaeser and Mare, 2001) or firm productivity (Combes *et al.*, 2010), and urban scale is measured either with metropolitan-area population or with population-density levels.

There are three standard challenges for interpreting agglomeration coefficients: unobservable firm attributes, unobservable worker attributes, and unobserved spatial attributes. Unobservable firm attributes mean that more productive firms may disproportionately locate in cities, perhaps to take advantage of large markets for products or workers. Unobservable worker attributes mean that more productive workers may be particularly present in cities. Unobservable spatial attributes mean that a large city may form in an area to take advantage of some innate locational attribute such as access to a harbour or a coal mine. In China, agglomerations have formed in special economic zones (SEZ) like Shenzhen, which the government established to encourage export, and so the SEZ status may be the spatial attribute that is increasing both the size and productivity of the agglomeration.

Typically, economists address these concerns both by controlling assiduously for observed personal characteristics and by instrumenting for city size with attributes, like natural geographic features, that seem to predict density but wouldn't impact productivity directly. Notably, controls for individual attributes do little to dispel the other two problems. Geographic instruments do little to address the concerns about unobserved firm and worker productivity. Here we simply note the problems with interpretation and proceed with simple, standard correlations.

Agglomeration effects are typically estimated by regressing the logarithm of wages on the logarithm of area population or area density, together with individual controls. In the US, the coefficient on either population or density is approximately 0.05, meaning that as area size or density doubles, wages increase by 0.05 log points or about 5 per cent. This correlation remains essentially unchanged if historical population or density levels are used to instrument for current population or density.

In comparison, the coefficient on area population is also about 0.05 for Brazil, and again essentially unchanged if historical population is used as an instrument. The coefficient on density is about 0.025. The coefficients for India are about 0.075 for both population and density, so it seems that agglomeration economies may be particularly strong in the sub-continent (Chauvin *et al.*, 2016).

The relationship between prefecture population and earnings in China is larger still, but it is statistically insignificant from zero. The insignificance may reflect our smaller Chinese data sample, or perhaps the population of China's prefectures doesn't mean all that much. These are political jurisdictions that may not reflect the boundaries of actual metropolitan areas. Prefecture density, measured as population per square kilometre, has an extremely strong relationship with earnings. As log density doubles, earnings rise by 0.19 log points. This finding remains when we instrument for current population density with historic population levels (Chauvin *et al.*, 2016).

(iii) Agglomeration economies across industries

These aggregate results mask considerable heterogeneity across industries. We now examine cross-industry differences in agglomeration, meaning the extent to which firms

locate near one another, and urbanization, meaning the extent to which firms locate in high-density prefectures. Industries differ significantly in both agglomeration and urbanization, and in the extent to which their productivity levels are correlated with agglomeration and density.

Table 1 shows the evolution of agglomeration across approximately 420 manufacturing industries. We use the Ellison and Glaeser (1997) index of agglomeration, which attempts to correct for establishment size. If an industry is concentrated in just a few large establishments, then naturally it will be concentrated in at most a few locations. The index corrects for this tendency, and can be interpreted through a dart-throwing metaphor. If industries choose locations by throwing darts at a map of prefectures, then the Ellison–Glaeser index can be interpreted as the probability that any industry's dart becomes welded to the immediately preceding dart instead of hitting the map at random.

In 2000, the mean agglomeration index was 0.0256, which is quite low relative to the US. The US mean in the original study was approximately double this amount. The range of agglomeration is considerable: for instance in 2004, the index runs from 0.0067 (essentially no agglomeration) at the 10th percentile to 0.0804 at the 90th percentile, with a standard deviation of 0.0439. This high dispersion may well be the legacy of the era of central planning in which some industries were consciously distributed across China.

Between 2000 and 2004, the index rose steadily, reaching a mean of 0.0396 by 2004. During this period of significant industrial growth, some regions seem to have attracted particular industries far more than others. Between 2004 and 2007, the level of agglomeration was steady or declined somewhat, and ended with a mean of 0.0373. The standard deviation followed roughly the same pattern, first increasing and then falling slightly, ending at 0.0386 in 2007.

The 10th percentile rose only slightly to 0.0065, which is still a very modest level of agglomeration. The 90th percentile increased to 0.0819, which reflects a fairly agglomerated industry. Ellison and Glaeser (1997) refer to any industry with a value over 0.05 as 'very localized'. One way of interpreting these trends is that China in 2000 had significantly less industrial concentration than the US, but ended up closer to US levels.

There are a few industries that have negative values of the index, meaning that they are more dispersed than they would be if locations were determined purely by chance. Silk-dyeing (a very traditional industry) and silver smelting both fall into this category. At the other end, artificial fibre manufacturing and electronic music equipment manufacturing have index values over 0.1.

Agglomeration theorists have identified a wide number of reasons why firms like to cluster with other firms of their own industry, including the ability to share inputs and ideas, proximity to customers, and access to a larger labour pool. We next test whether, on average, firms that locate in clusters of their own industry have higher levels of productivity. The first regression in Table 2 uses 2004 firm-level data and shows the association between an industry-prefecture's mean labour productivity and the prefecture's share of total employment in the industry. In all regressions in Table 2, we include industry-specific fixed effects, and cluster standard errors at industry level.


The estimated coefficient is 0.601, implying that as a prefecture's share of total industry employment increases from 5 to 15 per cent, labour productivity increases by 0.06 log points or about 6 per cent. Figure 3 illustrates this relationship using bin scattered

Table 1: Ellison–Glaeser index of agglomeration across industries in China

Year	N	Mean	SD	p10	p90
2000	413	0.0256	0.0284	0.0026	0.0592
2001	412	0.0266	0.0262	0.0031	0.0626
2002	417	0.0309	0.0351	0.0037	0.0664
2003	418	0.0336	0.0350	0.0053	0.0729
2004	420	0.0396	0.0439	0.0067	0.0804
2005	423	0.0380	0.0399	0.0064	0.0807
2006	418	0.0390	0.0432	0.0069	0.0801
2007	418	0.0373	0.0386	0.0065	0.0819

Table 2: Advantage of agglomeration at industry level

	(1)	(2)	(3)
	(ln) Prefecture mean labour productivity	Prefecture % of total employment in industry	(ln) Prefecture mean labour productivity
Prefecture % of total employment in industry	0.601*** (0.079)		
(ln) Prefecture population density (/km ²)		0.005*** (0.000)	0.059*** (0.003)
R-squared	0.145	0.154	0.164
N	36,034	34,976	34,791

Figure 3: Labour productivity and employment concentration across prefectures in China

data points. The relationship seems non-linear: areas with more than 2 per cent of an industry's employment typically have quite good productivity outcomes. This effect is somewhat difficult to interpret, as the idea of a spatial equilibrium for firms precludes the possibility that firms will be more productive in one place than another unless there are offsetting disadvantages of the locale, such as higher real estate costs.

We now turn from localization to urbanization. The two concepts are different, but they are also linked. The second regression in [Table 2](#) shows the relationship between the logarithm of a prefecture's population density and its share of an industry's total

employment. Locations with denser population also tend to have more industrial employment. The correlation is statistically significant, but relatively small, reflecting the non-urban nature of much manufacturing. Figure 4 shows the relationship visually, highlighting high levels of concentration in the densest prefectures.

The third regression of the table shows the relationship between population density and mean labour productivity across prefectures. We estimate a coefficient of 0.059, implying that as density doubles, labour productivity increases by about 6 per cent. This estimate is slightly higher than wage-based agglomeration estimates for the US, but lower than the wage-based estimates for China (Chauvin *et al.*, 2016). Figure 5 shows the connection, again using binned data. To test whether this reflects the impact of SEZ status, we rerun the regression excluding special economic zones, and still find a large and significant relationship.

Why should agglomeration economies be stronger in India and China than in Brazil and the United States? One possibility is that these relationships are largely spurious. We know that firm-level productivity differences in India and China are particularly large (Hsieh and Klenow, 2009). Perhaps more productive firms just select more into cities and this has a larger effect in China and India because these firms enjoy a particularly

Figure 4: Employment concentration and population density across prefectures in China





Figure 5: Labour productivity and population density across prefectures in China



large advantage. Another possibility is that there is more selection of able workers into Chinese and Indian cities.

Yet it is also possible that these measured agglomeration economies are real. Large cities can be far more connected to the outside world. Consequently, western technology may enter through cities and cause firms there to be more productive. If this is the case, then it is particularly important to understand the limits to urban growth in India and China—for instance, China’s *hukuo* system that constrains rural–urban migration.

While this overall connection is interesting, it is somewhat more interesting to better understand the heterogeneity across industries in the returns to agglomeration and urbanization. Figure 6 shows heterogeneity in the correlation between labour productivity and density across four-digit industries in China.

In some industries, the correlation between labour productivity and density is extremely high, while in others, such as crude oil processing, the correlation is small or even negative. Figure 7 shows the strong positive correlation between density and productivity for machine-made paper and cardboard manufacturing. In the US, we would think of this industry as being largely suburban, taking advantage of cheap land. But in

Figure 6: Heterogeneity in the correlation between labour productivity and density





Figure 7: Labour productivity and prefecture population density for machine-made paper and cardboard



China, this type of manufacturing is still largely urban and seems to reap productivity advantages from its urban location.

We now focus on the inter-industry variation in agglomeration, urbanization, and the observed productivity benefits from locating in areas with high industrial concentration and population density. Based on 2004 data, [Table 3](#) shows a correlation matrix between different measures across approximately 420 different industries. The first row and column show results related to the Ellison–Glaeser agglomeration index, our preferred measure of agglomeration.

The second row and column present results related to the correlation between a prefecture’s share of total industry employment and population density, our preferred measure of urbanization. The correlation between agglomeration and urbanization is quite small and insignificant. Some industries form industrial clusters and others disproportionately locate in cities, but the demand for urban density does not seem correlated with the demand for proximity to other firms in the industry.

The third row and column provide results related to the correlation between labour productivity and the share of total industry employment in a prefecture. There is really no tendency for highly agglomerated industries to have a higher productivity return to locating in an industrial cluster. This low correlation reminds us that these productivity relationships are equilibrium outcomes. If an industry has an innately higher return to agglomerating, then presumably agglomeration continues until those returns are eroded. Consequently, we do not see a tendency of extra agglomeration in industries with higher observed returns to agglomeration.

The fourth row and column give results related to the correlation between population density and labour productivity. In this case, there is a modest positive correlation with urbanization. Hence industries with higher observed returns to urbanizing do seem to urbanize more.

The next two rows and columns turn to industrial characteristics that could potentially drive up returns to either agglomeration or urbanization. First, we consider the share of employees with college degrees. Highly educated industries tend to be more urbanized within the US, perhaps because of the informational advantages of locating in a city. Second, we consider the share of an industry’s output that is exported. Export industries have fewer gains from proximity to Chinese customers, which might reduce the benefits of urban location. However, they also benefit from access to ports, which might increase the tendency to locate in ports and SEZs.

As shown in the fifth row and column, more-educated industries are not more likely to agglomerate or urbanize. They do, however, show significant productivity benefits from locating in industrial clusters. One interpretation is that higher-educated industries are more knowledge-intensive and benefit more from knowledge sharing, which industry clusters serve to facilitate.

Export-orientation, conversely, is strongly positively correlated with agglomeration and urbanization and negatively correlated with industry-level education. The agglomeration effects are actually stronger when we exclude special economic zones from our analysis, but the urbanization results are significantly weaker. This suggests that the high level of urbanization in these industries reflects partially the tendency to locate in SEZs. One interpretation of the results on agglomeration is that these industries have to learn what external markets want and they typically learn this from other firms in the industry.

Table 3: Industry-level pairwise correlations related to agglomeration advantage

	EG agglomeration index	[corr.] prfc. In(pop. density) and % employment	[corr.] prfc. % employment and In(labour productivity)	[corr.] prfc. In(pop. density and In(labour productivity)	% employees with college degree	In(weight-to-value ratio)	% sales from export
EG agglomeration index	1.000						
[corr.] prfc. In(pop. density) and % employment	0.074 (0.129)	1.000					
[corr.] prfc. % employment and In(labour productivity)	-0.018 (0.721)	0.3119*** (0.000)	1.000				
In(labour productivity)	0.028 (0.574)	0.1253** (0.010)	0.1274*** (0.009)	1.000			
% employees with college degree	-0.014 (0.780)	0.002 (0.966)	0.2146*** (0.000)	0.049 (0.318)	1.000		
In(weight-to-value ratio)	-0.1514*** (0.008)	-0.1205** (0.034)	-0.0992* (0.081)	-0.011 (0.854)	-0.3093*** (0.000)	1.000	
% sales from export	0.2880*** (0.000)	0.1594*** (0.001)	-0.0982** (0.045)	0.049 (0.322)	-0.3474*** (0.000)	-0.3256*** (0.000)	1.000

Notes: Approximately 420 4-digit industries; except for pairwise correlations involving weight-to-value (approximately 310 industries). p-value in parentheses; ***, **, *, indicate statistical significance at 1%, 5%, and 10% levels, respectively.

Table 4: Industry-level factors explaining agglomeration advantage

	(1)	(2)	(3)	(4)	(5)	(6)
	EG agglomeration index			[corr.] prfc. ln(pop. density) and % employment		
% employees with college degree	-0.024 (0.113)			0.011 (0.339)		
(ln) Weight-to-value ratio		-0.007*** (0.002)			-0.014** (0.006)	
% sales from export			0.083*** (0.015)			0.133*** (0.035)
Constant	0.039*** (0.006)	0.031*** (0.005)	0.020*** (0.004)	0.174*** (0.015)	0.157*** (0.012)	0.144*** (0.013)
R-squared	0.000	0.023	0.083	0.000	0.015	0.025
N	420	311	417	420	311	417
	(7)	(8)	(9)	(10)	(11)	(12)
	[corr.] prfc. % employment and ln(labour productivity)			[corr.] prfc. ln(pop. density) and ln(labour productivity)		
% employees with college degree	0.921*** (0.200)			0.212 (0.235)		
(ln) Weight-to-value ratio		-0.010 (0.006)			-0.001 (0.005)	
% sales from export			-0.066* (0.035)			0.034 (0.036)
Constant	-0.009 (0.011)	0.020** (0.010)	0.050*** (0.010)	0.133*** (0.012)	0.140*** (0.010)	0.135*** (0.011)
R-squared	0.046	0.010	0.010	0.002	0.000	0.002
N	420	311	417	420	311	417

Table 4 shows the results on industrial heterogeneity in regression form. We also include (log) weight-to-value ratio, a proxy of transportation cost per unit of output, but this measure is only available for a subset of industries and come from US Commodity Flow Survey (see Duranton *et al.*, 2014). The first three regressions show correlations with the agglomeration index. While education level does not predict the agglomeration index, weight-to-value and export orientation do.

The next three columns show the correlation between a prefecture's population density and its share of employment in an industry, our measure of urbanization. Similar to earlier results on agglomeration, urbanization is strongly negatively correlated with weight-to-value, and positively correlated with export orientation. This suggests that successful cities and industrial clusters may be particularly important to supporting industrial exports.

Regressions (7)–(9) show results related to the correlation between concentration and productivity. In this case, education is a strong positive correlate of gains to concentration. Our preferred interpretation is that industrial clusters facilitate knowledge sharing and thus benefit knowledge-intensive industries. Export orientation is negatively associated with productivity benefits from concentration, yet this coefficient becomes insignificant if we exclude firms in SEZs.

These results hopefully emphasize that urbanization and localization benefits differ across industries. The strongest results highlight the connection between export orientation and both agglomeration and urbanization.

(iv) Human capital externalities

One literature has focused on the connection between earnings and metropolitan area population or density. A second literature (Rauch, 1993; Moretti, 2004) emphasizes the links between earnings and the skill level in the metropolitan area. This literature aims at uncovering the size of human-capital externalities—the benefits from having skilled neighbours. In both cases, the basic premise is that people benefit from having human capital around. Agglomeration economies emphasize the quantity of workers, and human-capital externalities emphasize the quality of workers.

If anything, the omitted-worker characteristics problem is more severe with human-capital externalities than with agglomeration economies. Since areas are defined by their skill level, it seems reasonable to assume that there is sorting on unobservable human capital as well as on observable human capital. Consequently, it makes sense to take all estimated human-capital externality coefficients with a grain of salt. Nevertheless, it may still be useful to compare US coefficients with those for the developing world. The following discussion draws heavily on Chauvin *et al.* (2016).

Within the US, the human-capital effect on wages has risen over time. Using 2010 data, Chauvin *et al.* (2016) find a coefficient of approximately one. This implies that a 10 percentage point increase in the share of adults with college degrees is associated with an 0.1 log point increase in earnings, holding individual years of schooling constant. This coefficient remains unchanged when we instrument using historic schooling data, but that instrumentation does little to eliminate concerns about sorting on unobserved skills.

The coefficient is about 2 for India, 4.7 for Brazil, and 5.2 for China. The standard deviation of the human-capital measure is smaller for India, Brazil, and China, but even taking this into account, the effect of area-level skills on earnings appears to be larger than in the US.

Why should area-level skills be so important in the developing world? One natural hypothesis is that skills enable the spread of knowledge and the knowledge gap is particularly large in many developing-world cities. Bloom *et al.* (2016) find corroborating evidence for this view: they show a strong link between good management practices and proximity to universities throughout the world. Apparently, being close to centres of knowledge production increases the tendency to know how to run a business.

Skills predict not only earnings at a point in time, but also the growth in area-level earnings and population (Glaeser *et al.*, 1995). In the US, a 10 percentage point increase in the share of the adult population with college degrees in 1980 is associated with an increase in population growth of 0.21 log points and an increase in income of 0.09 log points between 1980 and 2010. Skilled areas have been growing far more rapidly in the US than unskilled areas.

In Brazil, the link between skills and area growth is even stronger. A 5 percentage point increase in the share of adults with college degrees in 1980 is associated with an

0.25 log point increase in population and an 0.6 log point increase in income over the next 30 years. These effects are almost too large to be plausible. In China, a 5 percentage point increase in the share of adults with tertiary education in 1980 is associated with more than a 1 log point increase in population, although surprisingly the impact on earnings is negative and statistically indistinct from zero. The Indian coefficient is positive but small for population growth, and we lack income growth data for India over the entire time period.

The Brazil results seem to suggest that education is closely tied to local success within that economy. The Chinese income-growth result is puzzling, but one possibility is that the extremely fast growth of high human-capital cities has attracted an abundance of less skilled workers who have pulled income levels downward. The Indian results may reflect the far more spatially static nature of the Indian economy.

Certainly, these results suggest that cities and skills are strong complements in the developing world. Perhaps the most basic implication of this finding is that investing in skills reaps considerable returns, which is corroborated by scores of other studies. A secondary implication is that public policy should ensure that skilled cities face no artificial barriers to their growth, such as artificial limits on housing supply or under-provision of urban services.

We now turn to a well-documented fact in the US: measures of local entrepreneurship predict local success. Unfortunately, the development literature has no equivalent results. Consequently, it remains unclear how important local entrepreneurship is to the success of developing-world cities.

III. Entrepreneurship and economic development

Both anecdotes and data support the importance of local entrepreneurship for US cities. In 1971, a billboard rose on the highway leaving Seattle proclaiming ‘Will the last person leaving Seattle – Turn out the lights’. Boeing, the region’s largest employer then, had been laying off workers, and just as no one could imagine a Detroit with a smaller General Motors, no one could imagine a Seattle with a smaller Boeing. In the 4 decades since then, new businesses have come to the city and reinvigorated its economy, including Amazon, Costco, Microsoft, and Starbucks. In some cases, like Microsoft, the entrepreneur had a long-term connection with Seattle. In other cases, like Amazon, the entrepreneur was attracted by attributes of the city, including its well-educated workforce.

In 1961, Benjamin Chinitz argued that New York City was more resilient than Pittsburgh during the 1950s, because New York City had a culture of entrepreneurship that meant that its business leaders were good at adapting to industrial decline. In modern language, we might describe New York as having a healthy endowment of entrepreneurial capital because its dominant industry, garment production, had limited-scale economies and few barriers to entry. In contrast, Pittsburgh had US Steel, and the steel industry had large-scale economies, which meant that Pittsburgh trained company men instead of entrepreneurs.

Subsequent empirical research has tested Chinitz’s hypothesis with various proxies for local entrepreneurship, including average establishment size and the share of

employment in new establishments at some initial time period. These variables are strongly correlated with subsequent employment growth both across cities and across industrial groups within cities. The effect is enormously robust and not just a reflection of either broad American regional patterns (e.g. the decline of the rustbelt) or industrial patterns (e.g. the decline of manufacturing). Glaeser *et al.* (2015) follow Chinitz directly and use the presence of mines in 1900, which explains Pittsburgh's steel industry, as an instrument for large establishments. The basic correlation between small establishment size and subsequent employment growth remains strong with this IV strategy.

One puzzle is that these proxies for local entrepreneurship are strongly correlated with employment growth, but not income growth. This seems quite reasonable within cities, where presumably the supply of labour across industries should be quite elastic. It is somewhat more surprising that entrepreneurship does not predict income growth across cities, which may also be explained by a sufficiently elastic labour supply. Alternatively, the job-creating entrepreneurs might be quite good at keeping the costs of labour low.

Modern variants of the Chinitz hypothesis essentially view entrepreneurship as yet another form of human capital. According to this view, just as some cities are endowed by their history with more formal education, the industrial past has left some places, such as New York, with more entrepreneurial human capital than others, such as Pittsburgh. While some entrepreneurial human capital is mobile, some of it stays put and provides an enduring economic advantage to its locale.

The public role in generating entrepreneurship is less clear. It seems quite reasonable to believe that local regulations can stymie entrepreneurship, although there has been little research using US data documenting such a relationship. While local governments do occasionally try to increase entrepreneurship by supporting specific 'innovation clusters', we know little about whether such clusters are really effective or whether other local policies, such as entrepreneurship training programmes, will bear fruit.

(i) Entrepreneurship in Africa

Despite the remarkable enterprise maps by John Sutton and his co-authors, we have no comparable literature documenting the effects of entrepreneurship in Africa. This dearth of research is a major shortcoming because there is a significant debate about the relative role of local entrepreneurship in many developing countries. One side of the debate affirms that local entrepreneurship is as important in the developing world as it is in the developed world. The other side claims that the human capital gap between developed and developing countries is now so wide that developing-world cities will only be able to export manufactured goods and services with the help of FDI.

The literature on FDI is well developed, but it yields somewhat ambivalent answers about economic growth. For example, Borensztein *et al.* (1998) find that FDI positively impacts growth only when the host country has at least a threshold level of human capital. This finding suggests that FDI might not be a solution for countries with particularly low levels of human capital, but the question still remains as to whether local entrepreneurship can be effective in those countries either. One possible view is that neither FDI nor local entrepreneurship will engender growth if the level of human capital is sufficiently low.

The Sutton enterprise maps provide a remarkable overview of businesses across a range of African economies, including Ethiopia, Ghana, Tanzania, and Zambia. The export businesses tend to skew towards natural resources, such as copper, and agricultural products, such as coffee, flour, and salt. Many of these firms, such as Zambia's Unity Garments, began as trading firms and expanded into production. Forty-eight per cent of the Ethiopian firms profiled in [Sutton and Kellow \(2010\)](#) began with trading.

There are also numerous businesses that specialize in retail trade, consumer goods, and transportation for the African market. Ethiopia's Belayneh Kindie began as a transport company in 2006 and has since branched out into businesses as varied as metals production and hotel construction. Zambia's Zambeef began when two partners leased a butchery and abattoir ([Sutton and Langmead, 2013](#)). While there may be far too few African entrepreneurship success stories, the Sutton maps document that they do exist.

The Sutton maps also document the importance of immigrant entrepreneurs in Africa. In Ghana, two Lebanese brothers founded Irani Flour, an Armenian founded Takoradi Flour Mills, and two Greek brothers founded Panbros. In Zambia, Mohammed Iqbal Patel, a Zambian citizen of Indian descent, founded Trade Kings as a bakery. It now has 1,600 employees, manufactures detergents and steel, and allegedly operates the largest lollipop line in the southern hemisphere. Most spectacularly, a Portuguese immigrant, Fernando Duarte, co-founded Nando's, an international casual-dining restaurant chain, in South Africa in 1987.

Immigrant entrepreneurship is common in the US and Europe as well. Google's Sergey Brin was born in Russia, and Intel's Andy Grove came from Hungary. The Franco-Israeli businessman Patrick Drahi, who founded the Altice Group, was born in Morocco. [Kerr and Kerr \(2016\)](#) estimate that over 35 per cent of new firms in the US have at least one immigrant among their top three earners.

The very act of immigration itself reflects risk-taking, so perhaps it is not surprising that immigrants are also disproportionately drawn to the risks of entrepreneurship. In some cases, existing businesses may discriminate against immigrants, and then entrepreneurship becomes an alternative for their talents. Immigrants, such as Chinese diaspora, often maintain global links that can support start-ups, such as trading firms. The sharp difference in entrepreneurship rates across ethnicities suggests that some ethnicities, like India's Gujaratis, have developed entrepreneurial human capital over generations. Higher-human-capital immigrants would have a particular entrepreneurship advantage in lower-human-capital countries.

Along with education, immigration does seem to provide one way for sub-Saharan Africa to obtain more entrepreneurial talent. In principle, Africa provides enormous opportunities for talented go-getters with a good understanding of the needs of global markets. Yet Africa also faces headwinds in attracting such entrepreneurs. Governments often place significant barriers to new businesses, and negotiating local politics can be more challenging for outsiders.

Moreover, relocating to Africa may be seen as a far less pleasant prospect than moving to Paris or Silicon Valley. The difficulty of living in many African cities is surely a handicap in the global war for talent. The benefits of attracting entrepreneurial outsiders remind us that improving the quality of life in Africa's cities should be seen as an economic development strategy that may be as important as reducing unnecessary regulations.

Education improvement seems another critical investment, and skills can mean more than just formal schooling. [Chinitz \(1961\)](#) himself suggested that entrepreneurial human capital was learned at the breakfast table and at the workplace. Many developing-world cities seem to be already well endowed with entrepreneurial human capital. To walk through the Dharavi slum of Mumbai is to be surrounded by small, scrappy businesses. Many African cities also have an abundance of energetic, small-scale entrepreneurs.

In a sense, the lack is not entrepreneurial talent, but rather the ability to produce goods for global customers. The entrepreneurs of Lusaka are largely making products for Zambians and perhaps a few tourists. The city has spread knowledge about local market opportunities, but not the tastes of customers elsewhere and certainly not the knowledge of how to produce for global markets.

(ii) Policies towards entrepreneurship in Africa

Another perspective suggests that FDI, immigration, and education are complements rather than substitutes. The intrusion of foreign companies into a developing-world city brings knowledge, and potentially also opens markets for domestic entrepreneurs. [Desai *et al.* \(2006\)](#) find that a 10 per cent increase in foreign activity within a country is associated with a 2.2 per cent increase in domestic economic activity, which supports the view that there are spillovers from foreign investment for domestic businesses. [Greenstone, Hornbeck and Moretti's \(2010\)](#) work on million-dollar plants also shows such positive spillovers from outside investment within the United States.

Immigrant entrepreneurs will find a country more attractive if its work force is more skilled. Skilled natives will also find it easier to partner successfully with immigrant entrepreneurs. While investing in education can result in myriad economic benefits, it is less clear how to encourage native entrepreneurship.

The US literature does indicate that local entrepreneurship has been important for local economic growth. There is not yet any comparable literature for developing-world cities, and there is little hard evidence—in either the US or elsewhere—on how public policy can potentially encourage local entrepreneurship. There are at least three public policy strategies aimed at increasing local entrepreneurship: training, clusters, and deregulation.

Business schools have tried to train entrepreneurs for decades, yet there is little rigorous evidence that such training works. There are cheaper programmes that try to provide disadvantaged youths, such as ‘The Possible Project’ in Cambridge, Massachusetts, but they have not yet been evaluated with randomized control trials. It consequently remains an open question whether cities can actually teach entrepreneurship.

A second approach focuses on the generation of entrepreneurial clusters, which presumably allow entrepreneurs to learn from each other. Boston’s Innovation District is one such public initiative. Private initiatives, such as co-working spaces for small start-ups, also provide scope for sharing entrepreneurial knowledge. In a sense, markets filled with small, individual merchants in the developing world, either with or without explicit public support, represent yet a third form of entrepreneurial cluster. Again, we have little firm empirical evidence on whether the formation of such clusters materially increases the overall level of entrepreneurship within a city.

A third approach starts with the view that at least some entrepreneurs are deterred by various regulations. Many US cities, for example, forbid food trucks to provide meals on city streets, which appears to deter at least one form of urban entrepreneurship. In the developing world, really small-scale entrepreneurs typically ignore labour- and product-market regulations, so deregulation seems unlikely to increase the number of really tiny firms. However, it seems more likely that these regulations prevent the growth of such firms, especially when they reach the point of employing non-family members. Such regulations may explain the dominance of small firms in the firm-size distribution in the many developing countries. Small firms can't grow into larger firms because they would then have to follow the rules.

FDI does seem like a sure-fire way to generate local employment, but it is less clear that it generates long-term growth, especially for low-human-capital societies. Most of the successful transitions from poverty to prosperity involved a significant number of home-grown entrepreneurs, like Soichiro Honda.

In many cases, however, these entrepreneurs did benefit from imported entrepreneurial talent. For example, the early development of Shenzhen relied heavily on Chinese businessmen living in nearby Hong Kong. The Bangalore culture of entrepreneurship got some help from Patni Computer Systems, which began as part of Boston's information technology cluster and then employed future entrepreneurs, like Narayana Murthy.

Across US cities, differences in employment growth are clearly linked to differences in the supply of entrepreneurship across space. This same pattern may hold in the developing world as well, but we lack both basic facts and serious evaluations of entrepreneur-related public policy. This topic remains a pressing area for future research.

IV. Taming the demons of density

If urbanization can play an important role in abetting economic growth, then one question is whether public policies should do more to increase city size. There are many reasons to be wary of explicit spatial policies that encourage migration to one region or another. Most obviously, it is unclear whether encouraging urbanization would enhance welfare overall. On average, workers in cities earn more, but they also pay more for housing and suffer other costs. The standard economic model of migration assumes a spatial equilibrium, so that the marginal migrant is indifferent between the city and the rural hinterland, which implies that there is no direct welfare benefit from encouraging migration. Certainly, there may exist externalities from moving to cities, but these can be both positive and negative and we currently cannot tell whether those external benefits on net favour cities (Glaeser and Gottlieb, 2009).

Moreover, accepting a role for spatial bias in policies sets an uncomfortable precedent. Spatially biased policies may well be used to favour politically powerful regions, rather than regions that should be subsidized. Loud voices will clamour for support for poorer regions, even if economic development suggests that people should leave such areas. A principle of spatial neutrality would seem to be the safest course, which would force regions to compete for capital and workers rather than relying on largesse from the national government.

A more sensible policy alternative is to focus on reducing artificial barriers to urban growth and improving the quality of urban life. If cities have benevolent economic effects, then it can be quite costly to impose land-use regulations that stymie urban construction, such as the stringent floor-area requirements that Mumbai has had for most of the past 50 years. In some cases, including Mumbai, these land-use controls were imposed to limit the growth of the city. Often, they have only prevented legal, safe housing and left a back door for the growth of sprawling slums.

City governments can also bring urban growth by becoming more effective at improving urban quality of life. Most of the downsides of density, such as contagious disease and congestion, are negative externalities that become magnified when people live close to one another. By reducing these externalities, developing-world cities can attract immigrant entrepreneurs and allow more people to enjoy the added productivity in cities.

In this section, we focus on the downsides of density and make four primary points. First, growing cities need infrastructure, but to get infrastructure right, we need to get institutions right. Second, incentives must accompany infrastructure. Third, property titling and the protection of private property are extremely valuable in urban contexts. Fourth, infrastructure, incentives, and institutions must be adapted to local conditions. The last point runs throughout our discussion and is not treated as a distinct topic.

We focus on four downsides of density: contagious disease, traffic congestion, crime, and high housing costs. The first two downsides are standard externalities that are exacerbated by urban proximity. Closeness enables bacteria and viruses to travel more readily over people. Crowded roads diminish mobility. Crime increases in cities because density increases the number of possible targets and makes solving crimes more difficult. High housing cost is not a classic externality, but still reflects a primary cost of urban life that can be reduced through sensible public policy.

(i) Infrastructure, institutions, and limited state capacity

Infrastructure, such as water pipes and sewers, can reduce water-borne illnesses like cholera that can plague cities. Traffic congestion is a particularly important urban cost; developing-world workers can easily spend 2 hours commuting each day. More roads can enable faster commutes, and so can bus transit and rail links.

Yet infrastructure is expensive and it is possible to waste billions on foolish infrastructure projects, especially if it is to be funded by subsidies rather than user fees. Transportation economics has documented many cases, especially in urban rail systems, where costs far exceed benefits. The life-saving social value of clean water makes white-elephant sewer lines less likely than white-elephant metro lines. Yet corruption in water and sewerage provision is also a real problem.

Achieving good urban infrastructure also depends on establishing appropriate institutions, for infrastructure does not emerge in a vacuum. Cities have four primary options when considering the right delivery mechanism for infrastructure, summarized in [Figure 8](#).

The first model for infrastructure delivery is a public system integrated into the executive branch of the national or local government. Whenever governments build roads or sewers directly, they are using this form of institution. It is also the oldest system,

Figure 8: Typology of infrastructure providers

<p>Public integrated</p> <p><i>Examples:</i> city departments of public works, Transport for London</p> <p><i>Features:</i> weak incentives, strict labour rules, political control, tax funding</p>	<p>Public independent</p> <p><i>Examples:</i> Port Authority of New York and New Jersey, Tennessee Valley Authority</p> <p><i>Features:</i> career-based incentives, increased labour flexibility, political independence.</p>
<p>For-profit independent</p> <p><i>Examples:</i> Cintra, ConEdison, Manila Water</p> <p><i>Features:</i> strong incentives, high labour flexibility, political independence but prone to corruption</p>	<p>Non-profit independent</p> <p><i>Examples:</i> turnpike trusts</p> <p><i>Features:</i> modest incentives, high labour flexibility, limited ability to raise capital from taxes or markets</p>

followed by the Tarquin Kings who built the Cloaca Maxima, ancient Rome's sewer system, around 600 BC. The costs and benefits of integrated public systems are the costs and benefits of all direct public provision.

Public entities have weak incentives to cut costs, which may lead to higher quality of services in the case of non-contractible quality, or just more expensive projects (Hart *et al.*, 1997). They can also be bound by strict rules regarding pay and personnel (Bold *et al.*, 2009). Since tax dollars can be used to fund projects, public provision can lead to white elephants that provide little large-scale benefit. Some discipline can be put on that risk if governments rely on high-quality cost–benefit analysis. Yet since wealthy developed countries like the US frequently ignore such analysis, it may be asking too much for developing-world governments to acquire and adhere to such analysis.

Any large-scale spending is an opportunity for large-scale theft, and infrastructure's costs tend to get particularly bloated because of corruption. In the case of public provision, the opportunities for corruption expand with the share of inputs provided by the private sector. The most standard form of corruption occurs when private providers charge too much for inputs and repay some of their profits to the politically powerful. The New York County Courthouse, built during the heyday of Boss Tweed in the 1860s, ended up costing over 30 times the original estimates, and payments to Tweed ring members might have exceeded \$10m.

Direct political control creates the risk that projects will be driven largely by political aims, such as employing large numbers of political supporters. Political control also generates some accountability. If the public is really unhappy with an infrastructure project, then elected officials can be held responsible. The Mayor of London, for example, is clearly held accountable for the functioning of London's transportation system.

Independent public authorities rose to prominence in nineteenth century America, during an era in which many voters did not trust local elected officials. The Erie Canal and the Croton Aqueduct were both examples of infrastructure built by independent public entities. The Port Authority of New York and New Jersey, the Triborough Bridge Authority, and the Tennessee Valley Authority are famous twentieth century examples of independent public authorities. The advocates of such authorities thought that they

provided means of avoiding the politicization of infrastructure and reducing corruption. Typically, such entities had leadership that was appointed by a higher level of government.

The de-politicization of these authorities carries both benefits and costs. Robert Moses, the legendary leader of the Triborough Bridge Authority, was able to build quickly and typically ignore short-run political objectives. However, he also famously ran roughshod over local objections to his projects. Moreover, independent authorities' political independence can certainly be compromised, as illustrated by the Port Authority's role in closing highway lanes in Fort Lee, New Jersey in September 2013, allegedly to punish the city's mayor for opposing Governor Chris Christie.

Bold *et al.* (2009) suggests that independent authorities may be able to provide stronger incentives for employees than standard integrated public entities that are bound by a variety of work rules. Certainly, the compensation packages at independent authorities have sometimes been far more generous than in standard public bureaucracies. Yet this freedom will only be valuable if top management itself has strong incentives, which likely come either from political overlords or career concerns.

At their best, independent authorities are led by competent executives, who are motivated by a desire for public acclaim and the potential to impress private-sector employers. At their worst, they become fully politicized fiefdoms of elected officials, and their independent status provides a cover for incompetence. They can be a solution when local government is malfunctioning, but they can also make matters worse.

A third approach is the public-private partnership, which occurs when a private for-profit operator is expected to build and operate a public service, such as a water system or transit mode. Private entities typically have strong incentives to cut costs, and the ability to flexibly motivate and hire workers. **Hart *et al.* (1997)** emphasize one downside of private provision: overly strong incentives to reduce quality in order to cut costs. Historically, their most important downside has been their tendency to corrupt governments. This corruption has either taken the form of bribery for cash subsidies or bribery for non-cash benefits.

New York City, for example, entrusted Aaron Burr's Manhattan Water Company with providing clean water to the city in 1799. The company was subsidized by the ability to do banking activities, which was quite valuable since New York City had only one other chartered bank at the time. Unsurprisingly, the water company pursued the lucrative practice of banking and did little water provision. The Bank of the Manhattan Water Company eventually became the Chase Manhattan Bank, later folded into J.P. Morgan Chase.

Private provision of water can certainly work, but it can also lead to excessive subsidies and under-provision of quality water. Werner **Troesken (2004)** documents how the move from private to public water provision in the US led to improved connections and health outcomes for African-Americans. This is partly because private water companies in southern US cities seem to have systematically favoured the richer, whiter neighbourhoods there.

In later years, America's private traction companies became practically synonymous with corruption, as they bribed city councils and state legislature to receive rights-of-way on city streets for decades or even centuries. London's electrified deep-level subways were built by Charles Yerkes, who had gotten his start with the Philadelphia Wideners, who turned corrupting the city council into an art, and then made his fortune using similar methods in Chicago.

Engel *et al.* (2014) survey public–private partnerships (PPPs) in the world today and find a wide range of experiences. In strong states, such as Chile, the experience has been largely benign. In weaker states, including some in sub-Saharan Africa, PPPs have often either subverted the government or suffered expropriation. The risks of corruption increase when public dollars are used to subsidize private companies, since larger subsidies are such a natural objective of private companies. Even when a well-run auction ensures an initially good deal for the public, later renegotiation is quite common and typically leads to settlements that are far more generous to the private provider.

PPPs can be a plausible institution in some circumstances, but they do not avoid the problem of weak capacity. A government that is unable to build a road competently on its own is also unlikely to be strong enough to withstand attempts at subversion by a capable for-profit entity. The case for PPPs and independent public authorities becomes stronger when those entities are actually made self-financing through user fees. Infrastructure that funds itself is more likely to provide social value. User-fee financing also increases the incentives for road maintenance, since user fees are only collected if drivers use the roads (Engel *et al.*, 2014). User-fee financing is less appropriate if the marginal cost of using the infrastructure is substantially below the average cost of providing the service, or if the infrastructure generates significant positive externalities, which may occur when bus rapid transit reduces congestion on standard car lanes.

A final option is non-profit transport provision, such as the turnpike trusts that operated in England during the eighteenth century. The robust non-governmental organization network in sub-Saharan Africa would seem to make non-profit provision an option today as well. Yet non-profit entities typically lack the financing capacity of either governments or for-profit entities, since they have neither the power to tax nor opportunities for future profit. Consequently, they are unlikely to provide the large-scale investments needed in developing-world cities.

In nations with weak states, the optimal institutional arrangement involves a choice between imperfect options. The case for private provision depends on the ability to put safeguards in place that limit corruption. For example, if the public sector can commit to a no-subsidy rule, and require the private provider to fund itself entirely through user fees, then the public–private option becomes safer.

The case for independent authorities depends on strong visible leadership. If the authority can attract a well-known figure with a reputation to lose, then its chances for success increase dramatically. For example, Bangalore’s high-technology cluster has produced a number of global business leaders who make ideal independent authority leaders, such as Nandan Nilekani who was, indeed, successful leading India’s Unique Identification Authority.

Direct state control remains the fall-back for services that require subsidies where no strong independent leaders are available. In that case, it is particularly necessary for different levels of government to monitor each other and for civil society to keep watch. Unfortunately, weak governments make it far more difficult to efficiently build the infrastructure that developing cities need.

(ii) Infrastructure and incentives

Even if infrastructure is built competently, it may still fail to deliver social value if individual actors don’t cooperate. When it comes to roads, a major problem is that driving

increases so that congestion makes the roads impassable. In the field of sewers and sanitation, users may be unwilling to pay for connections. In both cases, incentives are needed to accompany the infrastructure development.

More than 50 years ago, [Downs \(1962\)](#) proposed Downs' Law: peak-hour congestion rises to meet maximum capacity, which means that new roads will inevitably be flooded by new drivers. [Duranton and Turner \(2011\)](#) empirically document a 'fundamental law of road congestion': highway miles travelled increase roughly one-for-one with highway miles built. One reason for increased driving is that highways spur suburbanization ([Baum-Snow, 2007](#)). This implies that the behavioural response is likely to completely offset the extra lanes added by new construction. Braess's Paradox shows that in some circumstances adding lane capacity can even make traffic congestion worse.

For 50 years, economists have advocated a simple solution to traffic delays: congestion pricing. While this idea dates to Arthur Pigou, William Vickrey is credited with making the most plausible case for efficient road pricing in cities. Singapore adopted congestion pricing in 1975 when it was a relatively poor place, and today Singapore is moving towards congestion pricing based on global positioning system metres. London has congestion charging that also makes its streets more agreeable. Yet despite the potential of congestion pricing to make road traffic move more swiftly, popular opposition has generally limited its implementation.

The benefits of transportation are largely reaped by the user, but the costs include large externalities, and hence the policy problem is to ensure that roads do not get over-used. The benefits of public health infrastructure are largely external, especially with sewers, and hence the policy problem is to ensure that the infrastructure actually gets used. In New York City in the 1840s and in Africa today, there is a last-mile problem: while water- and sewer-related infrastructure is built, poorer households are unwilling to pay for connection to the system. In both cases, the poor understandably want to avoid the connection fees and are perfectly willing to stick with shallow wells and pit latrines.

In Manila, where septic tanks are standard, poor families are often unwilling to accept free de-sludging of their systems. The benefits of de-sludging are almost entirely external to the household, yet the installation of indoor septic tanks incurs much inconvenience. Inducing the adoption of water is somewhat easier than inducing adoption of sewers, since users receive a larger share of the benefits of water themselves.

The natural solution for road overuse is a congestion charge. The hard problems are enforcing this charge and getting the political will to get it through. In the case of sewers and sanitation, there are two plausible options: subsidizing connections to the water system or fining households that do not connect. Finding the finances for subsidies may be difficult, and a largely subsidy-financed system will be prone to waste. Enforcing penalties on households who use traditional shallow wells and pit latrines is also politically difficult and prone to abuse.

[Ashraf et al. \(2016\)](#) present a simple model in which either subsidies for adoption or penalties for non-adoption can be used to induce the spread of a sanitary technology. With weak institutions, subsidies can generate waste, while penalties can lead to corruption and even extortion of the innocent. The paper associates subsidy-related waste with weakness in government's executive functions and penalty-related waste with weakness in government's judicial function. When the executive is very weak relative to the judiciary, then penalties are optimal. When the judiciary is weak relative to

the executive, then subsidies make sense. When both are weak, then one solution is to use mild penalties that are too low to lead to extortion of the innocent, and do the rest of the work with subsidies. Once again, the appropriate institution depends on local conditions, such as the relative strength of the executive and judiciary branches of government.

New York's death rates did not start to fall after the Croton Aqueduct was built, but rather after 1866, when the Municipal Board of Health began imposing penalties for non-adoption. To reduce corruption, the inspectors were made independent of the police force and made to report directly to a board led by health professionals. Their task was made easier by the fact that property title was far better defined in 1866 New York than in most developing-world cities. Unclear property ownership makes it difficult to impose obligations that could reduce the downsides of density.

(iii) Property ownership in the city

At low levels of density, ownership is easier to assess. A farmer surrounded by empty land can, perhaps, be assumed to be the natural claimant to that space. In dense cities, the lines of possession become blurrier. Cities always have some form of shared space, like roads, and ownership over those spaces is inherently ambiguous. The public sector is typically the owner of shared spaces, but there are also many privately owned shared spaces, such as the plazas that appear at the feet of skyscrapers. Moreover, when informal settlements form on public ground, or on commandeered private space, it may be quite difficult to ascertain who has property rights over the space.

The Common Law of England evolved over centuries and paid a great deal of attention to land law, primarily because land was the primary pre-modern form of wealth. England consequently had well-developed tools for ascertaining ownership long before people started crowding into large cities. In much of the developing world, cities have formed before property-related institutions matured. In many cases, formal rules are borrowed from European countries, but the institutions to actually enforce those rules are weak. For instance, courts are not trusted in much of the developing world (Glaeser *et al.*, 2016).

Consequently, many developing-world cities have residents who lack any clear title to their land. In some cases, the demand for title is also weak because it is unclear what institutions will protect formal property rights. The costs of limited property assignment are quite significant. De Soto (2000) argues that 'because rights to these possessions are not recorded', it is impossible to use them to start businesses. Field (2007) shows that assignment of property rights increases labour supply, presumably because workers need to spend less time protecting their property. Unclear ownership also reduces the incentive to invest in property improvement, and as discussed above, it is hard to impose obligations when ownership is unclear.

The success of western cities depended on orderly transfer of property from one use to another. Tenements were built on former farm land. Mansions were destroyed to build commercial skyscrapers. When land isn't properly owned, it can't be easily sold, preventing the modernization of a city. In sub-Saharan Africa, slums often occupy the most centre land in the city. If those homes were owned, then the owners could benefit by selling the land and the city would benefit from a more centralized and efficient business district.

Property ownership is also crucial because it is a pre-condition for property or land taxation. An effective urban government is typically self-financing, which enables local initiative and avoids costly haggling over transfers from the national government. When local governments are dependent on property taxes, they are incentivized to deliver local amenities to increase property values. Even more importantly, property taxes are more efficient than taxes on earnings or sales: because property can't move easily, taxing land doesn't distort households' or firms' location decisions as much.

While the need for titling seems clear, the path towards clearly defined ownership is more ambiguous, but the English legal history does provide something of a guide. In the twelfth century, juries, a group of well-informed locals, were used to ascertain the long-term resident of a spot of land. Jury-like institutions still provide the natural tool for allocating property in developing-world slums.

In some cases, the long-term resident of a particular home may be obvious, and the allocation of property may be straightforward. In other cases, juries may be able to determine whether a family is a long-term resident of a particular slum, but not who lives in a particular home. In that case, the tools of mechanism design can be used to allocate homes across potential residents. For example, the jury-determined residents of a particular slum can be allowed to participate in a slum-specific auction, where units are sold in exchange for payments, which are then allocated against future property taxes. The government can, in principle, lend the money but it must be willing to take back the property if these payments are not made in the future. Particular residents, such as families with small children, can be loaned extra funds for bidding. Residents who are out-bid can be compensated with the funds paid by the winning bidders.

This mechanism sketch is only one possible means of allocating property. Implementation should depend on local conditions. The larger point is that more widespread titling has the capacity of making cities more functional. Naturally, that agenda also requires institutions that make ownership attractive.

Glaeser *et al.* (2016) argue that in a regime with weak judicial institutions, bright-line rules are less prone to subversion than more complex arrangements. In the case of property ownership, this implies that property should be protected with fixed penalties, or injunction, rather than flexible liability rules that are factually intensive and thus more vulnerable to judicial subversion.

This distinction is particularly clear in the case of 'eminent domain', which is typically administered through a liability rule today. Owners are compensated based on the assessed damage from losing their property. Such assessments are frequently disputed in both the developed and developing world, and many owners often argue that the courts have given them too little compensation. One natural alternative, which has been in Indian law for 5 years, is that eminent domain can only be used when a super-majority of residents agree. This super-majority rule ensures that most residents have been made better off by the taking, but also avoids the hold-out problem, where a single owner blocks the transaction for everyone else.

Protecting homes and land is part of the larger problem of battling crime in developing-world cities. In the developing world, crime prevalence differs substantially from continent to continent. In Latin America, violent crime is the dominant urban problem. Drug gangs often play a major role in this type of crime, and the police are frequently ineffective. Violent crime is also significant in parts of Africa, but is much less of a problem in the cities of India or East Asia.

The basic economics of crime-fighting has been understood since [Becker \(1968\)](#). Some combination of deterrence and incapacitation is necessary, but institutional weakness in the developing world makes crime-fighting particularly challenging. Often the police function essentially as a rival gang, rather than as a crime-fighting unit. Often the key question is how police leaders can be given the right combination of incentives and capacity to actually reduce crime.

One particularly dramatic tension in crime-fighting is the distinction between overwhelming force and community policing. Strategies like stop-and-frisk in New York or the UPPs⁶ in Rio de Janeiro follow an overwhelming-force approach that totally overwhelms a neighbourhood. Community policing means developing social relationships that enable police to leverage their knowledge of the neighbourhood. Both approaches have costs and benefits, but it seems most likely that community policing is a more attractive strategy in relatively functional neighbourhoods that are not overwhelmed by violence, and that overwhelming force is the only possible approach when homicide rates get sufficiently high.

Weak states create two crime-related problems. They produce a police force that is either incapable of enforcement or that extorts the public. Over the years, a variety of tools have come about to improve policing. Brazil, for example, has adopted a quasi-military option in which the uniformed police operate under strict discipline. Much earlier, the US and Chile improved police honesty through rotation. Moving police around reduced the ability to form corrupt local relationships. Data analysis has been an important part of New York's more recent attempts to reduce violence.

In this section, we have discussed a variety of strategies aimed at reducing the negative externalities associated with urbanization. These problems are not distinct from the larger goal of obtaining the higher urban productivity. The benefits of urban economic strength will not be fully realized unless developing-world cities are made more livable.

V. The larger benefits of urbanization

Section II discussed the cross-sectional differences in productivity across cities in the developing world. A simple static decomposition suggests that urbanization can lead to income growth by moving people from less productive places to more productive places. This section concludes by discussing the possibility that cities may generate positive externalities for the nation as a whole.

Perhaps the most obvious potential benefit of urbanization is that it lowers the number of people who are farming the land, which should increase the returns to agriculture. More speculative economic benefits may include technological improvements, which will then spread nationwide. These improvements can come from urban innovators themselves and from imported foreign technology that comes in through cities.

The correlation between urbanization in 1960 and subsequent growth that was discussed in section II may reflect such technological gains, but it may also reflect a connection between urbanization and the quality of government. Many authors have noted

⁶ Unidade de Polícia Pacificadora, or Pacifying Police Unit, is a law enforcement programme pioneered in the Brazilian state of Rio de Janeiro that aims at reclaiming territories controlled by gangs of drug dealers.

a link between urban density and democratic revolution. [Carp \(2007\)](#) shows the role of urban connectivity in enabling the coordinated efforts of the American Revolution. [Wallace \(2013\)](#) shows empirically that urbanization increases the probability of regime change for dictators: among dictatorial regimes, a one log point increase in the size of the largest city increases the number of regime changes by 0.5 log points.

[Glaeser and Steinberg \(2016\)](#) argue that there are several reasons why cities seem to have been the wellspring of successful democratic revolutions throughout history. Cities speed the flow of ideas, including those that emphasize the shortcomings of any current regime. Cities enable surreptitious organization, including the formation of anti-regime groups. Large populations in political capitals are able to put direct physical pressure on ruling elites, as the residents of Cairo did by taking control of Tahrir Square.

Yet, as the example of Tahrir Square illustrates, successful revolutions may not lead to successful democracies. The range of economic possibilities that exist in a city should help boost the demand for economic freedom. Yet the downsides of urban density also increase the demand for strong leadership that can tame congestion, crime, or contagious disease. It is unclear whether urbanites have the incentives to push for lasting democracy once a dictator has been toppled. Perhaps they will prefer authoritarian regimes that allow abundant economic freedom, but that take a strong hand against quality-of-life infractions, an example being the chewing gum ban in Singapore.

It is also unclear whether cities create the civic capital that is needed to promote good governance in any regime. There are certainly examples of urban groups that have fought corruption, such as New York's Committee of Seventy that battled the Tweed Ring in 1871. Yet many urbanites remain as politically supine as most subsistence farmers, and so it is unclear whether mass urbanization will fundamentally lead to better government. This remains yet another pressing topic for future research.

VI. Policy conclusions

This paper emphasized the dramatic productivity differences across space within developing countries. Urban productivity is much higher than rural productivity and urban productivity rises as cities grow. Better-educated places are far more productive than less-educated places. Here we discuss four implications of this analysis.

First, there are large potential costs from artificially restricting the growth of cities. Many developing-world leaders are concerned about the chaotic growth of urban areas and their instinct is to limit city growth. The high productivity of developing-world cities argues against this perspective. In some cases, cities have explicitly adopted draconian land-use controls, such as tight limits on density, partially to restrict urban growth. In many cases, these policies merely produce dysfunctional cities through sprawling out rather than going up. While this paper has analysed no particular land-use restriction, the benefits of agglomeration would seem to suggest that these policies should be carefully screened with cost-benefit analysis, and those restrictions that carry more costs than benefits should be scrapped.

Our second conclusion is that cities should rethink their approaches to the disamenities of urban life. These downsides of density are not merely nuisances. They reduce urban growth and make it harder for cities to play their larger economic role. In section

IV, we discussed a variety of different approaches to urban disamenities. In some cases, better incentives, such as congestion pricing, are needed. In other cases, institutional reform may be required to make urban service provision more effective.

Our third conclusion is that education is the bedrock of urban success. Urban density can spread more ideas when urbanites are more knowledgeable and have better communication skills. This underscores the value of human capital investments in developing countries.

Finally, while we know too little about urban entrepreneurship in the developing world, there is surely some basis for targeted policy interventions that aim at increasing effective innovation. Such innovations would help us to assess the importance of urban entrepreneurship. Similarly, it is worth re-examining the regulatory barriers to local entrepreneurship, including licensing rules and labour market regulations. In some cases, the costs of these rules may greatly exceed the benefits.

Cities have the capacity to empower economic change. The developing world needs policies that enable their cities to grow and become safer, healthier, and more functional.

Data appendix

Detailed firm-level data, based on which we calculate individual firms' productivity, come from the Chinese Industrial Census (CIC), also referred to as Annual Survey of Industrial Firms. An annual firm-level survey conducted by China's National Bureau of Statistics (NBS), the dataset covers mining, manufacturing, and public utility industries. In this paper, we focus on manufacturing firms, and the decade between 1998 and 2007 in which the surveys contain necessary variables for standard productivity measurements.

Firms presented in datasets are relatively large. According to the official documentation, the surveys include all state-owned enterprises (SOEs), and non-state firms with sales greater than 5 million CNY ('above-scale' firms). Yet, in fact, a significant number of below-5-million non-state firms, accounting for about 5 per cent of the unbalanced panel, are also included. In comparison with the 2004 Economic Census that covers the universe of industrial firms, the CIC excludes 80 per cent of firms, yet these 'below-scale' firms only accounted for 28.8 per cent of industrial workforce, 9.9 per cent of output, and 2.5 per cent of exports (Brandt *et al.*, 2012).

A substantial amount of data work is necessary before empirical analysis. Key issues include standardizing industry codes and address IDs, matching individual firms over time, deflating nominal variables, etc. Brandt *et al.* (2012) provide a comprehensive description of such data work in their online appendix. This paper follows their steps and has largely replicated their results.

Table A1 lists aggregates of individual firms' balance sheet items. In our sample period, the total number of manufacturing firms doubled from about 165k to 337k. Aggregate (nominal) industrial output, value added, and export experienced an impressive six-fold increase, while total employment grew by 40 per cent, suggesting considerable productivity gains.

Table A1: Summary statistics, CIC full sample, 1998–2007

Year	No. firms	Employment	Labour compensation	Net value of fixed assets	Intermediate input	Value added	Industrial output	Sales	Sales from export
1998	165,120	56.44	0.53	4.41	5.11	1.94	6.77	6.54	1.08
1999	162,034	58.05	0.54	4.73	5.42	2.16	7.27	7.06	1.15
2000	162,887	55.59	0.58	5.18	6.40	2.54	8.57	8.37	1.46
2001	171,259	54.41	0.65	5.54	7.11	2.83	9.54	9.32	1.62
2002	181,559	55.21	0.71	5.95	8.23	3.30	11.08	10.86	2.01
2003	196,222	57.48	0.82	6.61	10.58	4.20	14.23	13.95	2.69
2004	279,038	66.20	1.10	7.97	15.13	6.05	20.50	20.43	4.05
2005	271,836	68.96	1.29	8.95	18.79	7.22	25.16	24.69	4.77
2006	301,961	73.58	1.60	10.58	23.62	9.11	31.66	31.08	6.05
2007	336,769	78.75	2.02	12.34	30.18	11.70	40.51	39.76	7.34

Notes: Employment in millions of workers, all nominal values in trillion CNY.

References

- Arriaga, E. (1968), 'Components of City Growth in Selected Latin American Countries', *Milbank Memorial Fund Quarterly*, **46**(2), 237–52.
- Ashraf, N., Glaeser, E., and Ponzetto, G. (2016), 'Infrastructure, Incentives, and Institutions', *American Economic Review*, **106**(5), 77–82.
- Baum-Snow, N. (2007), 'Did Highways Cause Suburbanization?', *Quarterly Journal of Economics*, **122**(2), 775–805.
- Becker, G. (1968), 'Crime and Punishment: An Economic Approach', in N. G. Fielding, A. Clarke, and R. Witt (eds), *The Economic Dimensions of Crime*, Palgrave Macmillan, 13–68.
- Bloom, N., and Van Reenen, J. (2010), 'Why Do Management Practices Differ across Firms and Countries?', *Journal of Economic Perspectives*, **24**(1), 203–24.
- Sadun, R., and Van Reenen, J. (2016), 'Management as a Technology?', Working Paper, last updated 31 May 2016.
- Bold, T., Collier, P., and Zeitlin, A. (2009), 'Provision of Social Services in Fragile States', World Bank Working Paper, Washington, DC, World Bank.
- Borensztein, E., De Gregorio, J., and Lee, J.-W. (1998), 'How Does Foreign Direct Investment Affect Economic Growth?', *Journal of International Economics*, **45**(1), 115–35.
- Brandt, L., Van Biesebroeck, J., and Zhang, Y. (2012), 'Creative Accounting or Creative Destruction? Firm-level Productivity Growth in Chinese Manufacturing', *Journal of Development Economics*, **97**(2), 339–51.
- Carp, B. (2007), *Rebels Rising: Cities and the American Revolution*, New York, Oxford University Press.
- Chauvin, J., Glaeser, E., Ma, Y., and Tobio, K. (2016), 'What is Different about Urbanization in Rich and Poor Countries? Cities in Brazil, China, India and the United States', *Journal of Urban Economics*, in press.
- Chinitz, B. (1961), 'Contrasts in Agglomeration: New York and Pittsburgh', *American Economic Review*, **51**(2), 279–89.
- Combes, P.-P., Duranton, G., and Gobillon, L. (2008), 'Spatial Wage Disparities: Sorting Matters!', *Journal of Urban Economics*, **63**(2), 723–42.
- — — Roux, S. (2010), 'Estimating Agglomeration Economies with History, Geology, and Worker Effects', in E. L. Glaeser (ed.), *Agglomeration Economics*, Chicago, IL, University of Chicago Press, 15–66.
- Desai, M., Foley, C. F., and Hines, J. (2006), 'Do Tax Havens Divert Economic Activity?', *Economics Letters*, **90**(2), 219–24.
- De Soto, H. (2000), *The Mystery of Capital: Why Capitalism Triumphs in the West and Fails Everywhere Else*, New York, Basic Books, and London, Bantam Press/Random House.
- Downs, A. (1962), 'The Law of Peak-hour Expressway Congestion', *Traffic Quarterly*, **16**(3), 393–409.
- Durand, J., Peláez, C., Thomas, D. S., Carleton, R. O., Baumgartner, L., and Lee, E. S. (1965), 'Patterns of Urbanization in Latin America', *The Milbank Memorial Fund Quarterly*, **43**(4), 166–96.
- Duranton, G., and Turner, M. A. (2011), 'The Fundamental Law of Road Congestion: Evidence from US Cities', *American Economic Review*, **101**(6), 2616–52.
- Morrow, P., and Turner, M. A. (2014), 'Roads and Trade: Evidence from the US', *Review of Economic Studies*, **81**(2), 681–724.
- Ellison, G., and Glaeser, E. (1997), 'Geographic Concentration in US Manufacturing Industries: A Dartboard Approach', *Journal of Political Economy*, **105**(5), 889–927.
- Engel, E., Fischer, R., and Galetovic, A. (2014), 'Finance and Public–Private Partnerships', in *Financial Flows and Infrastructure Financing*, Reserve Bank of Australia.
- Fay, M., and Opal, C. (2000), 'Urbanization without Growth: A Not So Uncommon Phenomenon', Policy Research Working Paper No. 2412, Washington, DC, World Bank Publications.
- Field, E. (2007), 'Entitled to Work: Urban Property Rights and Labor Supply in Peru', *Quarterly Journal of Economics*, **122**(4), 1561–602.
- Glaeser, E., and Gottlieb, J. (2009), 'The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States', *Journal of Economic Literature*, **47**(4), 983–1028.

- Glaeser, E., and Mare, D. (2001), 'Cities and Skills', *Journal of Labor Economics*, **19**(2), 316–42.
- Steinberg, B. (2016), 'Transforming Cities: Does Urbanization Promote Democratic Change?', NBER Working Paper No. 22860, November.
- Kerr, S., and Kerr, W. (2015), 'Entrepreneurship and Urban Growth: An Empirical Assessment with Historical Mines', *Review of Economics and Statistics*, **97**(2), 498–520.
- Ponzetto, G., and Shleifer, A. (2016), 'Securing Property Rights', NBER Working Paper No. 22701, September.
- Scheinkman, J., and Shleifer, A. (1995), 'Economic Growth in a Cross-section of Cities', *Journal of Monetary Economics*, **36**(1), 117–43.
- Gollin, D., Jedwab, R., and Vollrath, D. (2016), 'Urbanization With and Without Industrialization', *Journal of Economic Growth*, **21**(1), 35–70.
- Greenstone, M., Hornbeck, R., and Moretti, E. (2010), 'Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings', *Journal of Political Economy*, **118**(3), 536–98.
- Hart, O., Shleifer, A., and Vishny, R. (1997), 'The Proper Scope of Government: Theory and an Application to Prisons', *Quarterly Journal of Economics*, **112**(4), 1127–61.
- Hsieh, C.-T., and Klenow, P. (2009), 'Misallocation and Manufacturing TFP in China and India', *Quarterly Journal of Economics*, **124**(4), 1403–48.
- Jones, P., D'Aoust, O., and Bernard, L. (2016), 'The Urban Wage Premium in Africa', Oxford/LSE Working Paper.
- Kerr, S., and Kerr, W. (2016), 'Immigrant Entrepreneurship', NBER Working Paper No. 22385, July.
- Moretti, E. (2004), 'Human Capital Externalities in Cities', ch. 51 of J. V. Henderson and J.-F. Thisse (eds), *Handbook of Regional and Urban Economics*, Vol. 4, 2243–91, Elsevier.
- Rauch, J. (1993), 'Productivity Gains from Geographic Concentration of Human Capital: Evidence from the Cities', *Journal of Urban Economics*, **34**(3), 380–400.
- Sutton, J., and Langmead, G. (2013), *An Enterprise Map of Zambia*, Vol. 4, International Growth Centre in association with the London Publishing Partnership.
- Kellow, N. (2010), *An Enterprise Map of Ethiopia*, International Growth Centre in association with the London Publishing Partnership.
- Kpentey, B. (2012), *An Enterprise Map of Ghana*, Vol. 2, International Growth Centre in association with the London Publishing Partnership.
- Troesken, W. (2004), *Water, Race, and Disease*, Cambridge, MA, MIT Press.
- Wallace, J. (2013), 'Cities, Redistribution, and Authoritarian Regime Survival', *Journal of Politics*, **75**(3), 632–45.
- Young, A. (2012), 'The African Growth Miracle', *Journal of Political Economy*, **120**(4), 696–739.