

Delineating metropolitan areas: Measuring spatial labour market networks through commuting patterns

Gilles Duranton[‡]
University of Pennsylvania

April 2013

ABSTRACT: This chapter first discusses the necessity of defining metropolitan areas and current practice in several countries. It argues for the use of a simple algorithm that exploits cross-municipality commuting patterns. Municipalities are aggregated iteratively provided they send a share of their commuters above a given threshold to the rest of a metropolitan area. This algorithm is implemented on Colombian data and its robustness is assessed. Finally, the properties of the resulting spatial labour market networks are explored.

Key words: metropolitan area definition, municipal aggregation, Colombian cities, Zipf's law

JEL classification: R11, R12, R14

[‡]Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104, USA (e-mail: duranton@wharton.upenn.edu; website: <https://real-estate.wharton.upenn.edu/profile/21470/>). Also affiliated with the Centre for Economic Policy Research, the Rimini Centre for Economic Analysis, and the Spatial Economic Centre at the London School of Economics. I am grateful to Rafael Cubillos for giving me data without which this project would have been impossible. James Bernard, Matthew Degagne, and Hongmou Zhang provided very able research assistance. I also thank Paul Cheshire and Yoshi Kanemoto for very useful conversations and for getting me interested in this subject many years ago. Finally, feedback from Alvaro Pachon, Rafael Cubillos, José Salazar and other seminar participants at Departamento Nacional de Planeación is gratefully acknowledged.

1. Introduction

This chapter proposes a methodology to define metropolitan areas by iterative aggregation of spatial units using daily commuting flows between them. In essence, a spatial unit A is aggregated to another spatial unit B if the share of the workers who work in B among all those that reside in A is above a given threshold. Another spatial unit C may next be aggregated to the union of A and B if, similarly, it sends a fraction of its commuters greater than the same threshold to this newly formed unit even though it may not have been possible to aggregate C directly to either A or B. This process of aggregation repeats until no further unit can be aggregated.

This algorithm is implemented with a threshold of 10% of commuters on municipal data in Colombia to define metropolitan areas for this country, which currently lacks well-defined metropolitan areas. Although aggregating spatial units iteratively using a minimum commuting threshold is not novel, our implementation is novel in two respects. First, we show that a careful implementation of an aggregation algorithm that relies solely on a minimum commuting threshold criterion is enough to define meaningful metropolitan areas and generate metropolitan cores endogenously. This is unlike the practice of many statistical institutes which usually predefine metropolitan cores and use a minimum commuting rule in conjunction with several other criteria. Second, we assess the robustness of the set of resulting metropolitan areas to changes in the minimum commuting threshold for aggregation.

Defining metropolitan areas is important for several reasons. Historically, as cities grew both in population and spatially, they would directly annex surrounding municipalities. In many countries, this process has stopped ; richer municipalities resist fiscal integration with their poorer neighbours ; mayors attempt to retain their jobs ; or, as in Colombia, there may be significant constitutional and administrative barriers to merging municipalities. As a result, administratively defined cities are typically restricted to an urban core and are no longer representative of their broader metropolitan environment.

Related to this, existing administrative units such as municipalities do not generally constitute functionally autonomous units. Instead neighbouring municipalities are often

economically integrated in all sorts of ways. This implies that an economic shock or a policy intervention in one municipality may have important spillover effects on its neighbours. Given the difficulty of keeping track of spillover effects, it is easier (and typically more efficient) for the policy to target functionally consistent units.

Being able to deal with functionally consistent units is also important for research. For instance, cities tend to grow geographically by spreading outwards, outside the boundaries of the core municipality. When looking at patterns of urban growth based on municipal data, one may conclude that large cities grow slowly. This is often far from being the case. A core municipality is often ‘full’ and its metropolitan area typically grows at its extensive margin via its peripheral municipalities. Hence, urban growth is most appropriately measured at the metropolitan level.

Finally, cities constitute interesting spatial networks of commuting workers, transacting firms, or interacting individuals. To be able to study these networks meaningfully it is fundamental to first be able to describe them.

The rest of this chapter is organised as follows. Section 2 provides some background about the situation in Colombia, current practice in other countries, and prior academic literature. Section 3 presents the data and our aggregation algorithm. Section 4 provides our list of metropolitan areas and metropolitan regions for Colombia. The robustness of these results is assessed in section 5. Finally, section 6 concludes.

2. Background, current practice, and literature

2.1 The current situation in Colombia

Although there exists an official set of ‘metropolitan areas’ in Colombia, these areas are mostly administrative units, constituted on a voluntary basis (Senado de la República, 2012). While there is certainly a strong case for associations of neighbouring municipalities to form broader formal institutions, these ‘political’ metropolitan areas are usually not

appropriate for analysis and decision-making by higher levels of governments.¹

For historical reasons and, perhaps, because of institutional rivalry, Bogotá, the largest municipality in Colombia, is not part of any metropolitan area even though there is no observable discontinuity between Bogotá and, for instance, its Southern neighbour Soacha. Cali, the third largest city in Colombia, is not part of any institutional arrangement with any of its neighbours either. Barranquilla is a less extreme case. Its official metropolitan area is composed of only five municipalities whereas we obtain a metropolitan area of already nine municipalities with an extremely conservative commuting threshold of 30% (which is three times as large as our preferred threshold of 10%). On the other hand, Medellín, the second largest city, has formed the ‘metropolitan area of the Aburrá Valley’ which corresponds exactly to the one generated by our algorithm with our preferred commuting threshold of 10%. However, Medellín is the exception, not the rule. A systematic and consistent set of metropolitan areas is needed for Colombia.

2.2 Current practice in the world

While details vary, there are two features that are common to most ordinarily used definitions of metropolitan areas.

The first is the preponderant role given to commuting patterns. Metropolitan areas are thus viewed as integrated labour markets. There are good reasons for this. Since Marshall (1890), economists usually think of cities as bringing benefits, in terms of ‘thick’ labour markets, greater diversity of available final and intermediate goods, and more intense individual interactions conducive to knowledge spillovers. Focusing on the first series of these benefits coming from local labour markets makes sense for two reasons. The first is that commuting patterns can easily be tracked. The census and many other sources of labour market data usually record both the place of residence and the place of work of workers. The variety of final and intermediate goods, input-output linkages, and

¹The French government defines ‘statistical’ metropolitan areas through its statistical institute (INSEE). At the same time, there are many ‘urban communities’ which are voluntary unions of neighbouring municipalities, i.e., political metropolitan areas. The two coexist to serve extremely different purposes.

knowledge spillovers are much more complicated to track (Holmes, 1999, Charlot and Duranton, 2004, Handbury and Weinstein, 2010). There is also a broad consensus among economists interested in cities that commuting patterns usually take place over distances that we naturally recognise as being 'metropolitan'. Instead, knowledge spillovers might take place over much shorter distances while input-output links often take place at a scale broader than the metropolitan area, as argued for instance by Krugman (1991).

In addition, there are other criteria that could be used to define metropolitan areas including non-economic criteria such as the sense of belonging to a place, etc. In practice, however, because they are easier to track and because their scale seems right, commuting patterns play an overwhelming role in the definition of metropolitan areas.

The second key feature of most official definitions of metropolitan areas is the use of an iterative approach to aggregate municipalities (or other basic geographical units such as counties in the us) into metropolitan areas. More specifically, a minimal threshold of commuters is chosen. As soon as the share of commuting flows from an origin municipality to a destination municipality is above this threshold, the origin municipality is aggregated to the destination municipality. We will refer to the aggregated municipality as a 'satellite' municipality and the one it is aggregated to as its 'core'. These two municipalities become part of the same metropolitan area. This procedure is then repeated until there remains no municipality to aggregate.

If employment in metropolitan areas were fully centralised at a unique central business district, there would be no need to use an iterative approach. However, in reality only a small proportion of jobs is concentrated in the centre of metropolitan areas. Glaeser and Kahn (2001) argue that less than 10% of employment in us cities is concentrated within 5 kilometres of their centre. This is far from the idealised description of monocentric cities where all the jobs are located in a central business district (Alonso, 1964, Muth, 1969, Mills, 1967). As a result, and given the gravitational nature of commuting where the number of commutes decreases with distance an iterative aggregation procedure is needed. Imagine a core municipality A, a 'first-ring' municipality B, and a 'second-ring' municipality C.

Municipality C may be sending lots of commuters to A and B but not enough to warrant immediate aggregation to A. As a result, B may be aggregated to A at the first round. Then C will be aggregated to the union of A and B at the second round.

Note that commuting thresholds are defined relative to the number of workers in the municipality at hand. This is because municipalities differ vastly in terms of their resident labour force. Using a relative threshold is important because it allows the aggregation of a small municipality that sends all its residents to the core municipality. Using an absolute threshold would not allow for this. Worse, on Colombian data it would lead to very misleading outcomes since there are many ‘commuters’ (in absolute terms) between the largest cities, including for instance the pair composed of Bogotá and Barranquilla, which are located several hundred kilometres from each other. Looking at absolute numbers of commuters is an interesting measure of ‘links’ between municipalities and could perhaps be instrumental in the circulation of knowledge. This, however, does not help aggregate nearby municipalities into metropolitan areas.

Aside from these two features which are used by most countries that define metropolitan areas, there are several others which are common to many countries.

The first of these features (which is used for instance in the us) is the pre-determination of a ‘core’. That is, the authority in charge of defining metropolitan areas will aggregate satellite units (counties in the American case) only around particular ‘core’ units which satisfy *ex ante* some particular properties in terms of population size and density. Put differently, a city needs to be ‘big enough’ and ‘dense enough’ to be considered as a potential nucleus for a metropolitan area. For instance, in the us, the core county must “(a) Have at least 50 percent of [its] population in urban areas of at least 10,000 population; or (b) Have within [its] boundaries a population of at least 5,000 located in a single urban area of at least 10,000 population. ”(us Office of Management and Budget, 2010).

While this type of criterion seems intuitive, our results for the Colombian case show that it is not needed in practice. First, pre-defined cores might be arbitrary. Instead the algorithm that defines metropolitan areas should also pick the cores endogenously.

Then, given the absence of *ex ante* cores, issues surrounding the minimum criteria that a core should satisfy become moot which is desirable. As will become clear below, it is best to avoid criteria that are either un-necessary or that can be manipulated to define metropolitan areas.

It is possible to think of some mostly rural municipalities that would attract a significant fraction of commuters from other larger municipalities. These rural municipalities would then be perversely tagged as ‘metropolitan cores’. One could also imagine large groups of rural municipalities with lots of cross-commuting giving rise to ‘metropolitan areas’. They would obviously be missing urban character. While such pathological situations are theoretically possible in the absence of pre-defined cores, the Colombian example shows that in all cases aggregation into metropolitan areas occurs around the largest municipality and there are very few cases of aggregation involving municipal cores with a small population. As argued below, these areas can always be selected out *ex post* by imposing a minimum population size for metropolitan areas.

Geographical contiguity may also be added as a criterion to define metropolitan areas. This seems natural. A highly integrated area is expected to be geographically continuous (also sometimes referred to as coterminous). While there might be esthetic reasons for imposing geographical continuity, there is no strong economic reason. Two municipalities separated by inhospitable terrain may form one economically integrated area and the area in-between may remain mostly rural. It is not clear why this area in-between should be forcibly integrated when it is not interacting with the other two municipalities. In any case, this is again a moot point because the algorithm used below to define metropolitan areas only aggregates contiguous areas with our preferred threshold of 10%. Again the gravitational nature of commuting implies that a municipality completely surrounded by a metropolitan area is unlikely to remain alone when all of its neighbours have been aggregated. In any case, rather than impose a contiguity constraint *ex ante* it is better to check for exceptions *ex post* and attempt to understand them.

Statistical authorities also sometimes add further criteria, including asking for ‘local

opinions' in the us. A related issue is whether the algorithm used to define metropolitan areas should be applied in a strict fashion or instead used more 'flexibly'. Conceptually, these two questions are separate. One may want to use a complicated algorithm to define metropolitan areas and apply it in a strict manner. Alternatively, it is possible to think of a simple algorithm subject to some 'operational adjustments' *ex post*. In practice, the issues of the number of criteria in the algorithm used to define metropolitan areas and whether this algorithm is applied flexibly or not are deeply intertwined. The use of many criteria (including fairly subjective ones that rely on local opinions) is probably a way to have some flexibility in the delineation of metropolitan areas. To make things worse, countries that use many criteria do not make their exact algorithm, the inputs into it, and its output public.

There are two reasons why one should use a simple and transparent algorithm that is applied strictly to define metropolitan areas. The first is that it really makes no sense to develop a methodology that defines metropolitan areas if it is to be renegotiated *ex post* because of a statistician's whims or because of political pressure. The second reason is that metropolitan areas are part of economic policies. Hence, the delineation of metropolitan areas affects the allocation of resources. It then becomes easy to see how and why the definition of metropolitan areas can become politicised. Policies that allocate resources to metropolitan areas whose definitions have been meddled with are biased. This means outcomes that are poorly measured, less efficient, and potentially unfair. To avoid political interference, it is crucial that the definition of metropolitan areas remains as simple as possible and that the task of defining them be given to an independent statistical institute. The advantages of doing so are overwhelming relative to the possibility of having one or two 'awkward' cases in the final list of metropolitan areas.

Statistical institutes also sometimes impose an *ex post* minimum population size criterion for metropolitan areas. This may not be needed if, for instance, the definition of metropolitan areas already imposes some minimum population constraint for the core municipality. For policy purpose, it is obvious that a minimum size threshold often

needs to be considered. This threshold may depend on the type of policy. Looking at the provision of university education for which the metropolitan area is arguably the relevant spatial scale, it is clear that a relatively high population threshold needs to be considered. We cannot expect 'metropolitan areas' with a few thousand inhabitants to be provided with universities. When looking at environmental issues such as the disposal of solid residuals, it is probably best to consider all metropolitan areas including small 'lone' municipalities. It is also the case that imposing a stringent threshold on the entire list of metropolitan areas suppresses useful information. As a result it is usually best to generate a complete list of municipalities and metropolitan areas. A cutoff can then be imposed for a particular analysis or a specific policy or set of policies. This has the added benefit of allowing for more relevant cutoffs to be chosen and forcing the policy makers to justify their threshold in a clear and transparent manner.

Another interesting feature of the definition of metropolitan areas in many countries is the fact that there are often several such definitions. For instance, France defines both 'urban areas' and 'urban units'. The latter are typically organised around a single core whereas the former are more standard (and broadly defined) metropolitan areas. The same situation is encountered in the US where there is a list of 'consolidated' metropolitan areas and a list of 'primary' metropolitan areas. Consolidated metropolitan areas are the union of several adjacent primary metropolitan areas. To be more concrete, Washington DC and Baltimore form two separate metropolitan areas but they also belong to the same consolidated metropolitan area. Again, like in France, the primary metropolitan areas appear to be core-based and correspond to integrated labour markets. Instead, consolidated metropolitan areas capture broader spatial units, and perhaps other forms of economic integration. Baltimore and Washington DC are certainly part of the same 'economic region' even though the proportion of workers that commute to DC from the Northern suburbs of Baltimore is probably quite low. There is a clear tradeoff here. Having multiple definitions will allow policy makers and analysts to capture different dimensions of economically integrated areas. At the same time, a multiplicity of definitions opens the door to arbitrary

decisions and political interference. There is also the issue of how to go on about several definitions and whether they should be based on different thresholds for commuting or defined by different principles. While we return to these issues in the Colombian case below, it is hard to deny that having two different definitions for two different spatial scales is attractive.

We draw the following conclusions from the preceding discussion. The case for defining metropolitan areas based on commuting flows and for using an iterative procedure is extremely strong. The case for having two definitions to capture two different scales is also strong. On the other hand, the justification for many other practices routinely used by statistical institutes appears weak. Defining ‘cores’ ex-ante seems unnecessary, prevents useful checks on the algorithm, and opens the door to political interference. The same arguments apply with respect to the use of other (i.e., non-commuting) criteria to define metropolitan areas. Finally, using a simple and transparent algorithm that can be replicated (or used by others) opens the door to a number of useful checks that can be done. The usual practice of statistical institutes of proposing ‘a list’ of metropolitan areas without the raw data and details of their algorithm is clearly unsatisfactory.

2.3 Prior literature

The necessity to define urban areas first became clear in the US during the 1950s. Strong urban expansion and suburbanisation was no longer accompanied by municipal annexation. This led to a divergence between the political boundaries of the urban core and the economic boundaries of metropolitan areas. To resolve this problem, the US bureau of the census defined its Standard Metropolitan Statistical Areas (SMSAs) in the early 1950s. Early discussions in Berry (1960) and Fox and Kumar (1965) were very much focused on defining metropolitan areas within a central place theory framework. Later Berry, Loble, Goheen, and Goldstein (1969) offered a remarkable early discussion that echoes many of the points made here and suggested relying solely on commuting patterns towards a predetermined urban core to define metropolitan areas. Following the US, other developed

countries also started defining their own metropolitan areas without seemingly much academic input. Their choices came under scrutiny in Hall and Hay (1980) and Cheshire and Hay (1989) who attempted to develop a broader perspective on European cities and needed a consistent set of units.

More recently, Kanemoto and Kurima (2004) have proposed an algorithm for Japan that has been widely used by subsequent research in absence of an official definition for metropolitan areas for this country. There is also a small stream of research that assesses how a range of local economic outcomes spatially autocorrelates across small spatial units to aggregate them into larger ones (see Cörvers, Hensen, and Bongaerts, 2009, for a recent example). Finally, geographers often propose lists of metropolitan areas but the definitions they propose are usually ad-hoc (see for instance Molina, 2001, for Colombia).

We also note that extant research sometimes defines its own zoning (Briant, Combes, and Lafourcade, 2010, Rozenfeld, Rybski, Gabaix, and Maske, 2011). The delineations currently used by researchers differ a lot. Using different zonings for policy purposes may be an issue because it is well known that the zoning that one adopts may drive some of the results.² At the same time, as already argued, there is nothing intrinsically wrong with using different zonings for different purposes since different problems may require a focus on different spatial scales. There is also a strand of literature (e.g. Duranton and Overman, 2005) that attempts to measure economic phenomena in continuous space doing away with spatial units altogether. This is not an option here given our perspective.³

3. A simple aggregation algorithm

In agreement with the argument above, our proposed algorithm is as simple as possible. It aggregates a spatial unit to another if the former sends a high enough fraction of its commuters to the latter. Subsequently, a third spatial unit is aggregated to the union of the first two provided it sends a high enough fraction of its commuters to this newly

²See for instance the well known ‘Modifiable Areal Unit Problem’ (MAUP). See Cressie (1993) for a presentation and a discussion.

³For instance, it is obvious that policies that allocate money to ‘places’ need discrete spatial units.

formed unit even though it may not have been possible to aggregate this third spatial unit to any of the first two individually. This process is repeated until no spatial unit remains to be aggregated.

3.1 Preliminary issues

Before going deeper into the details of the algorithm and its implementation to the Colombian case, it is useful to discuss the choice of commuting threshold. This choice is fundamentally arbitrary. Theory offers no reliable guidance here because economic integration between places follows a continuum. But since the objective is to delineate discrete units, there is no way around the necessity of a threshold. Taking a high threshold leads to the aggregation of very few satellite municipalities to urban cores, whereas taking a low threshold will lead to extremely large metropolitan areas. At the extreme, if each municipality sends at least one commuter to each of its neighbours, taking an arbitrarily low threshold will imply only one metropolitan area that covers the entire country. This is not helpful.

Adding to this, the choice of threshold is likely to depend on the size of the underlying units to be aggregated. Colombian municipalities are fairly large (on average more than 100 square kilometres). The gravitational nature of commuting implies that large municipalities will send on average only a small proportion of their commuters to work in other municipalities. Instead, France has more than 35,000 municipalities (and their average land area is only about 15 square kilometres). We thus expect much higher commuting flows between French municipalities because of this. Unsurprisingly, the threshold used by the French statistical institute is high at 40%.

Commuting distances also depend on the level of development. In developed countries where a large fraction of workers can commute by car or using well-developed public transportation systems, a large proportion of workers may be able to commute over long distances. In Colombia, where car ownership is still limited and public transportation underdeveloped, the fraction of commuters that can commute over long distance is much

lower than in Europe or North America. Hence, one may want to use different thresholds in developed and developing countries. This said, we also need to keep in mind that it is desirable to retain some consistency for the definition of metropolitan areas as a country develops.

A related problem associated with the choice of commuting threshold is the sensitivity of the delineation of metropolitan areas to small changes in the threshold. This can occur because of the iterative nature of the algorithm. Think of the following hypothetical example. Municipality D sends 12% of its workers to municipality C and 10% to B. Municipality C sends 12% of its workers to B and 10% to A. Finally, municipality B send 19% of its workers to A. With a commuting threshold of 20%, all four municipalities remain isolated since there is no flow above this threshold. For a threshold below 19%, B gets aggregated to A at the first round. Then C which sends $10\%+12\%=22\%$ to the union of A and B get aggregated at the second round. At the third round, D also get aggregated so that we end up with a metropolitan area made of all four municipalities. In this example, a small change of threshold from 20% to below 19% leads to a radically different zoning.

To the possibility of such perverse cases suggested by this example, there are two responses. The first is to choose a ‘natural’ threshold (typically a round number) to avoid any suspicion of interference. The second response is to assess the sensitivity of the delineation of metropolitan areas with respect to the choice of threshold by comparing outcomes for different values of the commuting threshold. We perform such robustness checks below.

3.2 Data

To define metropolitan areas for Colombia, we use the matrix of commuting flows from the 2005 Colombian census and 2010 population estimates from the Colombian statistical institute, DANE. This choice reflects two conflicting constraints: using consistent data (preferably from the same year) and using the most recent data. For population, the entire population of each municipality was considered. For each municipality, Colombian

statistics typically distinguish between an urban (or ‘head’) part and a rural part. Taking the entire population has the obvious drawback of aggregating rural populations to metropolitan areas. This drawback is minor in practice since most of the population of the municipalities that form large metropolitan areas is overwhelmingly ‘urban’. Discarding rural populations would also lead to some awkward choices to be made about how to compute commuting shares since data for commuting flows are only available for entire municipalities.

Census populations or population estimates based on censuses are the best available population data in most countries including Colombia. Commuting flows are measured from only a subsample of the population surveyed by the Colombian census.

This follows common practice in most countries where commuting questions (together with lots of other questions) are usually administered through the ‘long forms’ of the census given only to a fraction of the population for cost reasons. In our case, this suggests some minor imprecisions due to mismeasured commuting flows. The lack of precision becomes more important as lower commuting thresholds are considered since with a low threshold of say 1%, we may be well below the statistical margin of error in smaller municipalities. Results for low threshold are reported below but some care is needed in their interpretation given this reliability issue.

To delineate metropolitan areas for Colombia, we propose a commuting threshold of 10% which appears reasonable given that Colombian municipalities are fairly large.

3.3 Algorithm

The algorithm is available upon request. It was programmed in Stata. After cleaning up the original matrix of cross-municipality commuting flows and creating a number of working files, each loop of aggregation works as follows. Among all pairs of origin and destination municipalities the algorithm flags those for which the share of commuters from the origin is above the chosen commuting threshold. Before being aggregated to a destination, the algorithm verifies that in case a municipality could be aggregated to

several destinations, it is uniquely aggregated to the one it sends the most workers to. In case commuting flows between two municipalities are above the threshold in both directions, the algorithm also makes sure that the smallest municipality is aggregated to the largest.

At the aggregation stage, the name of the origin municipality is appended behind the name of the destination municipality and populations are added. The matrix of commuting flows is also appropriately aggregated and redefined. For instance, if municipality C sends 8% of its workers to municipality B and 9% to municipality A and if B is appended to A, then the commuting flows from C to B and C to A are aggregated into a unique flow of 17% from municipality C into the metropolitan area A+B. The process of selection of commuting flows and aggregation is then repeated until no municipality or group of municipalities remains to be aggregated to a metropolitan area.

As final output, the algorithm produces a list of metropolitan areas with its component municipalities (a 'core' and its 'satellites') and single municipalities. For verification purposes, the algorithm also keeps track of all origin municipalities which were aggregated and the destination municipalities they were aggregated to.

This algorithm generates a list of metropolitan areas and municipalities associated with a given commuting threshold.

We also propose to define broader units, urban regions. As argued above this is in-keeping with existing practice in many countries. Recall that, for instance, the us metropolitan areas of Washington DC and Baltimore are separate but they are also part of the same 'consolidated' metropolitan area. To define these broader urban regions, a natural approach would be to use the same principle as with metropolitan areas but adopt a lower commuting threshold. For these urban regions, we take a lower threshold of 5% but note that this change alone does not lead to dramatically larger units and clearly falls short of the notion of 'broad urban region'. The natural temptation would then be to lower this threshold even further. This is not a good idea since, as argued above, the aggregation exercise becomes fragile with very low commuting thresholds.

There is a deeper reason why even extremely low aggregation thresholds do not lead to broad urban regions. This is due to the self-reinforcing nature of the iterative aggregation process used to delineate metropolitan areas. To understand this subtle point, it is best to take a concrete example from Colombia. The 'coffee region' of Colombia is a confined high plateau between two branches of the Andes. It has three major cities which are fairly close to each other. The municipality of Pereira has around 450,000 people, that of Manizales is slightly below 400,000, and that of Armenia is slightly below 300,000. As small neighbouring satellite municipalities get aggregated to these three core municipalities, the three metropolitan areas that they form get more 'entrenched'. The municipalities that are in-between these three main cities may see a fair amount of cross commuting. But, as aggregation proceeds, these 'in-between' municipalities get aggregated to one of the three cores together with more peripheral municipalities. Given the gravitational nature of commuting, the aggregation of these peripheral municipalities lowers the tendency for their metropolitan areas to commute with each other. As a result, these metropolitan areas do not merge into a large single urban region even for a commuting threshold as low as 1%.⁴ However, it is interesting to observe that in many cases we obtain metropolitan areas that are contiguous with each other. Hence to define metropolitan regions, we propose to aggregate metropolitan areas that are contiguous with each other with a commuting threshold of 5%. Then, the three separate areas aggregated around Pereira, Manizales, and Armenia, which are contiguous, are also the main centres of the larger urban region of Pereira-Manizales-Armenia.

⁴This phenomenon is not unique to the coffee region. The same is observed in the region of the Caribbean coast where three of the main cities: Barranquilla, Cartagena and Santa Marta do not merge even for a low commuting threshold of 1%.

4. Results

4.1 Metropolitan areas

For the preferred commuting threshold of 10%, the list of the 45 resulting metropolitan areas with more than 100,000 inhabitants is provided in table 1. There are another 39 metropolitan areas with population above 50,000. In total, 99 satellite municipalities get aggregated to 22 cores, 19 of which have a population above 100,000. All the other municipalities remain stand alone municipalities. Metropolitan areas with a population above 100,000 are also depicted on the map of figure 1.

Before going further into the description of the list of metropolitan areas in table 1, a few important features related to the algorithm need to be discussed. First, the iterative nature of the algorithm is fundamental. With a 10% threshold, the algorithm goes through 7 rounds of aggregations before converging. In the case of the largest metropolitan area composed of Bogotá and 22 neighbouring satellite municipalities, only nine of them are added at the first round of aggregation.

It is also interesting to note that the algorithm always picks as core municipality the largest municipality of the metropolitan area. This demonstrates that defining cores *ex ante* is unnecessary. As can be verified on the map of figure 1, the metropolitan areas generated by the algorithm are also composed of contiguous municipalities. This shows that imposing contiguity is not needed. Finally, there is no set of small and rural municipalities that get aggregated into much larger ‘metropolitan’ areas. Actually, it is clear from the list given in table 1 that the aggregation of peripheral municipalities into broader metropolitan units occurs mostly for the largest municipalities.

The list of the 84 largest metropolitan areas contains 180 municipalities (of more than 1100 for the entire country). These 84 metropolitan areas host 32.1 million people, or about 71% of the population of Colombia. We note that peripheral municipalities are concentrated around the largest four cities. 55 of the 99 satellite municipalities are ag-


Table 1: Colombian metropolitan areas with population above 100,000

Core Municipality	Metropolitan Municipalities	Metropolitan Population	Core Population	Large satellite Municipalities
Bogotá DC	23	8,672,087	7,363,782	Cajicá, Chía, Funza, Mosquera, Soacha, Facatativá, Madrid, Zipaquirá
Medellín	10	3,544,703	2,343,049	Bello, Caldas, Copacabana, Envigado, Itagui, La Estrella
Cali	10	2,719,683	2,244,639	Candelaria, Jamundí, Yumbo, Florida, Pradera
Barranquilla	16	2,214,344	1,186,640	Baranoa, Malambo, Sabanalarga, Soledad
Cartagena	7	1,142,697	944,250	Arjona, Turbaco
Bucaramanga	4	1,074,929	524,112	Floridablanca, Girón, Piedecuesta
Cúcuta	4	773,659	618,310	Los Patios, Villa del Rosario
Pereira	3	717,383	457,103	Dosquebradas
Ibagué	1	526,547	526,547	
Santa Marta	1	447,857	447,857	
Villavicencio	2	441,906	431,476	
Manizales	2	439,630	388,525	
Armenia	4	430,749	288,908	Calarca
Pasto	2	416,224	411,706	
Montería	1	409,476	409,476	
Valledupar	1	403,414	403,414	
Buenaventura	1	362,625	362,625	
Neiva	1	330,487	330,487	
Rionegro	5	296,614	110,329	
Palmira	1	294,580	294,580	
Popayán	1	265,702	265,702	
Sincelejo	1	256,241	256,241	
Tuluá	2	217,189	199,244	
Riohacha	1	213,046	213,046	
Tunja	5	200,696	171,082	
Barrancabermeja	1	191,498	191,498	
San Andres de Tumaco	1	179,005	179,005	
Sogamoso	9	165,183	115,564	
Florencia	1	157,450	157,450	
Apartadó	1	153,319	153,319	
Uribia	1	144,990	144,990	
Maicao	1	141,917	141,917	
Turbo	1	139,628	139,628	
Girardot	3	139,155	101,792	
Ipiales	2	129,808	123,341	
Cartago	1	128,566	128,566	
Yopal	1	123,361	123,361	
Magangué	1	122,913	122,913	
Fusagasugá	1	121,535	121,535	
Guadalajara de Buga	1	116,105	116,105	
Quibdó	1	114,548	114,548	
Duitama	2	114,470	110,418	
Lorica	1	114,145	114,145	
Pitalito	1	113,980	113,980	
Ciénaga	1	103,066	103,066	

Sources: DANE and author's computations.

Notes: Name of satellite municipalities reported above a population of 50,000.

Figure 1: Map of Colombian metropolitan areas with population above 100,000



Sources: DANE and author's computations.

Notes: Name of metropolitan area reported above a population of 200,000.


gregated to one of the four largest Colombian municipalities. We also note that only 4 satellite municipalities are aggregated to core municipalities to form metropolitan areas with a population below 50,000 inhabitants. There is a strong rank correlation between the ranking of metropolitan areas in terms of population and the corresponding ranking of their core municipalities. For metropolitan areas with a population above 100,000, the correlation of log of population between the metropolitan area and the core municipality is 0.98. This said, there is some variation. The municipality of Medellín, the second largest in the country, has a population only 4% larger than that of the municipality of Cali, the third largest. However the metropolitan area of Medellín has a population 30% larger than that of metropolitan Cali.

Viewed differently, our aggregation into metropolitan areas corrects for the idiosyncracies of the delineation of Colombian municipalities. The municipality of Medellín is geographically relatively small whereas that of Cali is large. At one extreme, in the cases of Barranquilla or Bucaramanga, the metropolitan area has a population that is twice that of the core municipality. At the other extreme, some large municipalities like Santa Marta, Ibagué, or Villavicencio either remain isolated or only receive tiny satellite municipalities so that their metropolitan population roughly coincides with their municipal population. The near absence of satellite for these municipalities is unsurprising. Santa Marta is a coastal city in decline and residents of neighbouring municipalities will be more easily lured to work in Barranquilla which is located fairly close. Ibagué and Villavicencio are fairly large isolated cities located close to major geographical ‘ruptures’.


The four panels of figure 2 provide four magnified maps of the four most important concentrations of urban population where 16 of the largest 20 metropolitan areas are located including the largest five. These maps illustrate cases of contiguous metropolitan area such as Medellín and neighbouring Rionegro or the main cities of the coffee regions. These cases suggest that it is indeed interesting to consider a regional level of aggregation above metropolitan areas.

Overall the output generated by the algorithm appears to be highly consistent with both

Figure 2: Four regions of Colombia




Panel (a): Bogotá



Panel (b): Medellín



Panel (c): Barranquilla and the Caribbean coast



Panel (d): Cali and the coffee region

Sources: DANE and author's computations.

Notes: Core municipalities in dark blue (black) ; Satellite municipalities in light blue (grey).

Narrow boundaries between municipalities ; thick boundaries between metropolitan areas.

Metropolitan cores referenced with large fonts ; metropolitan satellite with population above 50,000 referenced with small fonts.

the underlying principles exposed above and qualitative features of the urban geography of Colombia.

4.2 Urban regions

We now turn to the delineation of broader urban regions. To define these regions we take a lower commuting threshold of 5% and aggregate the resulting metropolitan areas that are adjacent into urban regions.

The list of the 26 resulting urban regions composed of at least one metropolitan area of more 100,000 inhabitants is provided in table 2. These urban regions are also depicted on the map of figure 3, panel (a).

Several features stand out from table 2 and from the map of figure 3 (a). The most important is the emergence of several important urban regions composed of a number of metropolitan areas. The Northern part of the Caribbean coast appears as the second most important urban region of the country with more than five million inhabitants between Cartagena and Uribe at the extreme end of La Guajira region. There is also significant consolidation around Cali, Medellín, and the main cities of the Coffee region: Pereira, Manizales, and Armenia. A smaller urban region also occurs around Bucaramanga and Barrancabermeja. The urban region of Bogotá contains 12 more municipalities than the previously defined metropolitan area of Bogotá but its population of 8.9 million is only marginally larger than that of metropolitan Bogotá at 8.7 million.

The second important fact that comes out of table 2 is that, altogether, about 22 million people live in the four largest urban regions. This is just below half the population of the country.

We also note some interesting microfeatures about Colombian urban regions. Some regions like those around Bogotá or Medellín form highly compact regions. The urban regions of the main cities of the coffee region and that around Cali are less well formed and exhibit some 'holes'. These holes are even more apparent in the urban region of the Caribbean coast. We could choose to aggregate these unattached municipalities to the

Table 2: Colombian urban regions with at least one metropolitan areas with population above 100,000

Urban Region	Population	Municipalities
Bogotá Fusagasugá	8,909,613	35
Carribean coast	5,239,004	50
Medellín	3,914,585	18
Cali Buenaventura Tuluá Guadalajara de Buga	3,898,886	21
Pereira Manizales Armenia Cartago	1,848,224	14
Bucaramanga Barrancabermeja	1,309,812	7
Cúcuta	802,242	8
Ibagué	526,547	1
Montería	497,157	2
Villavicencio	441,906	2
Pasto	426,475	3
Apartadó Turbo	409,182	4
Sincelejo	386,560	5
Neiva	378,076	3
Sogamoso Duitama	300,580	14
Popayán	297,520	2
Tunja	206,336	6
San Andres de Tumaco	179,005	1
Florencia	157,450	1
Girardot	152,714	5
Ipiales	148,746	3
Yopal	123,361	1
Magangué	122,913	1
Quibdó	114,548	1
Lorica	114,145	1
Pitalito	113,980	1

Sources: DANE and author's computations.

Notes: The metropolitan areas of the Carribean coast urban region are Barranquilla, Cartagena, Santa Marta, Albania, Valledupar, Riohacha, Zona Bananera, Macaio, and Uribia. Despite its contiguity with the coffee cities of Pereira, Manizales, Armenia, and Cartago, the metropolitan area of Ibagué is left isolated because the geographical barriers that separates them is extremely important. Although the straight distance between Ibagué and Armenia is only about 50 kilometers, the crossing of the La Línea turnpike at 3,300 meters often takes several hours. The completion of the tunnel between Calarcá and Cajamarca will be an important first step towards integrating Ibagué with the coffee region.


urban region that surrounds them. That would hide some interesting evolutions. These holes reveal that these urban regions are still undergoing a process of formation. The regions around Bogotá or Medellín can be thought of as already mature urban regions organised around one dominant pole. The region around Cali is still under consolidation. The same happens to the Coffee or Caribbean urban regions which have the added complication of containing several cores of relatively even populations. We can also detect potential urban regions still under formation. For instance in the Boyacá region, Duitama and Sogamoso are already integrated. Tunja, the largest metropolitan area of the region remains isolated. These two areas will eventually be integrated, perhaps even into a much larger region with Bogotá. We can also see the basis of a future urban region around Montería on the Southern part of the Caribbean region going from Magangué to Turbo.

5. Robustness


To show the robustness of our approach, we duplicate our main analysis for a broad range of thresholds: 1%, 2%, 5%, 15%, 20%, 25%, and 30%. The two panels of figure 3 replicate the map of figure 1 for commuting thresholds of 5% and 20%. For most large Colombian cities, a higher threshold of 20% only makes minor differences. Of the largest 20 metropolitan areas with our preferred threshold of 10%, 19 are still in the top 20 with a threshold of 20% and the ordering of the top 10 is unchanged. Although the metropolitan area of Bogotá loses 15 municipalities in 23 with a higher threshold of 20%, the population remain very similar: 8.16 million instead of 8.72 million. The differences between these two rankings for the other core municipalities are even less important.

Moving to a lower threshold of 5% also makes little difference. The ordering of the largest nine cities is unchanged. The two most important changes are the disappearance of Rionegro and Palmira which ranked 19 and 20 with a threshold of 10%. Rionegro gets aggregated to its neighbour Medellín. The same happens to Palmira with its own neighbour Cali. Interestingly, there are no other changes among the largest metropolitan

Figure 3: Maps of Colombian metropolitan areas with population above 100,000



Panel (a): 5% commuting threshold



Panel (b): 20% commuting threshold

Sources: DANE and author's computations.

Table 3: Pairwise correlations for the log population of Colombian metropolitan areas

Threshold	1%	2%	5%	10%	15%	20%	25%	30%
1%	1							
2%	0.991	1						
5%	0.983	0.979	1					
10%	0.980	0.982	0.989	1				
15%	0.979	0.979	0.988	0.994	1			
20%	0.979	0.978	0.987	0.994	0.999	1		
25%	0.979	0.978	0.987	0.994	0.999	1	1	
30%	0.979	0.978	0.987	0.993	0.999	0.999	0.999	1

Note: Metropolitan areas with population above 50,000.

areas: the three main cities of the coffee region, Armenia, Manizales, and Pereira remain separate metropolitan areas despite their proximity. Similarly the three main cities of the Caribbean coast, Barranquilla, Cartagena, and Santa Marta also remain separate.⁵ These features persist even when we take an extremely low threshold of 1%.

More generally, table 3 reports log population size correlations for Colombian metropolitan areas defined according to the entire range of thresholds mentioned above. Among metropolitan areas that can be compared across thresholds (since for instance Rionegro disappears when lowering the threshold from 10 to 5%), the correlations reported in table 3 are extremely high, 0.97 or more. The correlations with our 10% reference threshold is at least 0.98. Repeating this table using population in level or population ranks instead of log yields even higher correlations.

Next, we assess how sensitive the number of municipalities in urban areas is with respect to the chosen commuting thresholds. Obviously the number of satellite municipalities is sensitive to the chosen threshold of commuting. Recall that with our reference threshold of 10%, 99 municipalities are satellites of an urban core. With higher thresholds

⁵We also start seeing satellite municipalities which are not geographically adjacent to the rest of their metropolitan areas. There are two such cases. The first is the Santanderian municipality of Sucre which get attached to Bucaramanga which is more than 200 kilometres far. Given that this municipality is not negligibly small and sends about 7% of its commuters to Bucaramanga, this corresponds to real flows, perhaps mostly students which are counted together with workers. The other case is Guacamayas, a tiny municipality at the North of the Boyacá region which gets attached to Bogotá which is nearly 400 kilometres away. Given that this case is driven by only 17 ‘commuters’, this may be a statistical glitch.

Table 4: Spearman rank correlations for the number of satellite municipalities of metropolitan areas

Threshold	1%	2%	5%	10%	15%	20%	25%	30%
1%	1							
2%	0.919	1						
5%	0.849	0.929	1					
10%	0.747	0.802	0.865	1				
15%	0.696	0.750	0.823	0.932	1			
20%	0.688	0.738	0.796	0.896	0.958	1		
25%	0.631	0.685	0.739	0.832	0.899	0.952	1	
30%	0.598	0.641	0.695	0.781	0.847	0.904	0.948	1

Note: Metropolitan areas with population above 50,000.

of 30% and 20%, this number falls to 25 and 41, respectively. With lower thresholds of 5% and 1%, the number of satellite municipalities increases to 180 and 616, respectively. With a threshold of 30%, the metropolitan area of Bogotá has only three municipalities instead of 208 with a low threshold of 1% even though population increases only by 37%.

To implicitly control for the large changes in the total number of satellite municipalities, we consider the Spearman rank correlation in the number of satellite municipalities as the threshold varies in table 4. Except for the highest thresholds for which very few metropolitan areas have satellites (only nine with a threshold of 30%), the correlations are generally high. For instance the Spearman rank correlation between our preferred 10% threshold and the two alternative thresholds of 5% and 20% are 0.86 and 0.90, respectively.

Another way to assess the robustness of our findings is to look at them through the perspective of Zipf’s law. This allows us to highlight the effect of the commuting threshold on the number of metropolitan areas. This exploration is also of independent interest because Zipf’s law is the subject of intense academic attention. See for instance Duranton and Puga (2013) for a recent review and Pérez (2008) for a recent contribution about Colombian cities.

Since Auerbach (1913), the distribution of city sizes has often been approximated with a Pareto distribution. To do this, a popular way is to rank cities in a country from the

largest to the smallest and regress log rank on log city population. Gabaix and Ibragimov (2011) highlight a possible small sample bias in the estimation of the coefficient on log city population and suggest instead using the log of the rank minus one half as the dependent variable:

$$\log(\text{Rank}-1/2) = \beta_0 - \xi \log \text{Population} + \epsilon.$$


The estimated coefficient ξ is the shape parameter of the Pareto distribution. Zipf's law (after Zipf, 1949) corresponds to the statement that $\xi = 1$. This implies that the expected size of the second largest city is half the size of that of the largest, that of the third largest is a third of that of the largest, etc.

Figure 4 provides a plot of the underlying data for Colombian municipalities, metropolitan areas defined according to our preferred commuting threshold of 10%, and metropolitan areas defined with a lower threshold of 2%.

For all Colombian municipalities in 2010, the estimated value of ξ is 0.85 suggesting a distribution more uneven than Zipf's law. We note however that this coefficient of 0.85 is mostly driven by a thin lower tail of small municipalities. It is reasonable to ignore extremely small municipalities since they are overwhelmingly rural. They are also exceptional since Colombian municipalities were designed to avoid extremely low population levels. Considering only municipalities with a population above 5,000 (or 84% of all municipalities hosting 98.7% of the population) yields a value of ξ of 1.02 and a higher R^2 of 98% instead of 92% for all municipalities. To make consistent comparisons with metropolitan areas, we can restrict our attention further to only large municipalities with a population above 50,000. In this case, the estimated value of ξ is 1.07 with a R^2 of 0.99. This value of 1.07 implies less disparities in population than implied by Zipf's law. However a relatively large standard error of 0.14 makes it impossible to reject a unit coefficient and Zipf's law.⁶

⁶First, because the dependent variable is computed directly from the explanatory variable, measurement error on the 'true' population also affects the rank and thus leads to a downward bias for the standard errors with OLS. Gabaix and Ibragimov (2011) show that the standard error on ξ is asymptotically $\sqrt{2/n} \xi$ where n is the number of observations. With our data, this implies a standard error of 0.14. The values of the standard errors for the other estimates of ξ reported here are of the same magnitude.

Figure 4: Zipf's law for Colombian metropolitan areas and municipalities



Sources: Author's computations with a minimum population threshold of 50,000.

Notes: The black triangles represent municipalities and the dotted line is the associated regression line (slope -1.07). The blue (light grey) dots represent metropolitan areas defined using our preferred 10% commuting threshold and the plain line is the associated regression line (slope -0.91). The red (dark grey) squares represent metropolitan areas defined using a 2% commuting threshold and the dashed line is the associated regression line (slope -0.81).

For Colombian metropolitan areas defined with our preferred commuting threshold of 10% and a minimum population size of 50,000, our estimate for ξ is 0.91 which suggests a distribution more uneven than implied Zipf's law. More generally, the estimate for ξ gets lower as lower commuting thresholds are considered. For a threshold of 30%, we estimate $\hat{\xi}_{30} = 1.00$; for 20% we get $\hat{\xi}_{20} = 0.95$; for 5%, we obtain $\hat{\xi}_5 = 0.88$; for 2% we have $\hat{\xi}_2 = 0.81$; finally for 1%, $\hat{\xi}_1 = 0.76$. Visual inspection of figure 4 confirms this.

The counterclockwise rotation of the Zipf line as lower thresholds are considered in figure 4 is easy to understand. On the one hand, a lower commuting threshold makes the largest metropolitan areas larger. On the other hand, there are more satellite municipalities so that the number of metropolitan areas decreases. In turn, this means that the smallest areas, just above the population threshold of 50,000, have a lower rank. Hence there is a downward shift of the left tail of the Zipf's regression line when lower commuting thresholds are used to define metropolitan areas. A combination of a shift rightwards for the largest areas and a shift downwards for the smallest areas obviously implies a flatter curve and a lower regression coefficient. We note that this would be observed even without censoring our observations at a population threshold of 50,000 since municipal aggregation overwhelmingly benefits large core municipalities and reduces the number of municipalities of a lower size.

This decline of ξ from 1.07 to 0.75 as lower commuting thresholds are considered shows that the estimates of the Pareto shape parameters for city populations are sensitive to how metropolitan areas are defined. Zipf's law is obtained exactly for a threshold of 30% but this is arguably too high a threshold to define meaningful metropolitan areas in Colombia. This result is in contrast with older findings by Rosen and Resnick (1980) that the size distribution of cities conforms better with Zipf's law when economically more meaningful definitions of cities are taken. This is also in contrast with more recent results by Rozenfeld *et al.* (2011) for the US and UK who find robust evidence for Zipf's law after defining cities using an aggregation criterion based on the geographical continuity of development.

To summarise, our findings suggest that the population of Colombian metropolitan

areas is fairly insensitive to the chosen commuting threshold. As lower commuting thresholds are considered, all the metropolitan areas that remain gain population but these increases tend to be small. Relative populations are even more stable since lower thresholds lead to population gains to all metropolitan areas. By contrast, the number of satellite municipalities is more sensitive to the chosen commuting threshold. As lower thresholds are considered, the number of satellite municipalities increases dramatically. Although lower thresholds lead to more satellite municipalities for most metropolitan areas, there is also growing heterogeneity with some metropolitan areas gaining a large number of satellites and some very few. In turn, these findings suggest that the physical extent of metropolitan areas is sensitive to the chosen threshold of commuting. In turn, the aggregation of municipalities also affects estimates of the size distribution of cities. Finally, we note that the stability of both population and the number of satellite municipalities is more marked around our reference commuting threshold of 10%.

6. Conclusions

In this chapter, we have proposed a simple way to define metropolitan areas relying exclusively on commuting patterns and implemented it on Colombian data. Aside from its simplicity, our approach offers two further advantages. It is fully transparent, which matters as soon as metropolitan area definitions affect policy interventions. The population of metropolitan areas is also highly robust to the details of the chosen threshold.

Our results also hold some interesting lessons for the description of large geographical networks such as labour market networks. Although their population is fairly insensitive to the details of the aggregation procedure, their physical extent is clearly much more sensitive to this. Our work also cautions against the use of simple summary statistics such as a Zipf exponent to characterise these networks.

References

- Alonso, William. 1964. *Location and Land Use; Toward a General Theory of Land Rent*. Cambridge, MA: Harvard University Press.
- Auerbach, F. 1913. Das Gesetz der Bevölkerungskonzentration. *Petermanns Geographische Mitteilungen* 59:73–76.
- Berry, Brian, Joe Loble, Peter G. Goheen, and Harold Goldstein. 1969. *Metropolitan Area definition: A re-evaluation of concept and statistical practice*. Washington, DC: US Bureau of the Census.
- Berry, Brian J.L. 1960. The impact of expanding metropolitan communities upon the central place hierarchy. *Annals of the Association of American Geographers* 50(2):112–116.
- Briant, Anthony, Pierre-Philippe Combes, and Miren Lafourcade. 2010. Does the size and shape of geographical units jeopardize economic geography estimations? *Journal of Urban Economics* 67(3):287–302.
- Charlot, Sylvie and Gilles Duranton. 2004. Communication externalities in cities. *Journal of Urban Economics* 56(3):581–613.
- Cheshire, Paul C. and Dennis Hay. 1989. *Urban Problems in Western Europe: An economic analysis*. London: Unwin Hyman.
- Cressie, Noel A. C. 1993. *Statistics for Spatial Data*. New York: John Wiley.
- Cörvers, Frank, Maud Hensen, and Dion Bongaerts. 2009. Delimitation and coherence of functional and administrative regions. *Regional Studies* 43(1):19–31.
- Duranton, Gilles and Henry G. Overman. 2005. Testing for localization using micro-geographic data. *Review of Economic Studies* 72(4):1077–1106.
- Duranton, Gilles and Diego Puga. 2013. The growth of cities. In Philippe Aghion and Steven Durlauf (eds.) *Handbook of Economic Growth*, volume 2. Amsterdam: North-Holland, forthcoming.
- Fox, Karl A. and T. Krishna Kumar. 1965. The functional economic area: Delineation and implications for economic analysis and policy. *Papers of the Regional Science Association* 15(1):57–85.
- Gabaix, Xavier and Rustam Ibragimov. 2011. Rank-1/2: A simple way to improve the OLS estimation of tail exponents. *Journal of Business Economics and Statistics* 29(1):24–39.
- Glaeser, Edward L. and Matthew Kahn. 2001. Decentralized employment and the transformation of the American city. *Brookings-Wharton Papers on Urban Affairs* :1–47.
- Hall, Peter G. and Dennis Hay. 1980. *Growth centres in the European urban system*. London: Heinemann Educational Books.
- Handbury, Jessie H. and David E. Weinstein. 2010. Is new economic geography right? Evidence from price data. Processed, Columbia University.

- Holmes, Thomas J. 1999. Localisation of industry and vertical disintegration. *Review of Economics and Statistics* 81(2):314–325.
- Kanemoto, Yoshitsugu and Reiji Kurima. 2004. Urban employment areas: Defining Japanese metropolitan areas and constructing the statistical database for them. Mimeographed, University of Tokyo.
- Krugman, Paul R. 1991. *Geography and Trade*. Cambridge, MA: MIT Press.
- Marshall, Alfred. 1890. *Principles of Economics*. London: Macmillan.
- Mills, Edwin S. 1967. An aggregative model of resource allocation in a metropolitan area. *American Economic Review (Papers and Proceedings)* 57(2):197–210.
- Molina, Humberto. 2001. *Análisis del Sistema Nacional de Ciudades. Aportes para una nueva regionalización del territorio colombiano*. New York and Bogota DC: UNDP and Ministerio de Desarrollo Económico.
- Muth, Richard F. 1969. *Cities and Housing*. Chicago: University of Chicago Press.
- Pérez, G. Javier. 2008. Población y Ley de Zipf en Colombia y la Costa Caribe, 1912 - 1993. Documentos de trabajo sobre economía regional 71, Banco de la República.
- Rosen, Kenneth and Mitchell Resnick. 1980. The size distribution of cities: An examination of the pareto law and primacy. *Journal of Urban Economics* 8(2):165–186.
- Rozenfeld, Hernán D., Diego Rybski, Xavier Gabaix, and Hernán A. Maske. 2011. The area and population of cities: New insights from a different perspective on cities. *American Economic Review* 101(5):2205–2225.
- Senado de la República. 2012. Informe de ponencia para segundo debate al proyecto de Ley número 141 de 2011. *Gaceta del Congreso* 137(21).
- US Office of Management and Budget. 2010. *2010 Standards for Delineating Metropolitan and Micropolitan Statistical Areas; Notice*. Washington, DC: Federal Register.
- Zipf, George K. 1949. *Human behavior and the principle of least effort: an introduction to human ecology*. Cambridge, Mass.: Addison Wesley.