

Agglomeration effects in Colombia

Gilles Duranton^{*‡}
University of Pennsylvania

May 2015

ABSTRACT: I estimate an elasticity of wages with respect to city population of about 5% for Colombian cities. This finding is robust to a number of econometric concerns. The second main finding is a negative effect of market access on wages. Third main finding regards stronger agglomeration effects in the informal sector. This third finding provides an explanation for a range of other negative findings, including only weak evidence in favour of human capital externalities, no evidence of a complementarity between cities and skills, and an absence of stronger agglomeration effects for older workers as would be predicted by the existence of learning effects. I do not find measurable effects of roads or amenities on wages either.

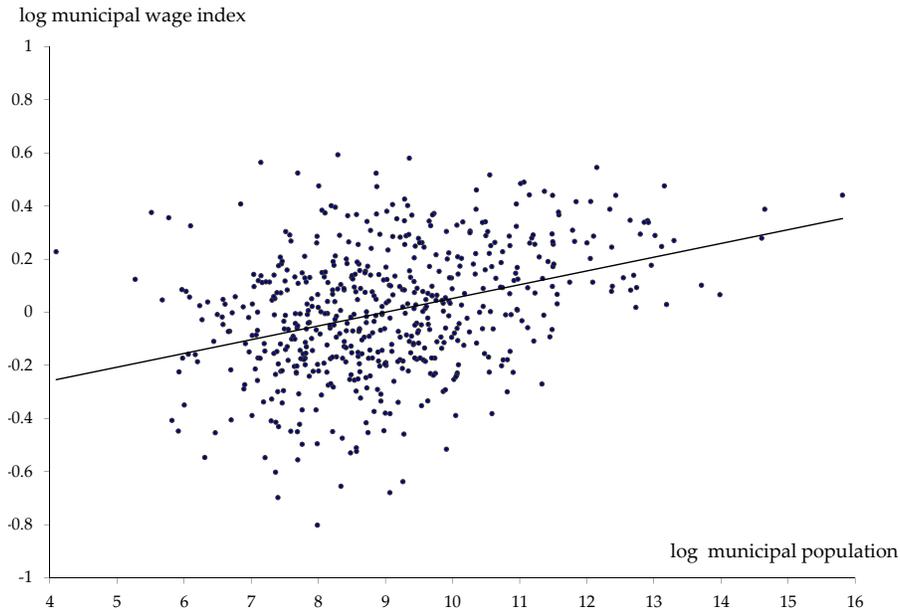
Key words: agglomeration, Colombia, city population, market access

JEL classification: R12, R23

^{*}Thanks to Prottoy Akbar for terrific research assistance. I am grateful to Jose Antonio Pinzon, Diana Lopez, Adolfo Meisel, and Fabio Sanchez for providing me data and helping with them. Thanks also to the editor, Gianmarco Ottaviano, three anonymous referees, Kristian Behrens, Matt Kahn, and Thierry Mayer for useful feedback and comments. Financial support from the Colombian Ministry of Planning (DNP) and UN-Habitat is gratefully acknowledged.

[‡]Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104, USA (e-mail: duration@wharton.upenn.edu; website: <https://real-estate.wharton.upenn.edu/profile/21470/>).

Figure 1: Population and wages in Colombian municipalities, 2008-2012



The vertical axis represents centered (natural) log municipal wages computed using 2008-2012 individual wage data after conditioning out year effects and individual characteristics. The horizontal axis represents log 2010 population in the urban part of municipalities. There are 583 municipalities. The slope of the regression line is 0.051 and the R-squared is 12%.

1. Introduction

In this paper I estimate agglomeration effects for cities in Colombia. My preferred estimate regarding the elasticity of wages with respect to city population is about 5%. Figure 1 provides a summary illustration for the recent years. It plots a log municipal wage index and log municipal population. The slope of the regression line in the figure is 5.1%. This finding is robust to a number of econometric concerns. My second main finding is a negative effect of market access on wages. Places with good external market access actually offer lower nominal wages. This is consistent with better market access allowing local workers to consume a broader variety of goods at a lower price coupled with a notion of spatial equilibrium leading to real wage equalisation. To my knowledge, this result is novel in the literature.

Third, I find evidence of stronger agglomeration effects in the informal sector. This finding provides an interesting unified explanation for a range of other negative findings. There is only weak evidence favour of human capital externalities, whereby wages in cities would increase with the local share of educated workers, all else equal. There is no evidence of a complementarity between city population and individual skills. I also find higher returns to city size for younger

workers. This is opposite of what would be predicted by the existence of greater learning effects in larger cities. I do not find measurable effects of roads or amenities on wages either. Given the precision of most estimates, these negative results should be thought as zeroes rather than inconclusive findings.

This investigation is important for several reasons. First, many surveys on urbanisation in developing countries lament the paucity of results about agglomeration economies in these countries (Henderson, 2005, Overman and Venables, 2005, Duranton, 2014). This analysis hopes to be part of what will hopefully be a broader knowledge base about urbanisation and its effects in developing countries. For instance, the elasticity of wages with respect to city population of 5% found for Colombia is slightly larger than comparable estimates for cities in developed countries as discussed in Puga (2010). At the same time, this elasticity of 5% is also smaller than recent findings for less advanced developing countries like China or India (Chauvin, Glaeser, and Tobio, 2013, Combes, Démurger, and Shi, 2015). This perhaps suggests that agglomeration effects are stronger in less advanced countries. In turn, if this is confirmed, understanding this feature in greater depth and confirming the conjecture proposed below should become an important research priority.

Second, this study wishes to take a *broad look* at agglomeration effects in Colombia. Most studies in the literature focus on specific dimensions of agglomeration economies or specific identification issues. For instance Ciccone and Hall (1996) and many others have focused on the possible endogeneity of city population. Rauch (1993), Moretti (2004) and many of their followers have explored the effect of city education on wages. Glaeser and Maré (2001) and De la Roca and Puga (2012) emphasised possible learning effects associated with larger cities. Glaeser and Resseger (2010) also suggest a complementarity between education and city population, etc. Although data and space limitations make it impossible to take into account everything that previous literature has considered, it is important to develop a more integrated approach that brings together most of the first-order issues highlighted by past research. This more integrated approach should also be useful to assess the robustness of many previous piecemeal findings and offer a broad assessment of agglomeration effects in Colombia.

Finally, knowing more about the benefits from cities is important for Colombia. As the Government of Colombia makes significant change to its system of municipal finance and ponders more broadly what sort of urban policy it intends to develop and implement (Samad, Lozano-Gracia, and Panman, 2012, Barco, 2013), it is important to know more about the economic benefit and

costs of cities in Colombia. This paper is part of that enterprise.

To fulfill the objectives of this paper, one key tradeoff must be resolved. Namely, available data preclude the use of some of the econometrically more advanced approaches that have been developed in the literature that are often extremely data-intensive and thus impossible to implement for most developing countries. For instance Combes, Duranton, Gobillon, and Roux (2010) use a rich panel of workers to control for unobserved individual characteristics that may be correlated with location choices. Such rich data are not available in Colombia, nor in most developing countries. Hence, expanding our knowledge base on agglomeration effects in developing countries will require some methodological compromises. At the same time, serious identification concerns should not be ignored. Put differently, although this study does not attempt to push the methodological frontier on the study of agglomeration effects, it endeavours to use the best available approaches given the data at hand.

Because comparability across countries is important, this study will focus mostly on widely available data.¹ It will also report both positive and negative results. Despite being often perceived unhelpful for publication, negative results should be reported. Otherwise it is impossible to know whether a relationship is absent because of a lack of positive result or because of a lack of investigation.

The rest of this paper is organised as follows. Section 2 describes the data. Section 3 discusses some simple agglomeration theory, its predictions, and the associated econometrics. Section 4 presents the main results regarding the elasticity of wages with respect to city population. Section 5 reports and discusses results about human capital externalities, market access, and infrastructure and amenities. Finally, section 6 concludes.

2. Context and data

At the end of our study period in 2010, Colombia had a population of about 44 million over a territory of 1.1 million square kilometers. It is located at the Northern tip of South America. Like many countries of this continent, it is already highly urbanised with 75% of its population living in cities. Hence, the magnitude of agglomeration economies is a concern for a vast majority of the

¹For instance, the innovative study of Greenstone, Hornbeck, and Moretti (2010) uses data from jurisdictions competing to host new large assembly plants. The ingenious design of the study circumvents serious identification concerns about unobserved local characteristics potentially correlated with employment changes. There is no hope of collecting such data for most countries. Nonetheless, there would be significant value replicating this study for a developing country but this represents a separate research project.

population. Colombia is also economically more advanced than developing countries in Africa and much of Asia with a GDP per capita of about \$8,000. While many developing countries are dominated by a primate city, Colombia has a balanced urban system with 60 municipalities with a population above 100,000. This provides useful variations for the estimation of agglomeration effects.

Note also that the geography of Colombia is extremely varied with much of its territory either dominated by the mountains of the Andes or, at the other extreme lowlands covered by an equatorial jungle. As a result, most of the Colombian population lives either along the Caribbean coast or on the high plateaus of the interior. We also keep in mind the existence of large spatial disparities within Colombia. The GDP per capita of the poorest Colombian department is only about a sixth of that of Bogotá. It is important to know how much of these disparities can be accounted for by agglomeration effects. Figure 2 provides for a map of the country and its main municipalities.

To measure agglomeration effects, much of extant research uses large-scale administrative data either at the worker or at the firm level. Such data only cover formal employment and would be inappropriate for cities in developing countries where half or more of the labour force is in the informal sector. I rely instead on the Colombian Great Integrated Household Survey (GEIH) administered by the Colombian National Statistical Agency (DANE) between 1996 and 2012.² This survey provides detailed information about a large cross-section of Colombian households including age, gender, year of education, municipality of residence, number of hours worked, and labour income.³ The initial data contain 5,437,894 observations ranging from 44,709 in 1996 to 647,050 in 2012. After eliminating individuals with no labour income and the 0.5% of workers with the lowest and highest wages every year, there are 2,739,768 observations left.

It is important to understand who the data cover and what is measured by labour income. First, and to repeat, the GEIH surveys household members that work in both the formal and the informal sector. In line with previous work on this survey, I use a broad definition of labour income that covers formal- and informal-sector wages as well as earned labour income for self-employed, be they street-sellers or taxi-drivers. My objective is to measure agglomeration effects across most of the workforce to avoid the potential selection issues one would face when focusing only on formal-

²Sometimes, the survey is conducted across consecutive years. That is why 2006 and 2007 appear as missing.

³The GEIH has been widely used by prior research, including Attanasio, Goldberg, and Pavcnik (2004) on the wage effects of trade reform or Angrist and Kugler (2008) on the effect of violence on wages, among many others. See panel A of table 1 for some descriptive statistics.

Figure 2: Map of Colombia and its main municipalities



Map using 2010 DANE population estimates

Table 1: Descriptive statistics

Panel A: Individual observations.	1996	2004	2012	all years	std. dev. (all years)
% male	62.67	57.13	55.30	56.80	49.53
Average age	35.55	37.71	39.03	38.06	13.31
Average years of schooling	7.52	8.50	9.40	8.88	4.71
Median municipal urban population	262,171	356,461	320,721	304,062	1,521,714
Panel B: 645 Municipalities.	p25	median	p75	mean	std. dev.
Total population (1993)	9,510	17,568	33,733	50,810	248,533
Total population (2005)	8,588	16,587	32,319	56,701	302,248
% pop. with postsecondary educ. (2005)	8.26	11.79	18.55	14.69	8.70
Market access index	30,941	123,036	280,977	201,870	239,181
Road length (km, 2012)	0	10.60	23.07	16.87	23.30
Distance to main road (km, 2012)	0	0	9.15	6.77	15.77

sector employees. This said, my coverage is less than 100% of participating workers because I need to exclude family helpers who report zero earnings to the survey. Because the earnings of self-employed may be differentially distorted across cities, I estimate some regressions on a more restricted sample that includes only wage earners. For similar reasons and to detect differential selection effects into occupations, I also verify the main results on a sample of only formal-sector workers.

Several other features of the GEIH are worth keeping in mind. First, as already mentioned the different waves of the GEIH are repeated cross-sections with no individual identifiers. This feature prevents us from including individual fixed effects in our estimations as Combes *et al.* (2010). Second, parts of the Colombian territory that were more recently administratively integrated with the rest of the country (the so called new departments) are not covered. This does not matter much for our purpose because these areas are extremely lightly populated and contain no significant city. In addition, this survey is not representative for smaller cities. Although this problem is likely to be much attenuated since I consider cities by size and not individually, I deal with it either by controlling for individual characteristics of workers or by restricting the sample to the larger cities. Finally, the exact definition of labour income in the survey changed twice during the study period. There is a first change in 2000 in how hourly wages are computed and a second in 2005 when several other surveys were integrated into the National Household Survey to form the GEIH. These changes were hopefully minor and uncorrelated with the questions at hand. However, and

because the number of surveyed households also greatly increased over time, the analysis will be duplicated for pairs of years separately. Finally, while the the GEIH contains information about occupations, it is hard to isolate public servants whose wage is not set by market mechanisms. As they might be over-represented in smaller cities, this may modestly bias downwards our estimates of agglomeration economies. See Departamento Administrativo Nacional de Estadística (2013) for further details about the GEIH.

This study uses the GEIH together with a wealth of municipal-level information. The first source of information are the Colombian censuses of 1993 and 2005 for population by municipality with a breakdown between the rural and the urban part of each municipality. Colombia is composed of 1,119 municipalities. Colombian municipalities are rather large and, on average, cover an area of 1,018 squared kilometres. This equivalent is to a disk with radius 18 kilometres. In most cases, municipalities are organised around one major settlement, referred to as the 'head' or urban part of the municipality. I complement the 1993 and 2005 population figures with DANE's municipal population estimates for 2010. The Colombian census also contains a wealth of socio-economic and education variables. Among them, the municipal share of the workforce with post secondary education is of primary interest.

To construct instruments for current population, I use population data collected from the 1843, 1870, and 1938 censuses. Although the first national census of Colombia dates back to 1812, the 1843 census, despite being probably subject to under-counting, is the first useable one (Bushnell, 1993). The details of the instrumentation strategy are postponed until later.

Other municipal characteristics are coming from several sources. Sánchez and Nuñez (2000) gathered a large number of geographical characteristics for Colombian municipalities including climate, availability of natural resources such as gold, coal, or oil, altitude, soils, or availability of water among others. This source is complemented with further municipal data collected by the National Planning Ministry (DNP) which cover land area, many socio-economic indicators, and some public finance variables. I also use a rich set of amenity and tourism variables collected by Meisel Roca and Pérez (2012) containing information about museums, public libraries, hotels, restaurants, and airport arrivals and departures.

I also use various measures of roads and market access for Colombian municipalities. Roads data are further described in Duranton (2015a) where I also compute market access variables. These data cover the Colombian road network from the country's colonial era to the current day

network and various market access variables measured either directly using wages in Colombian municipalities and travel time distance from Google Map or indirectly using a large-scale internal trade survey. See below for further details. Panel B of table 1 provides some descriptive statistics for the Colombian municipalities represented in the study.

While most of the analysis presented here uses the population of the urban part of Colombian municipalities, in some regressions I use total municipal population. I also sometimes aggregate municipalities into metropolitan areas (MA) as in Duranton (2015*b*) where metropolitan areas are defined using commuting patterns between municipalities.

3. Theory and estimation

The proposition that larger cities enjoy a productive advantage dates back to Smith (1776), who provided an insightful comparison between workers in Edinburgh and farmers in the Scottish highlands, and Marshall (1890), who argued that this productive advantage occurred because of more intensive input-output linkages, thicker local labour markets, and knowledge spillovers. Modern economic theory has provided a variety of microeconomic foundations for agglomeration economies (see Duranton and Puga, 2004, for a summary). Despite relying on different mechanisms, theories of agglomeration in production all predict that output increases more than proportionately with employment for a group of spatially proximate workers.

While there are different ways to define a given ‘group of spatially proximate workers’ such as the workers in narrowly defined sectors working within walking distance of each other, I will follow much of the recent literature and consider all workers in a city as defining the relevant group to measure agglomeration effects in my baseline analysis.⁴ To measure efficiency, I will also follow a dominant tradition in the literature and measure agglomeration benefits through wages. The prediction is then that, because of agglomeration effects, wages should increase with city population.

⁴A theoretical case can be made that agglomeration effects will, to some extent, take place within sectors of economic activity. These are referred to as localisation economies. In line with much of the literature, our focus here is on agglomeration effects taking place across sectors or irrespective of them. These are usually referred to as urbanisation economies. There are two reasons for this focus. First, tracking sectoral employment in Colombian cities is difficult given the existence of a large informal sector. Second, it is difficult to identify localisation effects as there is no obvious instrument for employment by sector. See Jofre-Monseny, Marín-López, and Viladecans-Marsal (2014) for a recent study on localisation and urbanisation effects.

More formally, I estimate variants of the following regression:

$$\log w_{ic(i)t} = \alpha \log Pop_{c(i)} + X_{it}\beta + A_{c(i)t}\gamma + \delta_t + \epsilon_{ic(i)t}, \quad (1)$$

where the dependent variable is the (natural) log of the hourly wage of person i who lives in city $c(i)$ in year t . The coefficient of interest is α , the elasticity of wages with respect to city population. The regression also includes a vector of individual characteristics X , a vector of city characteristics A , year indicators δ , and an error term ϵ .

Ideally one would like to compare the same workers across different, exogenously assigned locations. This is unfortunately impossible. Instead, I can only compare the wages of different workers who chose to be where they are.⁵ There are two main sources of bias in the estimation of equation (1). They both find their source in the fact that our explanatory variable of interest $\log Pop_{c(i)}$ is indexed by $c(i)$, that is the city c is *chosen* by worker i .

The first source of bias arises from the possibility that some city effects not included in the vector of city characteristics $A_{c(i)t}$ are correlated with city population and wages — a missing variable problem. Alternatively, workers may prefer to locate in cities where wages are higher — a reverse causation problem. Econometrically, the two problems are equivalent. I deal with them in two different manners. First, I use an extensive set of local control variables capturing the geography of Colombian municipalities, the education of their workforce, their infrastructure, their proximity to markets, their amenities, etc. Second, I implement a standard instrumental variable (iv) strategy and instrument current city population with long lags of population from 1843, 1870, and 1938 (Ciccone and Hall, 1996, Combes *et al.*, 2010). Given that past population is expressed in logs, this is equivalent to using past populations and long term population growth.

This iv strategy is valid if past populations are correlated with current wages only through current population. While this exclusion restriction is plausible, one can imagine a number of possible violations caused by alternative links between past population and current wages including, for instance, better market access, which may be extremely persistent, or some other geographical characteristics. These variables could affect wages and population at the same time and thus defeat the purpose of the instruments. To minimise potential problems, I consider a wide variety of possible control variables to preclude such correlations. An alternative is to use other instruments. In the spirit of Combes *et al.* (2010), I use soil characteristics, namely soil erodibility and fertility,

⁵This discussion builds on Duranton (2014). See also Combes and Gobillon (2015) for further details.

which predict current population but are not obviously linked to wages in the largest settlements where agriculture only plays a minor role.

The second main identification problem associated with the estimation of equation (1) is a possible correlation between the measure of city population and unobserved individual characteristics. More productive workers may stand to gain more from working in cities. This correlation between individual effects and city population may bias the estimate of the population elasticity α in equation (1). A possible strategy is to rely on a panel of workers and estimate α by comparing the same workers across several locations (e.g., Combes *et al.*, 2010). This strategy is not available here in absence of household identifiers. An alternative solution is to control for an extensive set of individual characteristics as in Bacolod, Blum, and Strange (2009), Bacolod, Blum, and Strange (2010) or Glaeser and Resseger (2010).

Much of the recent literature has been concerned with the heterogeneity of workers in their ability to generate and benefit from agglomeration effects. I explore many of these possible heterogeneous effects of agglomeration here.

The first form of heterogeneity regards the possibility of non-linear agglomeration effects. The literature has often focused on estimating a constant elasticity of wages (or any other measure of local efficiency) with respect to city population. A constant elasticity is often predicted by simple models of agglomeration economies (Duranton and Puga, 2004) but this prediction is mainly driven by the desire of theorists to rely on simple functional forms and keep their model tractable. Agglomeration benefits that differ with city size are interesting in themselves. They are also of importance for policy. As noted by Glaeser and Gottlieb (2008) and Moretti (2011), even though agglomeration effects rely on market failures, the distribution of population across cities need not be inefficient. With a constant agglomeration elasticity, relocating workers across cities is neutral from an efficiency perspective since the agglomeration gains in the city of arrival are fully offset by agglomeration losses in the city of departure.

There may also be some heterogeneity in agglomeration effects over time. The first half of the study period starting in 1996 was a time of considerable civil unrest in many parts of Colombia. Although armed groups were mostly present in rural areas, smaller cities were also greatly affected. In the second part of the study period after 2004, levels of civil unrest were much lower. This may have had a powerful effect on Colombian cities through some sort of peace dividend. More generally, as countries develop agglomeration benefits may change as well. It may also be

the case that agglomeration benefits are altered by technological progress which may alleviate or reinforce the need to co-locate and the benefits from such co-location. I assess time changes in agglomeration effects in what follows.

Different groups of workers may also benefit from agglomeration differently. There is a long standing conjecture that cities and skills may complement each other. There is some evidence of such effects being at play in the us (Bacolod *et al.*, 2009, Glaeser and Resseger, 2010). Whether a positive interaction between city population and a worker's education is also at work in Colombia is an open question. Because labour markets for male and female workers often differ considerably, particularly in a country like Colombia, urbanization may have different effects on male and female workers and more generally affect workers in different demographic groups differently. An analysis of individual heterogeneity in agglomeration effects is proposed below.

Another question of specific interest here is the possible difference in agglomeration effects between young and old workers. There is some evidence (e.g., Peri, 2002, Wheeler, 2006) that wage growth is stronger in larger cities in the us. There is also evidence that the length of tenure in large cities is also closely related to the ability of workers to reap the benefits of agglomeration (De la Roca and Puga, 2012). These outcomes are consistent with theories that emphasise greater learning in larger cities (Glaeser, 1999). Although the lack of an available panel of workers prevents a direct treatment of the issue of urban learning, note that these theories predict higher benefits from cities for workers that have lived there longer, that is older workers. Hence the issue of urban learning can be tackled indirectly.

There may be some heterogeneity, not only in the ability of workers to benefit from agglomeration, but also in their propensity to generate such effects. Following Rauch (1993) and Moretti (2004), a significant body of literature emphasises that agglomeration benefits may be generated mostly by skilled workers. Although the literature has focused on the correlation between average education (and more specifically the share of workers with higher education) and city wages, it is still unclear whether a measure of city education should be viewed as an alternative to city population to estimate the same agglomeration effects or as a way to capture other agglomeration effects. The analysis below will hopefully shed lights on these issues.

Going back to potential confounding factors that might be correlated with city size and explain wages, three of them are of particular relevance and interest: market access, the local road infrastructure, and amenities. Starting with market access, it is possible to view it as a form of agglomer-

ation effects taking place at a broader spatial scale. In the original specification of Krugman (1980), better market access allows firms to export more and at a lower cost. Put differently, better market access implies a stronger demand for the output of local firms and this increases the value of the marginal product of their labour and can thus be expected to lead to an increase in local wages. In a refinement of this idea, Krugman and Venables (1995) consider the importance of market access in a model where they add input-output linkages across locations. Better access to outside markets makes local firms more productive by allowing them to access a broader variety of goods at a cheaper price. In turn, this makes local firms more productive and, again, this should lead to higher wages. Following Redding and Venables (2004) there has been a lot of work highlighting the importance of market access including in particular the assessment of Fally, Paillacar, and Terra (2010) for Brazilian regions or that of Bartelme (2015) for the us.

These effects both lead to an increase in the nominal wage. Better market access also increases the real wage by allowing local consumers to consume a broader variety of goods at a lower price. Handbury and Weinstein (2015) provide strong evidence that larger cities (with perhaps better market access) provide a larger variety of grocery goods at a lower price.

The immobility of labour is an important assumption imposed by Krugman (1980) or Krugman and Venables (1995) in their models and much of the follow-up work. In the context of cities within a country, we need to allow for some labour mobility. Interestingly, labour mobility leads to very different predictions regarding the effects of market access. This is because labour mobility implies the equalisation of utility, often measured by real wages, in simple models. Then, if the wage divided by the local price index is equalised across cities, cities with better market access (which enjoy a lower price index) must also offer lower nominal wages.⁶ I explore these issues further in section 5.2.

There is also a recent literature interested in the effects of urban infrastructure, following in particular Baum-Snow (2007). This literature has looked at the effect of transportation infrastructure on a wide variety of outcomes including suburbanisation, the population growth of cities, or urban export and import patterns (see Redding and Turner, 2015, for a recent survey). Although some work has been devoted to the effects of transportation on rural wages (Michaels, 2008, Donaldson and Hornbeck, 2013), very little attention has been devoted to urban wages. In light of

⁶Although in absence of labour mobility, better market access raises wages as argued above, labour mobility, by increasing labour supply reduces wages. With free labour mobility, utility equalisation must thus imply lower nominal wages.

the discussion above, the expected effect of the road infrastructure (or urban wages) is ambiguous. More roads in a city may improve market access and thus have a negative effect on wages. The local transportation infrastructure may also act as an argument in the city production function and have a positive effect on wages. Finally, roads may affect the attractiveness of cities and may thus indirectly affect wages at the spatial equilibrium.

Related to this last point, we should expect any amenity to be potentially reflected in wages through, again, a spatial equilibrium mechanism (Rosen, 1979, Roback, 1982). Any city with more amenities will, at the spatial equilibrium, be a more expensive place to live and offer lower wages. Higher housing prices and lower wages are the two main margins that we expect to be at work to equate utility across cities (Glaeser and Gottlieb, 2009).

4. Results about city population

4.1 Baseline results

Table 2 reports the main OLS results regarding the elasticity of wages with respect to city population. In absence of any other control, the specification in column 1 reports an elasticity of 0.11 (or 11%). This indicates that in the cross-section of Colombian cities between 1996 and 2012, a 10% larger population is associated with 1.1% higher wages. Column 2 adds individual characteristics as control variables. This divides the coefficient on log city population by two at 0.054. This implies that about half the relationship between wages and city population is explained by worker characteristics. A prominent feature of the data is the greater representation of more educated workers in larger cities. Regressing the share of workers with postsecondary education on city population in 2010 yields a highly significant coefficient of 0.05 with a R^2 of 47%.⁷

While the rest of this section is devoted to asserting the robustness of this finding and examining the possible heterogeneity that underlies it, a few comments are in order at this stage. First, a city

⁷This coefficient of 0.05 may seem small but only small differences are needed here given the large effects of individual education on wages. For instance, workers with 18 years of education have an hourly wage more than six times as high relative to workers with only six years of education. Small differences in the education composition of the workforce across cities are enough to explain away about half the raw elasticity of wages with respect to city population of 0.11. Behrens, Duranton, and Robert-Nicoud (2014) estimate similar relationships for US cities. Using aggregate data, they estimate an elasticity of earnings with respect to city population of 0.082 without conditioning for the share of population with a post-secondary education degree and 0.051 after conditioning for this. They also estimate an elasticity of the share of educated workers with respect to city population of about 0.068. Hence the drop in the elasticity of city earnings with respect to population drops by a lower amount in the US when controlling for city population despite a slightly faster increase in the number of college educated workers as city size increases. This is to be expected given the much high college premium in Colombia and the fact that Behrens *et al.* (2014) use aggregate data.

Table 2: Agglomeration effects, OLS specifications for all workers

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Only pop.	Indiv.	Geog.	Access	Muni. pop.	MA	Non-lin.	Comp.
log population	0.11 ^a (0.0051)	0.054 ^a (0.0045)	0.060 ^a (0.013)	0.055 ^a (0.0062)	0.060 ^a (0.0053)	0.046 ^a (0.0070)	0.071 (0.045)	0.054 ^a (0.0046)
log population ²							-0.00072 (0.0020)	
log market access				-0.13 ^a (0.037)				
log market access ²				0.0068 ^a (0.0022)				
Educ. × log pop.								-0.0086 ^c (0.0045)
Observations	2,738,494	2,735,518	2,712,456	1,873,803	2,735,518	2,735,518	2,735,518	2,735,518
R ²	0.10	0.40	0.41	0.41	0.40	0.36	0.40	0.40
Municipalities	645	645	618	337	645	572	645	645

Notes: OLS regressions with year indicators in all columns. Robust standard errors clustered at the municipality level in parentheses. *a, b, c*: significant at 1%, 5%, 10%. The dependent variable is a log wage in all columns. Individual characteristics included in columns 2 to 8 are: gender, 26 education indicators, and imputed experience and its square. Geographic characteristics in columns 3 include three regional indicators (Andean, Carribean, and Oriental), the log of the urban area, two climate indicators (hot and moderate), a series of indicators variables for natural resources (coal, gold, and oil), the log of natural water flows, a precipitation index, and the log of altitude and its square. For a given municipality, market access is measured by multiplying 2005 population by average municipal wage for all trading partners with population above 40,000, dividing by Google map time distance, and summing. The log of market access and its square are included in column 4. Its exact computation is explained in section 5.2. Column 5 replicates column 2 using total municipal population. Column 6 uses metropolitan areas instead of municipalities.

population elasticity of wages of 5.4% is quantitatively important. Comparing a small municipality with an urban population of, say, 5,000 to Bogotá with a population above seven million, this elasticity implies a wage difference of nearly 50%. While city population is not the sole factor accounting for spatial wage disparities in Colombia, it is certainly a prime driver.

Second, this elasticity of 5.4% is somewhat larger than comparable findings for cities in more developed countries. Using a similar specification, Glaeser and Resseger (2010) report an elasticity of 3.4% for US cities while, Combes *et al.* (2010) report an elasticity of 3.3% for French cities using a more data-intensive approach. On the other hand, Combes *et al.* (2015) for Chinese cities and Chauvin *et al.* (2013) for Indian cities report much higher elasticities of 10 to 20%. Third, the drop in the coefficient on city population when controlling for individual characteristics is not unusual. As in most countries where this type of approach has been implemented, workforce composition

effects account for a sizeable fraction of spatial wage disparities (Combes and Gobillon, 2015).

Column 3 further includes a number of geographic variables including regional indicators, land area, climate and precipitation, indicators for natural resources, water flows, and altitude. The coefficient on city population for this specification is marginally higher. Note also that adding geographic controls does not increase the explanatory power of the regression by much. Although detailed results are not reported here, the coefficients on most geographic controls are not statistically significant. Wages are higher in the Andean and Oriental regions of Colombia relative to the Caribbean and Pacific regions. In this regression, climate and the presence of oil also have a significant coefficient but this relationship is not robust.

Relative to the specification of column 2, column 4 adds log market access and its square. In this regression, a theory-consistent measure of market access is computed as in Duranton (2015a). Interestingly the coefficient on log market access is negative. The coefficient on the squared market access term is positive but no municipality has a high enough market access for the overall effect to be positive. For now, we only note that introducing market access does not affect the coefficient on city population. The effect of market access on wages is subject to greater scrutiny below.

Column 5 duplicates column 2 but uses total municipal population instead of the population of the urban part of a municipality. The coefficient on log population increases marginally to 0.06. Column 6 duplicates column 2 again but uses metropolitan areas instead of municipalities. The coefficient on city population drops slightly to 0.046. This drop is small. It is caused by the fact that for small areas, the municipality and the 'metropolitan area' coincide whereas large metropolitan areas are formed of several municipalities. As a result, the dispersion of the dependent variable, population, increases. At the same time, the dispersion of the main explanatory variable, mean wages, is marginally reduced. This mechanically reduces the coefficient on population. Reassuringly the effect remains small. Interestingly, the R^2 of the regression also declines somewhat in this specification. This is perhaps a weak indication that it is better to work with municipalities than with metropolitan areas in the case of Colombia.

4.2 Heterogeneity in agglomeration effects

To start investigating possible heterogeneities in agglomeration effects, column 7 of table 2 attempts to detect non-linearities by adding the square of log city population as dependent variable to the specification of column 2. The point estimate for the coefficient on log city population is

marginally higher at 0.07 but insignificant while the coefficient on the quadratic term is insignificant and close to zero. In a supplementary appendix, I confirm this absence of non-linearities in agglomeration effects in Colombia by duplicating table 2 using only individuals living in municipalities with more than 20,000 inhabitants in 1993. Despite a much lower number of cities, the coefficients on log city population remain very close to those reported in table 2.

Finally, column 8 of table 2 adds an interaction term to the specification of column 2: the product of the worker's number of years of schooling by log city population. The coefficient on this interaction term is small, negative, and marginally significant. Adding further controls often makes the coefficient on the interaction term insignificant. This negative or insignificant interaction effect is nonetheless in contrast with results from extant literature, often highlighting the existence of higher returns to cities for more educated workers (see Wheeler, 2001, Bacolod *et al.*, 2009, Glaeser and Resseger, 2010, among others). A plausible explanation for this finding that relies on higher agglomeration effects in the informal sector is proposed below.⁸

To pursue this investigation on possible heterogeneous effects of agglomeration on different types of workers, I duplicate table 2 for prime-aged males only in a supplementary appendix. The coefficients on log city population are similar in both tables in that they never differ by more than 0.005 and are never statistically different. In the same supplementary appendix, I also perform the same exercise for prime-aged females. There are very little differences in the benefits from agglomeration accruing to men and women and there is no tendency for one gender to experience systematically higher coefficients. Hence, despite considerable differences in both the supply and demand for labour between males and females, both genders appear to benefit equally from agglomeration effects.

Turning next to the comparison of workers of different age groups, I duplicate table 2 for male workers aged 20 to 30 and aged 40 to 50. The differences are again small. If anything, younger workers appear to benefit more from agglomeration with a coefficient that is between 0.005 and 0.02 higher. This difference is sometimes statistically significant at 5%.

This result is somewhat surprising and at odds with a lot of extant literature. It is often assumed

⁸Another possible reason for this result is that in smaller cities more educated workers are often employed in the public sector where the urban agglomeration premium might be less. The GEH makes it hard to isolate public-sector workers and explore this conjecture further.

that workers learn more in larger cities.⁹ The learning-in-city hypothesis typically implies that the urban premium should be larger for older workers.¹⁰ This is not the case in Colombia.

Table 3: Agglomeration effects for workers with a written labour contract, OLS specifications

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Only pop.	Indiv.	Geog.	Access	Muni. pop.	MA	Non-lin.	Comp.
log population	0.073 ^a (0.0056)	0.037 ^a (0.0071)	0.0071 (0.016)	0.047 ^a (0.0076)	0.041 ^a (0.0081)	0.040 ^a (0.0083)	-0.049 (0.042)	0.036 ^a (0.0073)
log population ²							0.0036 ^b (0.0018)	
log market access				-0.12 ^a (0.035)				
log market access ²				0.0063 ^a (0.0020)				
Educ. × log pop.								0.0054 (0.0055)
Observations	742,387	742,310	735,237	536,314	742,310	742,310	742,310	742,310
R ²	0.02	0.51	0.52	0.52	0.51	0.51	0.51	0.51
Municipalities	568	568	544	304	568	497	568	568

Notes: This table duplicates table 2 restricting the sample to workers with a formal labour contract.

As it turns out both the absence of complementarity between city size and education and the stronger agglomeration effects for younger workers may have a common cause: higher returns to agglomeration in the informal sector where workers are younger and less educated. Table 3 duplicates table 2 for a sample of workers who state that they have a written labour contract. These workers are much more likely to work in the formal sector than those who do not have a written labour contract. The first noteworthy result of table 3 is that the coefficient on city population is always below its corresponding coefficient in table 2. This difference is statistically significant in a majority of cases. The other results are essentially the same except that the coefficient on the interaction between individual education and city population is now positive (but insignificant).

⁹See Glaeser (1999) or Duranton and Puga (2004). Indirect evidence of this has been proposed by Glaeser and Maré (2001) and Combes, Duranton, Gobillon, and Roux (2012) among others. They show that workers retain some benefits from agglomeration after they leave their city. In a slightly different spirit, Peri (2002) and Wheeler (2006), among others, shows that wage growth is stronger for younger workers in larger cities. In the most sophisticated treatment of the issue to-date, De la Roca and Puga (2012) show that workers benefit more from agglomeration as their tenure in a city gets longer.

¹⁰Unless older urban workers massively move to rural areas and are replaced by old workers from rural areas, in which case older workers in cities would be less skilled. This situation is deeply implausible.

In results not reported here, I also confirmed that measured agglomeration effects are larger for workers without a labour contract.

Smaller agglomeration effects in the formal sector provide a nice resolution to some of the more puzzling results obtained so far. This explanation is also consistent with some results obtained below about human capital externalities. We can also note that in table 3, the elasticity of wages with respect to city population is closer to the same elasticity estimated in more developed countries when only workers in formal employment are considered.

A worry with this explanation is that having a written labour contract is positively correlated with the wage. Having a labour contract may also be positively correlated with city size. It is weak but there is such a correlation in the data. As a result, workers with a written labour contract may not be comparable across cities of different sizes. More specifically, workers with a labour contract may be less skilled in larger cities and this may not be fully accounted for by our individual controls. While it is difficult to fully dispel this concern in a supplementary appendix, I duplicate again table 2 focusing this time on workers that are either ‘employees’ or ‘manual workers’. This selection excludes domestic employees and self-employed. For that alternate sample, the estimated elasticity of wages to city population is again lower than in table 2, and, if anything, even lower than in table 3. This is suggestive that the smaller agglomeration effects tend to be due to the nature of the occupation and the type of activity more than just the result of the formal sector casting a wider net in large cities.

Although they point at smaller agglomeration effects, the results of table 3 may understate the true magnitude of the decline when focusing only on the formal sector. The reason is that larger cities tend to be more unequal (Glaeser, Resseger, and Tobio, 2009, Behrens and Robert-Nicoud, 2014). A conclusion often reached by prior work is that more skilled workers benefit more from being in larger cities. By focusing on formal sector workers, we also focus on more skilled workers for whom agglomeration effects should be stronger.

In a supplementary appendix, I propose another duplication of table 2 that considers only workers in the top quartile of labour income in their city. Unsurprisingly, this table shows that the measured agglomeration effects are stronger for workers that are higher up in the wage distribution. This greater wage inequality in larger cities is in line with extant results for developed countries.

Agglomeration effects might also be heterogeneous over time. Many popular discussions un-

derscore either beliefs in how agglomeration effects are becoming irrelevant in an era of cheap telecommunications or, instead, how proximity increasingly matters. While a full analysis of this issue is clearly beyond the scope of this paper, it is interesting to look for any possible time heterogeneity behind the results reported so far. A table in a supplementary appendix duplicates the specification in table 2 column 3 for eight separate time periods, namely pairs of consecutive years since 1996 except for 2012, which is considered alone in the last column. Overall the changes in the coefficient on city population are limited. The small changes that can be observed coincide exactly with changes in the data collection process of the GEIH.

4.3 Further identification of agglomeration effects

Now that the stability of agglomeration effects over time is established, it is tempting to attempt to identify the elasticity of city wages with respect to city population using city population and wage changes over time. In practice, this implies estimating equation (1) using municipality fixed effects. To rely solely on cross-sectional variation and avoid using any time variation, the results reported in table 2 and all its variants imputed the same population to all municipalities for all years using an average of 1993, 2005, and 2010 population data. To estimate equation (1) using only the time variation in the data, a different approach is needed. Urban municipal population for 1996 to 2004 can be created by extrapolation using 1993 and 2005 population data. Population for 2008, 2009, 2011, and 2012 can also be imputed using a similar approach.

Because of these imputations of the key explanatory variable, city population, estimating equation (1) with city fixed effects suffers from a major measurement problem. The other identification issue that afflicts this estimation is that variations in population are likely to be sensitive to changes in wages. More specifically, in a fixed-effect or first-difference setting, population is likely to be ‘more endogenous’ than in a pure cross-sectional estimation.

This said, adding municipal fixed effects to the baseline specification of table 2 column 2 estimates an insignificant coefficient of 0.04 for city population. (Results are not reported here.) Using only the years for which population is better measured and using overall municipal population, which is perhaps better measured than population of the urban part, implies a coefficient of 0.081 significant at 1%. Many other specifications estimate coefficients close to zero with large standard errors. After extensive experimentations, my two main conclusions are that fixed effect

specifications do not contradict cross-sectional estimations but that they do not provide enough precision to be really useful.

Table 4: Agglomeration effects, TSLS specifications

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Only pop.	Indiv.	Geog.	Access	Muni. pop.	MA	Non-lin.	Comp.
log population	0.092 ^a (0.011)	0.048 ^a (0.011)	0.047 ^a (0.029)	0.056 ^a (0.012)	0.053 ^a (0.012)	0.042 ^a (0.014)	-0.12 (0.10)	0.077 ^a (0.014)
log population ²							0.0068 ^c (0.0039)	
log market access				-0.12 ^a (0.044)				
log market access ²				0.0061 ^b (0.0026)				
Educ. × log pop.								-0.71 (0.55)
Overid. p-value	0.38	0.25	0	0.050	0.25	0.33	0.39	0.27
First-stage Stat.	101	97.2	108	71.2	106	132	14.1	8.35
Obs.	2,029,310	2,026,723	2,015,610	1,444,433	2,026,723	2,063,528	2,026,723	2,026,723
Municipalities	316	316	311	165	316	264	316	316

Notes: This table duplicates table 2 using two-stage least-square (TSLS). In all columns, log population is instrumented by the corresponding log populations for 1843, 1870, 1938. The square of these instruments are also used in column 7. The log population term interacted with education in column 8 is a projection from log population in 1843, 1870, and 1938. Unlike in table 2, the set of geography controls in column 3 does not include the log of land area which makes the instruments weak since thinly populated places in the past led to the creation of large municipalities.

To alleviate worries about identification effects, I now turn to TSLS estimations and instrument contemporaneous populations with corresponding 1843, 1870, and 1938 populations. As argued above, these three instruments will obviously prevent any possible bias caused by reverse causality. The validity of the instruments rely on these three past populations being uncorrelated with unobserved drivers of contemporaneous wages conditional on the other controls. While a good case can be made for this (Combes *et al.*, 2010), it is by no means decisive. This type of iv estimation is probably best thought of as further evidence supportive of the main result and part of a series of robustness checks rather than the last word on the issue.

The first series of iv results is reported in table 4. This table replicates table 2 instrumenting contemporaneous population by the three historical populations described in the previous paragraph. In column 1, the coefficient on log city population is marginally lower than its OLS counterpart. Given the absence of individual characteristics, the informational content of this regression is

probably limited. In column 2, the TSLS elasticity of wages with respect to city population is 0.048 instead of 0.054 for its OLS counterpart. This difference is small, less than one standard deviation, and insignificant. The ancillary tests show that the instruments are strong and imply similar coefficients. Although the OLS-TSLS difference obtained with this specification is not significant, it appears again with the same sign and magnitude in most of the specifications discussed below. This is suggestive of perhaps a small upward bias in the OLS coefficients of table 2. Note also that estimating a marginally lower coefficient for city population with this sort of instrument is a standard finding for this type of regression (e.g. Combes *et al.*, 2010).

Turning to the subsequent columns of table 4, note that introducing a large number of geographic controls makes no change to the coefficient on city population in column 3. Columns 4 and 5 estimate slightly higher coefficients whereas, in keeping with OLS results, column 6 estimates a slightly lower coefficient. Column 7 estimates a negative coefficient for city population. The coefficient on squared log population is however higher than its OLS counterpart and significant. These two estimates imply a marginal agglomeration effect of 0.048 (as in column 2) for a city of 240,000. The coefficient on city population in column 8 is also close to its corresponding OLS coefficient.

The main conclusion to be drawn from table 4 is the marginally lower agglomeration elasticity with TSLS. To assess the robustness of this result further, a table in a supplementary appendix experiments with the instrumentation strategy and the estimation technique. It reports results for each instrument used separately. It also uses alternate instruments built from tourism and hospitality data (the number of air passengers and number of hotels). It also makes an attempt using soil characteristics in the spirit of Combes *et al.* (2010), although they turn out to be weak. Finally this table also experiments with alternate estimation techniques for instrumental variables such as the generalised method of moments (IV-GMM) and with limited information maximum likelihood (LIML) instead of TSLS. The important takeaway for that robustness table is that it essentially confirms the results obtained in table 4.

Overall, it appears that the instrumented coefficients on city population are generally statistically similar to our baseline OLS of 0.055 but with a marginally lower point estimate. This conclusion is generally robust to the choices of specification, instruments, and estimation technique.

5. Results about human capital externalities, market access, and infrastructure and amenities

I now extend the previous approach and focus on four sets of variables: city human capital, market access, infrastructure, and amenities. The results that follow should be viewed as providing robustness checks for our main results so far. It is also an exploration into alternative ways to measure and think about agglomeration effects.

5.1 *Human capital externalities*

Human capital externalities arise from the notion that having more education (or skills) in a city generates social benefits to urban workers over and beyond the already large private benefits of education. Put differently, the notion of human capital externalities recognises that, perhaps, not all workers contribute to agglomeration effects equally. More skilled workers may contribute more. This is in contrast with what we have done so far when regressing wages on overall city population. We have implicitly assumed that all workers contribute equally to agglomeration effects. The key distinction is thus that human capital externalities differ from the agglomeration effects measured so far as they concern the skill composition of the local workforce, not its absolute size.

Before going any further, and despite the obvious complementarity between the two concepts, it is important to recognise a key tension between the ‘agglomeration’ literature and the ‘human capital externality’ literature. The key explanatory variable of interest (city size or share/number of educated/skilled worker) of each literature is typically ignored (missing) in the other.¹¹ To reconcile these two strands of literature, in the regressions that follow I include both a measure of city scale (i.e., city population just like in previous regressions) and a measure of the share of more educated workers. This allows me to account for the skill composition of the workforce at the same time as its scale. To measure city skill/education, I use the share of workers with some university education since the importance of education at the upper end is underscored by extant results in the literature.

¹¹The pioneering efforts in the estimation of human capital externalities are due to Rauch (1993) and Moretti (2004). Their findings of strong external returns to education have been reproduced in many countries. See Duranton (2014) for a short review and discussion of the issue for developing countries.

Table 5: Agglomeration and human capital externalities, OLS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Only pop.	Indiv.	Geog.	Access	Muni. pop.	MA	Non-lin.	Comp.
log population	0.057 ^a (0.011)	0.052 ^a (0.0089)	0.055 ^a (0.013)	0.060 ^a (0.014)	0.047 ^a (0.0096)	0.035 ^a (0.010)	0.068 (0.049)	0.053 ^a (0.0089)
share educated	1.43 ^a (0.33)	0.050 (0.26)	0.47 ^a (0.12)	-0.14 (0.35)	0.30 (0.23)	0.35 (0.22)	0.036 (0.27)	0.047 (0.26)
log population ²							-0.00066 (0.0020)	
log market access				-0.13 ^a (0.042)				
log market access ²				0.0070 ^a (0.0025)				
Educ × log pop.								-0.0082 ^c (0.0043)
Observations	2,737,326	2,734,355	2,711,293	1,872,901	2,734,355	2,734,355	2,734,355	2,734,355
R ²	0.10	0.40	0.41	0.41	0.40	0.36	0.40	0.40
Municipalities	643	643	616	336	643	570	643	643

Notes: This table duplicates table 2 adding the share of population with some university education.

Table 5 duplicates table 2 adding the share of population with some university education, which can be loosely referred to as city education. Recall that individual education is controlled for so that the coefficient on city population really measures an external effect of education.¹²

Two results emerge. First, the coefficient on city population is unaffected by the introduction of this alternative measure of agglomeration. Its significance and its magnitude are the same as before. This reinforces all the conclusions drawn above. Second, the results regarding human capital externalities are mixed. The coefficient on city education is significant only in columns 1 and 3. In column 3, this coefficient is large but of the same magnitude as uncovered by Moretti (2004) in the US. In columns 5 and 6, the magnitude remains the same but the coefficient is insignificant. In other columns, the coefficient on city education is much smaller (even negative once) and insignificant.

When the same regressions as in table 5 are estimated without city population, the coefficient on

¹²There are of course serious identification worries. The coefficient on the share of educated workers will also capture complementarities between skilled and unskilled labour in the production function. In addition, more educated workers flock to cities with higher wages. As will become clear in the next paragraph, this worry is moot in the Colombian case. In any case, a full treatment and identification of human capital externalities is beyond the scope of this work. The share of university-educated labour is introduced to assess the robustness of prior finding and compare with previous work on this topic.

city education is high and highly significant. However in the horse race between city population and city education, city population emerges as a clear winner. Since both population and education are measured from the same source (the Colombian censuses of 1993 and 2005 and more recent population projections), it is hard to argue that this result is caused by a measurement issue where the effect of city education would be captured by population.

To verify that these results were not driven by any of the choices made when estimating the regressions in table 5, I made a number of checks. Maybe the share of educated workers in city population is not used with an appropriate functional form. In a supplementary appendix, I duplicate the regression of table 5 using the log share of university educated workers instead of the level. The results are the same. Maybe the explanatory variable is not appropriate. I experimented with various alternative measures of post-secondary education going from the more restrictive (university degree) to the least restrictive (any form of post-secondary education) and obtained again similar results. Using average years of schooling does not give stronger results either. Maybe using cross-sectional variation over a long period of time is not appropriate. Using only data for 2008-2010 with 2010 education data does not improve the results. Maybe using data for a large number of municipalities is not appropriate as human capital externalities may occur only in larger cities. Using municipalities with a population above 20,000 in 1993, which drastically reduces the number of municipalities, does not yield stronger results.

In a supplementary appendix, a table reports results where the share of education is instrumented by the number of local higher education institutions and the number of the positions for the local branch of the national technical university (SENA). I use a similar instruments when regressing population growth on city education in Duranton (2015c). Although the instruments are strong, the coefficients on city education is in most cases insignificant but typically around 0.50. From this, I conclude that the evidence on human capital externalities in Colombia is mixed.

While these weak results are in contrast with those obtained in more developed countries, they are perhaps consistent with the interpretation given above regarding the role of the informal sector. If agglomeration economies are in large part driven by the informal sector, this implies that they are driven by workers that are on average younger and less educated. Although there might be some human capital externalities coming from the most educated segment of the population, these two phenomena may offset each other.

5.2 Market access

As argued above, the effects of market access on wages are potentially ambiguous and need not be positive as often believed. To understand the reasons for this possible ambiguity, consider the following heuristic.¹³ Indirect utility in municipality i is given by $(w_i - H_i)/P_i$ where H_i is expenditure on housing and local transportation and P_i is the local price index for consumption goods. In this simple example, indirect utility is given by real income once the cost of local housing and transportation has been paid. Better market access can have four effects. The first two are higher demand for the output from local firms (because buyers can buy at a lower price inclusive of transportation costs) and higher productivity (because the firm can purchase intermediate inputs at a lower price). These two effects are expected to translate into higher wages. In addition, better market access also implies lower prices for consumption goods. Finally, better market access may lead to population growth and thus higher local housing and transportation costs. At a simple spatial equilibrium with labour mobility, indirect utility $(w_i - H_i)/P_i$ is equalised across municipalities. As a result, better market access can imply higher or lower wages depending on which effects dominates and what the main adjustment margin is. In particular, if better market access primarily means a lower price index locally, wages will need to adjust downwards for indirect utility to be equalised across municipalities.

Table 6 reports results for three series of regressions with three different measures of market access. Columns 1 to 3 use a theory-based measure of market access as computed in Duranton *et al.* (2014), itself a simplified version of the measure proposed by Redding and Venables (2004). Importantly, market access is computed over a municipality's trading partners. The shipment data from 2011 used to identify a municipality's trading partners were collected by verifying the content of travelling trucks along Colombia's major roads. In practice, trading partners are cities with a population above 40,000 that are in most cases not immediate neighbours. For instance, the shipment data used to compute market access does not record any trade between the capital city Bogotá and its immediate suburbs.¹⁴ As explained below, this selection on trade partners that are not neighbours may be an advantage.

¹³See Bartelme (2015) for a full-fledged model consistent with this heuristic.

¹⁴ The primary objective of the data is to collect information about traffic along major roads in Colombia, not to measure trade across places.

Table 6: Agglomeration effects, OLS specifications with market access

	'Theory' market access			Income market access			Population market access		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
log population	0.044 ^a (0.0064)	0.040 ^a (0.0087)	0.059 ^a (0.013)	0.055 ^a (0.0062)	0.060 ^a (0.014)	0.071 ^a (0.020)	0.054 ^a (0.0057)	0.052 ^a (0.0092)	0.052 ^a (0.014)
log market access	-0.079 ^b (0.031)	-0.080 ^a (0.030)	-0.016 (0.018)	-0.13 ^a (0.037)	-0.13 ^a (0.042)	-0.010 (0.025)	-0.11 (0.096)	-0.10 (0.099)	-0.0067 (0.046)
log market access ²	-0.028 ^a (0.011)	-0.027 ^b (0.011)	-0.019 ^a (0.0066)	0.0068 ^a (0.0022)	0.0070 ^a (0.0025)	0.00074 (0.0014)	0.0041 (0.0040)	0.0040 (0.0041)	-0.000030 (0.0019)
share educ.	N	Y	Y	N	Y	Y	N	Y	Y
Geography	N	N	Y	N	N	Y	N	N	Y
Observations	2,608,300	2,607,137	2,586,549	1,873,803	1,872,901	1,870,884	2,573,455	2,573,455	2,559,814
R ²	0.40	0.40	0.41	0.41	0.41	0.41	0.40	0.40	0.41
Municipalities	630	628	603	337	336	329	470	470	457

Notes: OLS regressions with year indicators and individual characteristics in all columns. Robust standard errors clustered at the municipality level in parentheses. *a*, *b*, *c*: significant at 1%, 5%, 10%. The dependent variable is a log wage in all columns. Geographic characteristics are the same as in table 2. Market access in columns 1 to 3 is computed as in Duranton, Morrow, and Turner (2014) using the fixed effects estimated from a trade gravity regression. Market access in columns 4 to 6 is computed for municipality *j* as $\sum_j Pop_j \bar{w}_j / d_{ij}$ where the *j* are all the trade partners of municipality *j* identified in the shipment data used in Duranton (2015a), *Pop_j* is population of municipality *j*, \bar{w}_j is the wage of municipality *j* net of individual characteristics estimated as above, and *d_{ij}* is road distance between *i* and *j*. Market access in columns 7 to 9 is computed for municipality *j* as $\sum_j Pop_j / d_{ij}$.

Column 1 adds this variable to the baseline specification of column 2 in table 2. Column 2 also adds the log share of educated workers. Column 3 further includes geographic controls and market access. The same patterns is repeated in columns 4 to 6 with a measure of market access based on the aggregate incomes of a municipality's trade partners. Unlike the measure of market access used in columns 1 to 3, this one is constructed directly from local incomes but it nonetheless uses the same sample of city pairs. Its main drawback is that it does not account for local prices unlike the theory-driven measure of market access estimated from trade data used in columns 1 to 3. For this reason, it should be viewed as more *ad hoc*. It is sometimes referred to as market potential since it only measures the size of nearby markets. Finally, columns 7 to 9 use an even simpler measure of market potential where only population is considered.

In columns 1 and 2, the coefficient on market access and on its square are both negative and significant. In column 3, both coefficients remain negative but only that on squared market access is significant. The more *ad hoc* market potential variable used in columns 4 to 6 that weights population by a measure of income has negative and significant coefficients in two of the three specifications. The positive coefficient on the squared term does not make the marginal effect of

market access positive over the range of observed market access. The simpler market access variable used in columns 7 to 9 is always insignificant. Note that this is due to larger standard errors. The point estimates are generally consistent with those obtained with the other two measures of market access.

In a table in a supplementary appendix, I compute market potential for population, local labour incomes, or local GDP using all other municipalities and not only trade partners. These measures of market potential are arguably even more *ad hoc*. They yield a coefficient on market access that is insignificant and poorly estimated. I also experimented with a variety of other specifications and obtained similar outcomes.

Another alternative is to compute the market access of a municipality including the municipality itself in the computation. This typically yields a positive effect of market access on wages. This is for instance what Fally *et al.* (2010) find when they assess the importance of market access for Brazilian regions.¹⁵ Their approach shares some similarities with what is done here. In a first regression, they compute a local wage conditioning out the effect of individual characteristics very much like here. In a second regression, they regress bilateral trade between Brazilian jurisdictions on distance and fixed effects for importers and exporters. Although details differ, this parallels what we use here except that we do not consider a municipality's own market access.

Although subject to some caveats and of limited robustness, the market access results of table 6 are in stark contrast with results in extant literature. This coexistence of negative, insignificant, and positive results may seem contradictory. Perhaps it is not. To understand this, it is important to note that the theory-consistent measure of market access used above is negatively correlated with city population with a correlation of about -0.3. This is in part explained by the fact that cities that export a lot to other cities in Colombia tend to be fairly isolated cities that export either natural resources or agricultural commodities. This negative correlation is also consistent with many theoretical results in the new economic geography literature (Fujita, Krugman, and Venables, 1999) where market access is typically a bell-shaped function of distance to the main population centre. The more *ad hoc* measures of external access used in table 6 and in the supplementary appendix, which have little to no effect on wages, have a positive correlation with city population between 0.1 and 0.3.

Finally, the measures of market access used in the literature, which rely on a municipality's own

¹⁵See also recent work by Bartelme (2015).

population or income, are more highly correlated with city population with a correlation typically above 0.5. As a result, these measures of market access also positively correlate with wages.

Separating between what one might call external market access (i.e., market access computed on other municipality) and ‘internal’ market access (i.e., a municipality’s own population) is important for two reasons. The first is that internal market access is essentially the aggregate income of a municipality and introducing it in a wage regression raises an obvious endogeneity problem. Second, internal market access will capture input-output linkages within a municipality, which are part of the agglomeration economies we try to identify. We also expect internal market access to be highly correlated with other benefits from agglomeration that do not percolate through input-output linkages but are instead associated with thick local labour markets and learning spillovers. Although the literature that attempts to assess the effects of market access on wages usually fails to separate between different possible benefits from agglomeration, it makes sense to separate between internal and external market access and acknowledge that internal market access will capture various benefits from agglomeration.

We can now go back to the simple equilibrium condition stated above: $(w_i - H_i)/P_i = \text{constant}$. Then, it is clear from the results obtained above that a larger local population raises local wages. Larger cities are also more expensive places to live, in part due to higher housing and transportation costs. In the context of this simple spatial equilibrium, equilibration is probably occurring through the adjustment of urban costs when it comes to ‘internal’ market access (i.e., city population). On the other hand, a higher ‘external’ market access may predominantly affect the consumption side (i.e., by lowering the price index P) and may thus lead to lower nominal wages especially if the population response is weak and does not lead to higher urban costs.

5.3 *Transportation infrastructure*

Related to market access, the local transportation infrastructure may also matter. In a related paper (Duranton, 2015a), I show that a better provision of roads or a better accessibility to the main road network of the country has a sizeable effect on both the volume and the value of exports of Colombian municipalities. Whether roads also have an effect on wages is, of course, another question.

Table 7 reports results for two series of regressions using different measures of roads. Columns

Table 7: Agglomeration effects and roads, OLS specifications

	Road index 1			Road index 2		
	(1)	(2)	(3)	(4)	(5)	(6)
log population	0.049 ^a (0.012)	0.050 ^a (0.014)	0.043 ^a (0.0054)	0.049 ^a (0.012)	0.050 ^a (0.014)	0.043 ^a (0.0054)
Roads	0.0061 (0.016)	0.0058 (0.016)	0.0067 (0.0075)	0.0061 (0.016)	0.0058 (0.016)	0.0067 (0.0075)
share educ.	N	Y	Y	N	Y	Y
Geography+access	N	N	Y	N	N	Y
Observations	2,399,834	2,399,834	1,724,221	2,399,834	2,399,834	1,724,221
R ²	0.40	0.40	0.41	0.40	0.40	0.41
Municipalities	243	243	130	243	243	130

Notes: OLS regressions with year indicators and individual characteristics in all columns. Robust standard errors clustered at the municipality level in parentheses. *a, b, c*: significant at 1%, 5%, 10%. The dependent variable is log wage in all columns. Geographic characteristics are the same as in table 2. Market access is as in table 2. The road index in columns 1 to 3 sums the log of the number of major national roads in the municipality, the log of the number of exits in and out of the municipality, and the log of kilometers of national roads minus the log of kilometers of distance to the closest national road for municipalities with no roads. Logs are normalised to zero in absence of roads within a municipality. The road index in columns 4 to 6 only sums the first terms of road index 1.

1 to 3 use an index summing the number of national roads, their number of entries into the municipality, the mileage and, for municipalities with no national roads, the (negative) distance to the closest national road. Columns 4 to 6 use a similar index that does not consider the distance to the closest road. The first result in table 7 is that the estimated elasticity of wages with respect to city population is not affected by the inclusion of road variable. The second main finding in table 7 is that the road variable is insignificant. Because of the way the two road indices used in table 7 are constructed, the roads coefficients can be loosely interpreted as elasticities. They are always less than 1%. Importantly, the standard errors around these coefficients are also extremely small. Consequently, we can rule out OLS elasticities as low as 3%. I experimented with these findings using the measures of roads that compose these indices one at a time and found similar results.

Following extant literature on the effects of roads on urban outcomes (Baum-Snow, 2007, Duranton and Turner, 2012), the worry is that road allocation may be endogenous. When assessing the effects of roads on the population growth of us cities, Duranton and Turner (2012) provide evidence of such endogeneity as their instrumental variable results indicate much stronger effects of roads than are estimated directly from the cross-section. Here, I follow the identification strategy used in Duranton and Turner (2012) and instrument the road indices used above by the

measures of the 1938 road network and colonial roads ('caminos reales'). Because population is also potentially endogenous and correlated with current roads, I also instrument population by the same population lags as in table 4. These instruments will rule out any form of reverse causality. However, old roads may also be correlated with contemporaneous characteristics of Colombian municipalities that determine their wages. For this reason, I also estimate iv regressions where the share of university educated workers and geographical characteristics are also included.

Table 8: Agglomeration effects and roads, LIML specifications

	Road index 1			Road index 2		
	(1)	(2)	(3)	(4)	(5)	(6)
log population	0.023 (0.021)	0.0061 (0.031)	-0.010 (0.026)	0.028 (0.020)	0.011 (0.029)	-0.0095 (0.026)
Roads	0.042 (0.031)	0.043 (0.034)	0.0025 (0.0069)	0.027 (0.029)	0.031 (0.031)	0.0054 (0.0061)
log share educ.	N	Y	Y	N	Y	Y
Geography+access	N	N	Y	N	N	Y
Overid. p-value	0.37	0.21	0	0.42	0.21	0
First-stage Stat.	3.49	2.31	34.5	5.66	4.86	216
Observations	1,852,423	1,852,423	1,389,553	1,852,423	1,852,423	1,389,553
Municipalities	137	137	73	137	137	73

Notes: This table duplicates table 8 using limited information maximum likelihood LIML. In all columns, log population is instrumented by the corresponding log populations for 1843, 1870, 1938. The road index is instrumented by the corresponding indices computed from the 1938 road network and from the colonisation roads ('caminos reales').

The results are reported in table 8. Unfortunately, the instruments are weak and much caution is needed when interpreting these results. It can nonetheless be noted that some of the specifications estimate a coefficient on the road index between 0.03 and 0.05. On the other hand, the coefficient on city population becomes insignificant. Interestingly, the sum of the coefficients on the road index and city population in a majority of cases remains the same as before at about 0.05-0.06. This suggests that the regressions cannot identify the effects of roads on wages separately from those of population when using instrumental variables. This is due to the fact that past populations and past road networks predict current roads better than current population. If we think that the true elasticity of wages with respect to city population is around 0.05, that only leaves a tiny positive effect of roads on wages.¹⁶

¹⁶In results not reported here, I also confirm the tiny effect of roads on wages in regression where I instrument for roads but not for city population.

Remember also that whatever effect of roads on wages is measured, this is an overall effect. As shown in Duranton (2015a) roads in a municipality foster exports. In itself, this should have a positive effect on wages. Roads may also lead to higher wages because they improve local productivity by facilitating local trades. At the same time, roads may be thought of as improving market access. In line with the results of the previous section, a better access, which lowers the price of imported goods, may lead to lower wages if labour is mobile across places. Hence documenting a zero or very small coefficient is not synonymous with the absence of an effect. Rather, roads may have countervailing effects on wages.¹⁷

5.4 Amenities

Models of spatial equilibrium involving utility equalization across locations usually imply that amenities will be capitalised into property prices through a higher demand for land and through wages. To live in locations with better amenities, workers are willing to accept a combination of higher rents and lower wages. The channel through which amenities capitalise into property prices is straightforward. Since housing is in limited supply in any location, a greater demand for housing will lead to higher prices. The mechanism through which amenities capitalise into wages is indirect. Firms that face higher land prices use less land, which, in turn, reduces the marginal product of labour and thus the wage (Roback, 1982, Glaeser and Gottlieb, 2009).

There is a large literature using data from developed countries that convincingly argues that climate is an important amenity (Rappaport, 2007, Cheshire and Magrini, 2006). There are major differences in climate between Colombian cities, caused mainly by differences in their elevation. This said, from the regressions estimated in section 4 the effect of climate variables on wages in Colombia is weak and it is unclear whether the climate affects wages indirectly through the channel just described or directly as a factor of production.

Table 9 performs three series of regressions with three types of 'human' (i.e., non-natural) amenity variables. Columns 1 to 3 consider leisure amenities: libraries, museums, and hotels, and restaurants. Column 1 adds these three variables to the baseline specification of column 2 in table 2. Column 2 also adds the log share of educated workers. Column 3 further includes geographic controls and market access. The same patterns is repeated in columns 4 to 6 with the log number

¹⁷In Duranton (2015c), I find that roads have a positive effect on city growth.

Table 9: Agglomeration effects and amenities, OLS specifications

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
log population	0.067 ^a (0.012)	0.075 ^a (0.016)	0.12 (0.074)	0.043 ^a (0.012)	0.043 ^a (0.016)	0.13 ^a (0.042)	0.053 ^a (0.0062)	0.054 ^a (0.0098)	0.066 ^a (0.018)
log libraries p.c.	0.058 ^b (0.025)	0.061 ^b (0.025)	-0.038 (0.031)						
log museum p.c.	-0.0026 (0.027)	0.0015 (0.027)	0.0076 (0.012)						
log hospitality p.c.	-0.047 ^c (0.027)	-0.045 (0.028)	-0.0042 (0.041)						
log homicide p.c.				0.012 (0.025)	0.012 (0.034)	-0.015 (0.014)			
log violence 1							0.086 ^b (0.036)	0.087 ^b (0.038)	-0.019 (0.021)
log violence 2							-0.053 (0.035)	-0.054 (0.037)	0.020 (0.017)
share educ.	N	Y	Y	N	Y	Y	N	Y	Y
Geography+access	N	N	Y	N	N	Y	N	N	Y
Observations	2,326,863	2,326,863	1,637,316	2,340,482	2,340,482	1,639,045	2,735,289	2,734,126	1,870,884
R ²	0.40	0.40	0.41	0.40	0.40	0.41	0.40	0.40	0.41
Municipalities	102	102	39	118	118	44	644	642	329

Notes: OLS regressions with year indicators and individual characteristics in all columns. Robust standard errors clustered at the municipality level in parentheses. *a, b, c:* significant at 1%, 5%, 10%. The dependent variable is log wage in all columns. Geographic characteristics are the same as in table 2 and also include log market access. Library per capita is computed as the 1+ the number of libraries in 2009 divided by 2005 population. One is added to accommodate the log transformation. Museum per capita is computed as the 1+ the number of museums in 1998 + the number of museums in 2009 divided by 2005 population. Hospitality per capita is computed as the 1+ the number of restaurants in 2005 + 20 × the number of hotels in 2005 divided by 2005 population. Homicide per capita is computed as the 1+ the number of homicides in 2002 + the number of homicides in 2009 divided by 2005 population. Violence 1 is computed as 1+ the sum over 1993 to 2005 of all operations by the army and any of the unofficial armed groups. Violence 2 is 1+ the sum of all deaths attributed to unofficial armed groups over 1993-2005.

of homicide per capita and in columns 7 to 9 with two variables associated with political violence between 1993 and 2005: the number of military operations plus attacks by unofficial armed groups and the number of death attributed to armed groups.

Among the three leisure activities, the coefficient for the number of libraries is only significant with a positive sign in columns 1 and 2, whereas hospitality has a negative significant coefficient only in column 1. The sign of the coefficient on the number of libraries may be viewed as ‘perverse’ since we expect good amenities to be associated with lower wages at the spatial equilibrium. However, this is probably due to reverse causation where richer municipalities can afford more libraries. In columns 4 to 6, the number of homicides is never significant. Finally, the first political

violence variable has a positive coefficient that is significant at 5% in columns 7 and 8 but becomes insignificant in column 9. The other political violence variable is always insignificant.

Overall, the patterns are weak since none of the amenity variables considered here (nor in experiments not reported here) appears to have a robust effect. A first possible explanation for these weak patterns is that the 'human' amenities considered in table 9 are highly correlated with the geographic controls, some of which may be viewed as natural amenities. Then, collinearity may prevent the appropriate identification of the effects of the human amenities considered here. This is not the case. The correlations between the amenities considered in table 9 and the geographic controls are low. Most of them are below 0.1 and the highest ones are only about 0.25. Interestingly, it is the inclusion of regional indicators that makes amenities insignificant when we control for geography in table 9. One would expect the effect of amenities to be robust to their inclusion.

On the other hand, it is true that the amenities considered in table 9 are strongly correlated with log population. We note nonetheless that the standard errors on the amenity coefficients are small. This argues against a collinearity problem. It is better to think about these results on amenities as being zeroes that are fairly precisely estimated rather than as having a complete lack of pattern. It is also noteworthy that this precision is obtained despite the amenity variables used in columns 1 to 6 being available for only a small subset of municipalities.

A worry might be then be that population is driven by these amenities. This conjecture is examined in Duranton (2015c) where I examine the determinants of city growth in Colombia between 1993 and 2010. I only find a modest role for amenities. That amenities play no discernible role in urban growth is consistent with the fact that Colombian households spend only a minimal fraction of their income on leisure (Meisel Roca and Pérez, 2012).

There remains of course the worry that these coefficients are biased by reverse causality. While this remains a possibility, it would require two things. First, the effect of all positive amenities should be overestimated (i.e., be zero instead of negative) because they are caused by greater local incomes. Second, the effect of negative amenities should be underestimated because they are caused by local poverty. It is hard to imagine opposite biases affecting positive and negative amenities with just this exact pattern.

To go deeper on this, it is possible that amenities affect more educated individuals more either because they are more mobile or because they value these amenities more, positively or negatively. Duplicating table 9 by adding some interactions between amenities and education does not change

anything and the coefficients on the interaction terms are erratic in their sign and lack precision. On the other hand, the coefficient on log city size remains positive, significant, and of the same magnitude as previously in all columns except column 3 where significance is lacking with only 39 municipalities (and essentially 17 explanatory variables at the municipality level).

6. Conclusions

This paper provides a broad exploration of agglomeration benefits in Colombian cities. The main finding is an elasticity of wages with respect to population of about 5%. This coefficient is extremely robust to the inclusion of a wide variety of controls and to the use of now standard instrumental variable strategies. This elasticity suggests sizeable gains from 'urbanisation'. Moving from a city with 10,000 inhabitants to Bogotá with more than 7 million is associated with about 40% higher wages. This elasticity of about 5% is somewhat larger than the same elasticity estimated across us and French cities but smaller than for Chinese and Indian cities.

Consistent with this cross-country comparison, I also find stronger agglomeration effects for workers in the informal sector. As these workers tend to be younger and less educated, this provides a unified explanation for a range of negative findings including weak evidence regarding human capital externalities, a lack of measurable learning effects, and an absence of complementarity between city size and individual education.

I also find suggestive and novel evidence of a negative effect of market access on wages. This effect is probably best understood in a spatial equilibrium framework with labour mobility. The evidence about roads suggests no effect of roads on wages (or tiny positive effects at best). There is also no effect of urban amenities on wages.

These findings should be useful for two reasons. First, they provide a rather comprehensive documentation of agglomeration effects for Colombian cities. This should be helpful at a time when the country rediscovers its cities and is starting to rethink the organisation of its territory after having considered it for many years mainly through military lenses.

This broad set of findings, positive and negative, should also be useful to develop a broader knowledge base about agglomeration effects in developing countries. There is now a lot of consistent evidence about agglomeration economies in more advanced countries. The evidence for developing countries still lags behind despite their crucial importance at a time when many of these countries are fast urbanising.

References

- Angrist, Joshua D. and Adriana D. Kugler. 2008. Rural windfall or a new resource curse? Coca, income, and civil conflict in Colombia. *Review of Economics and Statistics* 90(2):191–215.
- Attanasio, Orazio, Pinelopi K. Goldberg, and Nina Pavcnik. 2004. Trade reforms and wage inequality in Colombia. *Journal of Development Economics* 74(2):331–366.
- Bacolod, Marigee, Bernardo S. Blum, and William C. Strange. 2009. Skills in the city. *Journal of Urban Economics* 65(2):136–153.
- Bacolod, Marigee, Bernardo S. Blum, and William C. Strange. 2010. Elements of skill: Traits, intelligences, education, and agglomeration. *Journal of Regional Science* 50(1):245–280.
- Barco, Carolina. 2013. Misión para el fortalecimiento del sistema de ciudades de Colombia. Documento técnico de soporte política del sistema de ciudades de Colombia. Report for the Colombian National Planning Ministry (Departamento Nacional de Planeación, Dirección de Desarrollo Urbano).
- Bartelme, Dominick. 2015. Trade costs and economic geography: Evidence from the U.S. Processed, University of California Berkeley.
- Baum-Snow, Nathaniel. 2007. Did highways cause suburbanization? *Quarterly Journal of Economics* 122(2):775–805.
- Behrens, Kristian, Gilles Duranton, and Frédéric Robert-Nicoud. 2014. Productive cities: Sorting, selection, and agglomeration. *Journal of Political Economy* 122(3):507–553.
- Behrens, Kristian and Frédéric Robert-Nicoud. 2014. Survival of the fittest in cities: Urbanisation and inequality. *Economic Journal* 124(581):1371–1400.
- Bushnell, David. 1993. *The Making of Modern Colombia: A Nation in Spite of Itself*. Berkeley (CA): University of California Press.
- Chauvin, Juan Pablo, Edward L. Glaeser, and Kristina Tobio. 2013. Urban economics in the US and India. Manuscript in progress, Harvard University.
- Cheshire, Paul C. and Stefano Magrini. 2006. Population growth in European cities: weather matters-but only nationally. *Regional Studies* 40(1):23–37.
- Ciccone, Antonio and Robert E. Hall. 1996. Productivity and the density of economic activity. *American Economic Review* 86(1):54–70.
- Combes, Pierre-Philippe, Sylvie Démurger, and Li Shi. 2015. Migration externalities in China. *European Economic Review* 76(0):152–167.
- Combes, Pierre-Philippe, Gilles Duranton, Laurent Gobillon, and Sébastien Roux. 2010. Estimating agglomeration economies with history, geography, and worker effects. In Edward L. Glaeser (ed.) *The Economics of Agglomeration*. Cambridge (MA): National Bureau of Economic Research, 15–65.
- Combes, Pierre-Philippe, Gilles Duranton, Laurent Gobillon, and Sébastien Roux. 2012. Sorting and local wage and skill distributions in France. *Regional Science and Urban Economics* 42(6):913–930.
- Combes, Pierre-Philippe and Laurent Gobillon. 2015. The empirics of agglomeration economies. In Gilles Duranton, Vernon Henderson, and William C. Strange (eds.) *Handbook of Regional and Urban Economics*, volume 5A. Amsterdam: Elsevier, 247–348.

- De la Roca, Jorge and Diego Puga. 2012. The dynamic earnings premium of dense cities. Processed, CEMFI and IMDEA Social Sciences.
- Departamento Administrativo Nacional de Estadística. 2013. *Great Integrated Household Survey (GEIH) Methodology*. Bogota, Colombia: DANE.
- Donaldson, Dave and Richard Hornbeck. 2013. Railroads and American economic growth: A “market access” approach. Processed, Harvard University and MIT.
- Duranton, Gilles. 2014. Growing through cities in developing countries. *World Bank Research Observer* 39(1):39–73.
- Duranton, Gilles. 2015a. Roads and trade in Colombia. *Economics of Transportation* 4:forthcoming.
- Duranton, Gilles. 2015b. A proposal to delineate metropolitan areas in Colombia. *Economia & Desarrollo* forthcoming.
- Duranton, Gilles. 2015c. Determinants of city growth in Colombia. Mimeographed, Wharton School, University of Pennsylvania.
- Duranton, Gilles, Peter M. Morrow, and Matthew A. Turner. 2014. Roads and trade: Evidence from the us. *Review of Economic Studies* 81(2):681–724.
- Duranton, Gilles and Diego Puga. 2004. Micro-foundations of urban agglomeration economies. In J. Vernon Henderson and Jacques-François Thisse (eds.) *Handbook of Regional and Urban Economics*, volume 4. Amsterdam: North-Holland, 2063–2117.
- Duranton, Gilles and Matthew A. Turner. 2012. Urban growth and transportation. *Review of Economic Studies* 79(4):1407–1440.
- Fally, Thibault, Rodrigo Paillacar, and Cristina Terra. 2010. Economic geography and wages in Brazil: Evidence from micro-data. *Journal of Development Economics* 91(1):155–168.
- Fujita, Masahisa, Paul R. Krugman, and Anthony J. Venables. 1999. *The Spatial Economy: Cities, Regions, and International Trade*. Cambridge (MA): MIT Press.
- Glaeser, Edward L. 1999. Learning in cities. *Journal of Urban Economics* 46(2):254–277.
- Glaeser, Edward L. and Joshua D. Gottlieb. 2008. The economics of place-making policies. *Brookings Papers on Economic Activity* (1):155–253.
- Glaeser, Edward L. and Joshua D. Gottlieb. 2009. The wealth of cities: Agglomeration economies and spatial equilibrium in the United States. *Journal of Economic Literature* 47(4):983–1028.
- Glaeser, Edward L. and David C. Maré. 2001. Cities and skills. *Journal of Labor Economics* 19(2):316–342.
- Glaeser, Edward L., Matt Resseger, and Kristina Tobio. 2009. Inequality in cities. *Journal of Regional Science* 49(4):617–646.
- Glaeser, Edward L. and Matthew R. Resseger. 2010. The complementarity between cities and skills. *Journal of Regional Science* 50(1):221–244.
- Greenstone, Michael, Richard Hornbeck, and Enrico Moretti. 2010. Identifying agglomeration spillovers: Evidence from winners and losers of large plants openings. *Journal of Political Economy* 118(3):536–598.

- Handbury, Jessie and David E. Weinstein. 2015. Goods prices and availability in cities. *Review of Economic Studies* 82(1):258–296.
- Henderson, J. Vernon. 2005. Urbanization and growth. In Philippe Aghion and Steven N. Durlauf (eds.) *Handbook of Economic Growth*, volume 1B. Amsterdam: North-Holland, 1543–1591.
- Jofre-Monseny, Jordi, Raquel Marín-López, and Elisabet Viladecans-Marsal. 2014. The determinants of localization and urbanization economies: Evidence from the location of new firms in Spain. *Journal of Regional Science* 54(2):313–337.
- Krugman, Paul. 1980. Scale economies, product differentiation, and the patterns of trade. *American Economic Review* 70(5):950–959.
- Krugman, Paul R. and Anthony J. Venables. 1995. Globalization and the inequality of nations. *Quarterly Journal of Economics* 110(4):857–880.
- Marshall, Alfred. 1890. *Principles of Economics*. London: Macmillan.
- Meisel Roca, Adolfo and Gerson Javier Pérez. 2012. Las ciudades colombianas y sus atractivos. Report for the Urban Development Division of the National Planning Ministry (DNP-DDU).
- Michaels, Guy. 2008. The effect of trade on the demand for skill - Evidence from the Interstate Highway System. *Review of Economics and Statistics* 90(4):683–701.
- Moretti, Enrico. 2004. Estimating the social return to higher education: Evidence from longitudinal and repeated cross-sectional data. *Journal of Econometrics* 121(1):175–212.
- Moretti, Enrico. 2011. Local labor markets. In Orley Ashenfelter and David Card (eds.) *Handbook of Labor Economics*, volume 4B. Amsterdam: North-Holland, 1237–1313.
- Overman, Henry G. and Anthony J. Venables. 2005. Cities in the developing world. Report, Department for International Development (UK).
- Peri, Giovanni. 2002. Young workers, learning, and agglomerations. *Journal of Urban Economics* 52(3):582–607.
- Puga, Diego. 2010. The magnitude and causes of agglomeration economies. *Journal of Regional Science* 50(1):203–219.
- Rappaport, Jordan. 2007. Moving to nice weather. *Regional Science and Urban Economics* 37(3):375–398.
- Rauch, James E. 1993. Productivity gains from geographic concentration of human capital: Evidence from the cities. *Journal of Urban Economics* 34(3):380–400.
- Redding, Stephen and Anthony J. Venables. 2004. Economic geography and international inequality. *Journal of International Economics* 62(1):63–82.
- Redding, Stephen J. and Matthew A. Turner. 2015. Transportation costs and the spatial organization of economic activity. In Gilles Duranton, Vernon Henderson, and William C. Strange (eds.) *Handbook of Regional and Urban Economics*, volume 5B. Amsterdam: Elsevier, 1339–1398.
- Roback, Jennifer. 1982. Wages, rents and the quality of life. *Journal of Political Economy* 90(6):1257–1278.

- Rosen, Sherwin. 1979. Wage-based indexes of urban quality of life. In Peter Mieszkowski and Mahlon Straszheim (eds.) *Current Issues in Urban Economics*. Baltimore: Johns Hopkins University Press, 74–104.
- Samad, Taimur, Nancy Lozano-Gracia, and Alexandra Panman. 2012. *Colombia Urbanization Review. Amplifying the Gains from the Urban Transition*. Washington (DC): World Bank.
- Smith, Adam. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations*. London: Printed for W. Strahan, and T. Cadell.
- Sánchez, Fabio and Jairo Nuñez. 2000. Geography and economic development: A municipal approach for Colombia. *Revista de Planeación y Desarrollo* 31(3/4):141–156.
- Wheeler, Christopher H. 2001. Search, sorting, and urban agglomeration. *Journal of Labor Economics* 19(4):879–899.
- Wheeler, Christopher H. 2006. Cities and the growth of wages among young workers: Evidence from the NLSY. *Journal of Urban Economics* 60(2):162–184.