DOTS TO BOXES: DO THE SIZE AND SHAPE OF SPATIAL UNITS JEOPARDIZE ECONOMIC GEOGRAPHY ESTIMATIONS?

Anthony Briant, Pierre-Philippe Combes, Miren Lafourcade

► To cite this version:

Anthony Briant, Pierre-Philippe Combes, Miren Lafourcade. DOTS TO BOXES: DO THE SIZE AND SHAPE OF SPATIAL UNITS JEOPARDIZE ECONOMIC GEOGRAPHY ESTIMATIONS?. 2008. halshs-00349294

HAL Id: halshs-00349294

<https://shs.hal.science/halshs-00349294>

Preprint submitted on 28 Dec 2008

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

GREQAM

**Groupeement de Recherche en Economie
Quantitative d'Aix-Marseille - UMR-CNRS 6579
Ecole des Hautes Etudes en Sciences Sociales
Universités d'Aix-Marseille II et III**

**Document de Travail
n°2008-61**

DOTS TO BOXES: DO THE SIZE AND SHAPE OF SPATIAL UNITS JEOPARDIZE ECONOMIC GEOGRAPHY ESTIMATIONS?

**Anthony BRIANT
Pierre-Philippe COMBES
Miren LAFOURCADE**

August 2008

Dots to boxes:
Do the size and shape of spatial units
jeopardize economic geography estimations? *

Anthony Briant[†] Pierre-Philippe Combes[‡] Miren Lafourcade[§]

August 2008

Abstract

This paper evaluates, in the context of economic geography estimates, the magnitude of the distortions arising from the choice of zoning system, which is also known as the Modifiable Areal Unit Problem (MAUP). We consider three standard economic geography exercises (the analysis of spatial concentration, agglomeration economies, and trade determinants), using various French zoning systems differentiated according to the size and shape of spatial units, which are the two main determinants of the MAUP. While size matters a little, shape does so much less. Both dimensions seem to be of secondary importance compared to specification issues.

JEL classification: R12, R23, C10, C43, O18.

Keywords: MAUP, concentration, agglomeration, wage equations, gravity.

*This paper was prepared as part of the PREDIT research contract n°04-MT-5036, whose financial support is gratefully acknowledged. The authors are also involved in the Marie Curie Research Training Network 'TOM - Transnationality of migrants' funded by the European Commission (contract No: MRTN-CT-2006-035873). Financial support from CNRS (Combes) is also gratefully acknowledged. We thank Alain Sauvart and Micheline Travet for crucial help in collecting the data. We are specially grateful to Sébastien Roux for his very kind cooperation on wage data. We are indebted to Yasushi Asami, Andrew Clark, Gilles Duranton, Laurent Gobillon, Keith Head, Thierry Mayer, Daniel McMillen and Henry Overman for fruitful discussions and insightful comments. Conference and seminar participants in Kyoto, Toronto, Barcelona, Lyon and Pau also provided us with very useful feedback.

[†]Paris School of Economics (PSE). PSE, 48 Boulevard Jourdan, 75014 Paris, France. anthony.briant@m4x.org

[‡]GREQAM-Aix-Marseille Université and PSE, also affiliated with the CEPR. GREQAM, 2 rue de la Charité, 13236 Marseille cedex 02, France. ppcombes@univmed.fr. <http://www.vcharite.univ-mrs.fr/pp/combes/>.

[§]University of Valenciennes and PSE. PSE, 48 Boulevard Jourdan, 75014 Paris, France. lafourcade@pse.ens.fr. <http://www.enpc.fr/ceras/lafourcade/>.

1 Introduction

Most empirical work in economic geography relies on scattered geo-coded data that are aggregated into discrete spatial units, such as cities or regions. However, the aggregation of spatial dots into boxes of different size and shape is not benign regarding statistical inference. Up until recently, economists paid little attention to the sensitivity of statistical results to the choice of a particular zoning system, known as the Modifiable Areal Unit Problem (henceforth MAUP). Our main objective here is to assess whether differences in results across empirical studies are really sparked by economic phenomena in the process under scrutiny, or rather just by different zoning systems. We first investigate whether changes in either the *size* (equivalently the number) of spatial units, or their *shape* (equivalently the drawing of their boundaries) alter any of the estimates that are usually computed in the economic geography literature. Second, we address the important question of whether distortions due to the MAUP are large compared to those resulting from specification changes.

Disentangling these two effects is essential for policy. For instance, much work has tried to check empirically whether agglomeration enhances economic performance at the scale of countries, European regions, US states or even smaller spatial units such as US counties or French employment areas. The size of the resulting estimates differs between papers, but we do not know whether this reflects zoning systems or real differences in the extent of knowledge spillovers, intermediate input linkages, and labor-pooling effects on firm productivity. The resulting economic policy prescriptions regarding cluster-formation strategies will be affected accordingly. In the same vein, a large body of literature has evaluated the degree of spatial concentration, but does not check whether the conclusion that some industries are more concentrated than others results from the chosen zoning system or from more fundamental differences in the size of agglomeration and dispersion forces across industries at different spatial scales.

This paper is based on three standard empirical questions in economic geography, although many others could have been considered.¹ We start by evaluating the degree of spatial concentration under three types of zoning systems: administrative, grid and partly random spatial units. We then compare the differences between concentration measures (Gini vs. Ellison and Glaeser) with those between zoning systems. We then turn to regression analysis. Not only is the measure of any spatial phenomenon likely to be sensitive to the MAUP, but also its correlation with other variables. We estimate the impact of employment density on labor productivity and compare the size of agglomeration economies across zoning systems and econometric specifications. Finally, we run gravity regressions. As trade determinants are highly sensitive to distance, and are hence a priori more exposed to the MAUP, it is particularly relevant to evaluate the impact of MAUP relative to mis-specification in this context.

All of the empirical exercises suggest that changing the size of spatial units only slightly al-

¹For comparison purposes, we use the same specifications as those typically found in the literature (see Combes, Mayer, and Thisse (2008)), even though we do not necessarily think that they are the most apt.

ters economic geography estimates, and changing their shape matters even less. Both distortions are secondary compared to specification issues.

The remainder of the paper is organised as follows. Section 2 provides a simple illustration of the possible size- and shape-dependency of spatial statistical inference, along with a brief review of the MAUP literature. Section 3 lists the zoning systems for which our estimations are carried out. As a first sensitivity test, Section 4 is dedicated to the study of French spatial-concentration patterns. Sections 5 and 6 investigate the extent to which changing econometric specifications and zoning systems affect the size and significance of wage and trade determinants respectively. Section 7 concludes and suggests further lines of research.

2 The Modifiable Areal Unit Problem: A Quick Tour

The Modifiable Areal Unit Problem is a longstanding issue for geographers. In their seminal contribution, Gehlke and Biehl (1934) were the first to emphasize that simple statistics such as correlation coefficients could vary tremendously across zoning systems. They note that, in the United States, the correlation between male juvenile delinquency and the median equivalent monthly housing rent increases monotonically with the size of spatial units. Openshaw and Taylor (1979) pursued this line of investigation and, drawing on correlations between the percentage of Republican voters and the percentage of the population over sixty, standardized what they called the “Modifiable Areal Unit Problem”.²

2.1 A simple illustration of the MAUP

Spatial statistics may vary along two dimensions: firstly, the level of aggregation, or the *size* of spatial units, and secondly, at a given spatial resolution, the drawing of their boundaries, or their *shape*. Figure 1 illustrates these two related issues via the employment density - labor productivity relationship.

Black points display the location of skilled workers, whose individual productivity is denoted $\bar{y}$, while empty dots stand for unskilled workers, with productivity $y < \bar{y}$. In the top figure, space is divided into four rectangles, each consisting of three skilled and two unskilled workers. The spatial distribution of workers across units is uniform and average productivity is the same across units. To illustrate the shape effect, consider the bottom-left figure. Spatial concentration emerges here, with two clusters of six high-skilled workers and two clusters of four low-skilled workers. Average productivity is higher in the former due to the spatial sorting of labor skills. Hence, agglomeration economies, defined here as the positive correlation between productivity and employment density, are zero in the first zoning system but positive in the second. We now turn to the size effect. In the bottom-right figure, we consider smaller rectangles with the same proportions as in the top figure. Spatial concentration is also found

²See Fotheringham and Wong (1991) for an extended review of the earliest MAUP contributions.

Figure 1: The size and shape issues

here, but the relationship between productivity and density is less marked than in the bottom-left case. Indeed, the difference in productivity between low- and high-productivity regions remains the same (except for empty boxes), whereas the density gap increases in the bottom-right case.

The extent and scope of agglomeration economies change with the size and shape of units, even though the underlying spatial information -the location and productivity of workers- remains the same. Even so, the issue is more subtle than at first sight. While it is clear that changing the zoning system will likely alter the perception of a particular phenomenon, the crucial question is whether the measurement errors this induces are systematic or random. If the former, extreme caution is warranted in interpreting the results, because systematic errors imply that, for a given zoning system, part of the underlying process has been omitted from the estimation. Consequently, policy prescriptions from a particular zoning system may not be very informative. In the latter case, as the economic phenomenon is neither size- nor shape-dependent, the MAUP is not an issue *per se*, since most estimation procedures are robust to random measurement errors. This paper aims to disentangle the two effects.

2.2 Related literature

A number of authors have provided detailed analyses of the MAUP. Using simple univariate statistics, such as the mean and the variance, Amrhein and Reynolds (1996) and Amrhein and Reynolds (1997) show that size and shape distortions depend on the aggregation process, and more precisely on whether information is either averaged or summed, as well as on the spatial organization of raw data, as reflected for instance in its spatial autocorrelation coefficient. Ac-

cording to Arbia (1989), size and shape distortions are minimized (although never eliminated) under two restrictive conditions: the exact equivalence of sub-areas (in terms of size, shape and neighboring structure) and the absence of spatial auto-correlation. Clear theoretical underpinnings are more difficult to come by for other statistics such as correlations or regression estimates. Fotheringham and Wong (1991) consider a multivariate analysis of the determinants of mean household income for various zoning systems, and come to an alarming conclusion: “The MAUP [...] is shown to produce highly unreliable results in the multivariate analysis of data drawn from areal units”. They also find a sizeable range for correlation and regression coefficients, which are positively (or negatively) significant for certain data configurations, but insignificant for others, suggesting that inference is not robust to the aggregation process.

Amrhein (1995) was the first to suggest separating aggregation effects from other types of discrepancies, such as model mis-specification. One of his appealing results is that well-suited models, such as Amrhein and Flowerdew (1992), do not produce distortive aggregation effects, whereas others, for instance Fotheringham and Wong (1991), are contaminated by size and shape.

We extend this literature in a number of ways. First of all, we systematically assess the magnitude of size and shape distortions relative to mis-specification biases. Secondly, we examine different aggregation processes to test the robustness of economic inference to the MAUP, since for agglomeration economies, raw information is averaged over spatial units, while for trade flows it is summed. Finally, we extend the work of Fotheringham and Wong (1991) by comparing the estimates from six different administrative and grid zoning systems to those from a hundred random equivalent zoning systems.

3 Zoning systems and data

The first zoning system we consider is that composed of 341 continental “Employment areas” (EA). These spatial units are underpinned by clear economic foundations, being defined by the French National Institute of Statistics and Economics (INSEE) so as to minimize daily cross-boundary commuting, or equivalently to maximize the coincidence between residential and working areas. This zoning system was designed to reduce the statistical artefact due to boundaries, which is why it is widely used in France. As can be seen on the left-hand side of Figure 2, the average employment area is fairly small, covering 1,570 km², which is equivalent to splitting the U.S. continental territory into over 4,700 units.

Shape distortions can be identified from spatial units that are similar in size (or number) to differently-shaped employment areas. Conversely, size distortions can be highlighted with partitions of France involving units that are larger than the 341 EAs. Hence, to disentangle the two faces of the MAUP, we appeal to three other sets of zoning systems.

Figure 2: Small zoning systems

3.1 Administrative zoning systems

The first set refers to French administrative units. Continental France is partitioned into 21 administrative “Régions” (RE), depicted on the left of Figure 3, which are themselves split into 94 “Départements” (DE), shown on the left of Figure 4. All such units are aggregates of municipalities, the finest French spatial division for which data are available.³

It can nonetheless be argued that administrative boundaries do not capture the essence of economic phenomena that often spill over boundaries, which is one of the reasons why EAs were created. To circumvent this drawback, some authors, especially geographers, prefer to work with (often arbitrarily-drawn) grids. The rationale is that, even if they do not necessarily better match the “true boundaries” of economic phenomena, grid zoning systems provide a greater degree of homogeneity of spatial units than do administrative zoning systems.⁴

³The French metropolitan area is covered by 36,247 municipalities. The division of France into départements was adopted simultaneously with the first French constitution in 1790, replacing the old “provinces”, which more or less represented dioceses. However, these latter exhibited significant variation in tax systems, population and land areas, and the new division aimed to create more “regular” spatial units under a common central legislation and administration. Their size was chosen so that individuals from any point in the département could make the round trip by horse to the capital city in no more than two days, which translated into a radius of 30 to 40 km. Regarding the shape, a lively debate opposed the defenders of grid zoning against the partisans of keeping alive “the tight links created long ago by moral standards, customs, production, language and nature”. The former strategy proposed dividing France into 80 grid units, but the latter division was finally adopted, which resulted in 83 fairly homogenous (but not geometrically identical) départements, the number of which was later increased to 94. In 1956, départements were grouped together into régions, in order to lead some policies at a larger spatial scale.

⁴Another argument refers to intertemporal comparisons when using fixed grid zoning systems. These do not change over time, while administrative zones may do so. See ESPON (2006) for an overview of this issue.

3.2 Grid zoning systems

We therefore construct a second set of zoning systems purely based on grid units. We first enclose France into the smallest possible rectangle. We then divide this rectangle into smaller equally-sized sub-rectangles (based on longitude and latitude). As France is not rectangle (and closer to an hexagon, actually), several sub-rectangles which map onto water are obviously left out. We obtain the final grid by aggregating municipalities according to the sub-rectangle in which they have their centroid. The resulting spatial units are not perfect “rectangles” as their boundaries follow those of real municipalities. We choose the size of the sub-rectangles to produce three different zoning systems analogous to administrative zoning systems: 22 (non-empty) large rectangles (LR), 91 medium rectangles (MR) and 341 small rectangles (SR). It is worth noting that the largest zoning systems (LR and MR in Figures 3 and 4) include several rectangles which are partially truncated due to French national boundaries. The finest grid such as SR (Figure 2) circumvents this pitfall at the expense of geometry, since the units boundaries become increasingly ragged at the very fine scale. Therefore, overtly enlarging or tightening the units alters both their symmetry and regularity.

Figure 3: Large zoning systems

A comparison of the results obtained under respectively RE, DE and EA or LR, MR and SR gives a flavor of any size distortions. We capture the impact of shape by comparing the results obtained across zoning systems involving units of similar size (RE to LR, DE to MR, and EA to SR). While these comparisons tell us whether MAUP distortions exist, they do not indicate whether the differences in the results are systematic and significant, however, which is why we propose a third set of zoning systems.

Figure 4: Medium zoning systems

3.3 Partly random zoning systems

Our third set of zoning systems involves arbitrarily-drawn spatial units. We define a set of 100 different partitions of France, by randomly aggregating the 4,662 French “Cantons”,⁵ into zoning systems that have a number of units strictly equivalent to those of administrative ones (341 units for EA, 94 for DE and 21 for RE): we call these REA, RDE and RRE respectively. These are constructed using the following algorithm. We randomly draw one canton, called the seed, within each administrative unit. We then aggregate each seed to a second canton randomly drawn from those contiguous to it. We continue with a third canton and so on, until all existing cantons have been drawn. We run the algorithm 100 times at each scale. Broadly speaking, this procedure produces, for each scale, a partition of France with *jiggling* borders.

3.4 Characteristics of zoning systems

Our empirical analysis builds on sectoral time-series data at the municipal level. The aggregation into the aforementioned larger zoning systems yields a three-dimension panel of employment, number of plants and wages for 18 years (within the 1976-1996) period and 99 industries (at the two-digit level for both manufacturing and services). For 1996, we match this panel to a trade data set for manufactured goods. More details are provided in Appendix.

As can be seen in Table 1, zoning systems differ sharply in their economic features. The spatial variation in land area is smaller for small grid units than for employment areas, a property

⁵We use this intermediate grouping of French municipalities to reduce the computational time without losing too much spatial variability in the randomization process.

Table 1: Summary statistics

Zoning system		(EA)	(SR)	(DE)	(MR)	(RE)	(LR)
Number of units		341	341	94	91	21	22
Land area (km ²)	Av.	1569.8	1580.4	5733.3	5922.3	25663.4	24496.7
	Cv.	0.63	0.35	0.34	0.50	0.43	0.53
Employment (workers)	Av.	2012	2019	7300	7541	32678	31193
	Cv.	4.6	1.5	12.3	1.7	1.8	1.3
Employment density (workers/km ²)	Av.	4.6	1.5	12.3	1.7	1.8	1.3
	Cv.	8.7	3.1	6.3	1.7	1.8	0.8
Wage	Av.	1.3	1.2	1.3	1.3	1.3	1.3
	Cv.	0.2	0.2	0.1	0.1	0.1	0.1
Inter-area distance (km)	Av.	396	421	392	454	402	486
	Cv.	0.48	0.48	0.47	0.47	0.43	0.42
Trade flow (tons×1000)	Av.	30.96	35.55	84.03	112.09	778.84	918.56
	Cv.	3.1	3.8	3.0	3.1	1.5	1.9

Notes: (i) (EA): employment areas, (SR): small rectangles, (DE): départements, (MR): medium rectangles, (RE): régions, (LR): large rectangles. (ii) 1976-1996 average, except for trade flows (1996 value). (iii) Av. is the mean. Cv is the Coefficient of variation (standard deviation divided by mean). (iii) No units for wage because detrended and centered around individual mean.

that does not hold for larger administrative units. This reflects two opposite effects. On the one hand, grid units are more regular, which reduces the variance. On the other hand, the share of truncated grid units increases with size, which increases the variance. The latter effect dominates for medium and large units. A clear drawback of the grid strategy is that, when units are not small enough, the gains of reducing the variance of land area cannot be attained due to the irregularity of national borders. Conversely, this also shows that the French authorities were fairly successful in designing quite homogenous administrative units.

Regarding the other variables, an important distinction concerns the way in which information is aggregated. The values of *summed* information (employment and trade flows) increase with the size of the units, which is straightforward. By way of contrast, the overall picture should vary less for *averaged* information, insofar as boundaries do not create too many non-random errors. For instance, employment density differs only little across grid zoning systems, regardless of the size of their units, while it varies more for administrative units, which reflects that the design of administrative zoning systems was not based on this variable. The suspicion that the MAUP could therefore bias density estimations motivates the exercise carried out in Section 5. Average wages are little affected by zoning system.

Finally, variations across zoning systems are smaller for distance and trade than for employment density. Note however that distance increases with size and the shift from administrative to grid units. Therefore, boundaries do seem to affect the measurement of distance, which seems upward-biased for large-size and grid units. This other source of potential non-random errors is explored in detail in Sections 5 and 6.

4 Spatial concentration

Before turning to regression analysis, we carry out the most basic exercise in economic geography, which consists in measuring the extent of spatial concentration, an issue widely-covered in the literature. Apart from a small number of continuous approaches, such as Duranton and Overman (2005), work in this area is based on discrete zoning systems. While some work has focussed on the comparison of spatial concentration across industries, such as Ellison and Glaeser (1997), only little has assessed the legitimacy of comparing results across zoning systems that differ in the size and shape of spatial units. In this section, we compare the variability in concentration due to the zoning system with that from different concentration indices.

4.1 Gini indices

We compute the spatial Gini index associated with every administrative and grid zoning system by industry. We then rank industries by spatial concentration and compute Spearman rank correlations across zoning systems. The results are shown in Table 2.

Table 2: Spearman rank correlations of Gini indices (1976-1996 average)

	(EA)	(SR)	(DE)	(MR)	(RE)	(LR)
(EA)	1	0.99	0.99	0.99	0.95	0.95
(SR)		1	0.98	0.99	0.96	0.96
(DE)			1	0.99	0.97	0.97
(MR)				1	0.98	0.98
(RE)					1	0.98
(LR)						1

Note: (EA): employment areas, (SR): small rectangles, (DE): départements, (MR): medium rectangles, (RE): régions, (LR): large rectangles.

Rank correlations across zoning systems that are similar in size (EA and SR, DE and MR, and RE and LR) are very high, with values of at least 0.98 (see the sub-diagonal elements in Table 2). The ranking of industries is therefore virtually unaffected by changes in the shape of units. Size has a slightly greater effect on concentration. For instance, the rank correlation between employment areas and régions is 0.95, which remains high. Making shape more homogeneous across scales leads to similar results, with the correlation between small and large rectangle zoning systems being 0.96.

With respect to partly random zoning systems, we restrict ourselves to the year 1996 for computational reasons. Figure 5 plots the average (over 100 runs) of the Gini obtained in each industry, for respectively the REA, RDE and RRE zoning systems. Industries are ranked by ascending Gini at the REA level (the upper dark points).

Two comments are in order. First, regardless of the industry, the Gini index falls with aggreg-

Figure 5: The size-dependency of the Gini index

Note: (REA): Random employment areas, (RDE): Random départements, (RE): Random régions.

ation level, the difference being less pronounced at both lower and higher values. The intuition is that smaller units have more areas with no registered employment, which raises the Gini index mechanically. Second, Figure 5 shows that ranks are not consistent across zoning systems. For instance, industry hierarchy is not the same for REA as for RDE and RRE. However, rank correlations remain high, at respectively 0.95 for REA-RRE (on average, over the 10,000 possible combinations of zoning systems), 0.99 for REA-RDE and 0.97 for RDE-RRE. This further suggests that size matters little. Moreover, the weak shape-dependency observed for administrative and grid zoning systems is confirmed by the average industry rank correlations across random zoning systems at each scale, which are respectively 1, 1, and 0.98 for REA, RDE and RRE.

4.2 Ellison and Glaeser indices

It is well known that the spatial Gini index is contaminated by industry structure. Given total industry employment, industries with fewer plants will have higher Ginis, even with random plant location. Ellison and Glaeser (1997) develop a measure of concentration that is purged of this plant size effect. Table 3 depicts the Spearman rank correlations for this index.

The rank correlations are generally lower than those for the Gini indices. Hence, any distortions due to the MAUP are more pronounced when spatial concentration is measured via the EG index. In particular, size distortions are clearly aggravated, even though the rank correlations remain fairly high (0.80 for instance between employment areas and région).

Further support comes from partly random zoning systems. With respect to shape, the average rank correlation is 0.99 for both REA and RDE, and 0.88 only for RRE. Regarding size, the related correlation is 0.88 for the REA-RDE pair, 0.77 for REA-RRE and 0.81 for DE-RE. Hence, the EG index, although more sophisticated, is somewhat more sensitive to the MAUP. Such size effects clearly stand out in Figure 6, which depicts average (over 100 runs) EG indices for re-

Table 3: Spearman correlations for EG indices (1976-1996 average)

	(EA)	(SR)	(DE)	(MR)	(RE)	(LR)
(EA)	1	0.78	0.92	0.81	0.80	0.77
(SR)		1	0.72	0.86	0.84	0.82
(DE)			1	0.81	0.81	0.77
(MR)				1	0.93	0.89
(RE)					1	0.93
(LR)						1

Note: (EA): employment areas, (SR): small rectangles, (DE): départements, (MR): medium rectangles, (RE): régions, (LR): large rectangles.

spectively the REA, RDE and RRE zoning systems. Contrary to the Gini index, the EG index increases with the size of units. Further, aggregation discrepancies are not uniform, being more pronounced in the right-hand tail of the distribution.

Figure 6: The size-dependency of the EG index

Note: (REA): Random employment areas, (RDE): Random départements, (RRE): Random régions.

4.3 Comparison between the Gini and the EG

The success of the EG index over the Gini coefficient lies in its alleviation of concentration due to the location of big plants. The crucial question we address here is whether the zoning system affects the ranking of industries more than does the choice of the index itself. To answer, we turn to a between-index rank correlation analysis.

Tables 4 and 5 show that the between-index Spearman rank correlations are definitely smaller than their within counterparts. Even within each zoning system (the diagonal elements of Table 4), the rank correlation is 0.77 at best (for the RE zoning system), with the lowest correlation being 0.52 (EA or SR).

Table 4: Spearman rank correlations between Gini and EG indices (1976-1996 average)

	EG (EA)	EG (SR)	EG (DE)	EG (MR)	EG (RE)	EG (LR)
Gini (EA)	0.56	0.49	0.59	0.58	0.63	0.59
Gini (SR)	0.56	0.52	0.59	0.60	0.65	0.61
Gini (DE)	0.60	0.52	0.64	0.61	0.67	0.63
Gini (MR)	0.59	0.54	0.62	0.64	0.69	0.64
Gini (RE)	0.64	0.61	0.67	0.70	0.77	0.71
Gini (LR)	0.64	0.61	0.68	0.70	0.75	0.73

Note: (EA): employment areas, (SR): small rectangles, (DE): départements, (MR): medium rectangles, (RE): régions, (LR): large rectangles.

Table 5: Rank correlations between Gini and EG indices (Average over 100 runs)

	EG (REA)	EG (RDE)	EG (RRE)
Gini (REA)	0.56	0.55	0.60
Gini (RDE)	0.61	0.61	0.64
Gini (RRE)	0.65	0.66	0.75

Note: (REA): Random employment areas, (RDE): Random départements, (RRE): Random régions.

There is considerable evidence that index choice, which we can consider as a specification issue, produces greater distortions than the choice of zoning system, in terms of both size or shape. It should thus be of greater concern than the MAUP.

To further gauge the extent of the bias induced by mis-specification compared to the MAUP, a detailed look at particular industries is useful. The left-hand side of Figures 7 and 8, shows the concentration patterns in the “Textile” and “Hotels and Restaurants” industries: ⁶ it is clear that the Gini index exhibits more concentration at low-scale resolutions. Turning to EG estimates (Figures 7 and 8), the reverse pattern holds: concentration now increases with scale. Moreover, the time patterns of concentration are fairly consistent for a given index, regardless of the scale, while greater differences are found between indices. For instance, in Figure 7, the movement in the Gini index suggests an upwards trend in the textile industry from 1995 onwards, whereas no such trend is seen in the EG index at medium or large size. By way of contrast, the differences due to shape (that is, between the dashed and plain curves in Figures 7 and 8) are less pronounced than the differences between indices, and smaller than those between scales. We cannot therefore reject that size affects the level and, to a lesser extent, the time pattern of spatial concentration. However, the conclusions depend critically on the choice of the index (Gini or EG). Consequently, thinking about which index is the most relevant is more important than the MAUP: the finding of rising or falling concentration likely substantially affects the ensuing policy prescriptions.

⁶The (similar) results for other industries are available upon request .

Figure 7: Textile Manufacture Industry

Notes: (i) EA: employment areas, SR: small rectangles, DE: départements, MR: medium rectangles, RE: régions, LR: large rectangles.

Figure 8: Hotels and Restaurants

Notes: (i) EA: employment areas, SR: small rectangles, DE: départements, MR: medium rectangles, RE: régions, LR: large rectangles.

5 Agglomeration economies

While the MAUP only slightly distorts spatial concentration patterns, it might have a greater effect on the explanation of the spatial distribution of economic variables. We therefore now consider the incidence of the MAUP in the context of multivariate regression analysis. In this section, we focus on the estimation of agglomeration economies. Evaluating the magnitude of the benefits reaped from spatial proximity is important for policy, and much work, such as Ciccone and Hall (1996), has been devoted to the estimation of the productivity gains resulting from dense clusters of activities. The benefits from proximity to large markets and the local composition of labour skills are generally simultaneously estimated.⁷

We regress local wages, a frequently-used measure of local labor productivity, on local employment density. Let w_{at} denote the wage in area a at date t , computed as average earnings of all workers located in a at date t (henceforth the gross wage), and Den_{at} employment density (per square Km). The benchmark specification we run is the following:

$$\log w_{at} = \alpha \log Den_{at} + \gamma X_{at} + \varepsilon_{at}, \quad (1)$$

where X_{at} is a vector of control variables. We compare the estimated elasticity of wages to employment density across zoning systems. As above, we then check whether the choice of zoning systems matters less for the size of agglomeration economies than the biases from choice of controls in the wage equation, which is a specification issue.

5.1 A simple correlation

It is useful to briefly look at simple gross wage / density correlations to have an idea of agglomeration economies. Given the panel structure of the data, we estimate equation (1) with no controls other than time dummies. Table 6 shows the resulting elasticities.

Table 6: Gross wages and density: simple correlations

Zoning system	Dependent Variable: Log of gross wage (pooled across years)					
	(EA)	(SR)	(DE)	(MR)	(RE)	(LR)
Density	0.071 ^a (0.001)	0.07 ^a (0.002)	0.073 ^a (0.001)	0.05 ^a (0.002)	0.09 ^a (0.003)	0.099 ^a (0.006)
Time dummies	yes	yes	yes	yes	yes	yes
Obs.	6138	6118	1692	1638	378	396
R ²	0.47	0.24	0.73	0.38	0.76	0.55

Notes: (i) (EA) = employment areas; (SR) = small rectangles; (DE) = départements; (MR) = medium rectangles; (RE) = régions; (LR) = large rectangles. (ii) All variables in logarithms. (iii) Standard-errors in brackets. (iv) ^a, ^b, ^c : Significant at the 1%, 5% and 10% levels respectively.

⁷See Combes, Mayer, and Thisse (2008).

The elasticity of wages with respect to density is not significantly affected by shape. In both EA and SR, the value is about 0.07, which lies within the range of [0.04, 0.11] reported in Combes, Mayer, and Thisse (2008), drawn from the analysis of US and European data. Even though some differences result from the move to a larger scale, the shape effect remains small. Size differences do not really matter when moving from small to medium units, although larger differences occur as we move to the largest units. However, this is not necessarily due to the size distortion induced by the MAUP. At small spatial scales density economies mainly work through technological spillovers and labor pooling; at larger scales, they are mainly generated by sharing markets for final and intermediate goods. Hence, there is no reason why the elasticity of productivity with respect to density should be the same. However, we do not know how much of the difference reflects real economic phenomena and how much the MAUP.

It is worth noting that the explanatory power of employment density is significantly lower (almost halved) for the grid than for the administrative units. Therefore, boundaries which do not reflect administrative/economic realities do actually generate measurement errors, possibly in both the left-hand and right-hand side variables. However, the good news is that these errors seem to be largely randomly distributed: even though density loses explanatory power, the overall picture with respect to elasticity is one of stability.

As a second step, we compare the two MAUP effects to those resulting from the inclusion in the wage equation of controls for skills (Section 5.2), and market potential (Section 5.3).

5.2 Controlling for skills and experience

Our empirical analysis uses rich individual wage information from a large panel of workers followed across time and jobs. We are hence able to apply a sophisticated procedure to control for local skills, and to check whether the greater productivity observed in dense areas is partly due to the spatial sorting of workers as suggested by Combes, Duranton, and Gobillon (2008). In a first stage, we calculate local wages net of individual skills and experience, as follows:

$$w_{it} = \theta_i + \nu_{j(i,t)} + X_{it}\beta + \epsilon_{it},$$

where w_{it} is the wage of worker i at date t . This is a function of θ_i , an individual fixed-effect capturing the impact of both time-invariant observed and unobserved skills, $\nu_{j(i,t)}$, an effect specific to the firm j where i is employed at date t , and X_{it} a set of controls for worker's i experience at date t (age, age-squared, and number of previous jobs interacted with gender). Based on the estimates provided in Abowd, Creedy, and Kramarz (2003), and following Combes, Duranton, and Gobillon (2008), we define a wage net of any individual observed and unobserved skills and experience effects, $(w_{it} - \hat{\theta}_i - X_{it}\hat{\beta})$. We then compute the average of this net wage over all individuals living in the same area a , at date t (henceforth net wage). This yields a measure of local labor productivity purged of individual skills and experience. We proceed by regressing

net wages on employment density. The results are shown in Table 7.

Table 7: Net wages and density: simple correlations

Zoning system	Dependent Variable: Log of net wages (pooled across years)					
	(EA)	(SR)	(DE)	(MR)	(RE)	(LR)
Density	0.033 ^a (0.001)	0.028 ^a (0.001)	0.029 ^a (0.001)	0.023 ^a (0.002)	0.048 ^a (0.003)	0.052 ^a (0.004)
Time dummies	yes	yes	yes	yes	yes	yes
Obs.	6138	6118	1692	1638	378	396
R ²	0.22	0.10	0.34	0.24	0.62	0.57

Notes: (i) (EA) = employment areas; (SR) = small rectangles; (DE) = départements; (MR) = medium rectangles; (RE) = régions; (LR) = large rectangles. (ii) All variables in logarithms. (iii) Standard-errors in brackets. (iv) ^a, ^b, ^c : Significant at the 1%, 5% and 10% levels respectively.

The elasticity of net wages with respect to employment density is half of that for gross wages, which is a difference of an order of magnitude greater than that due to the MAUP. We therefore reach the same conclusion as previously: differences due to the size and shape of spatial units are small compared to the upward bias induced by the omission of workers' skills and experience in the wage equation. Moreover, shape distortions are even attenuated in many cases (between DE and MR, and RE and LR, for instance), once these controls are included.

5.3 Market potential as a new control

Not only local density and skill composition affect labor performance, but so does the proximity to large economic centers outside the area. A major drawback of the above wage specifications is that there are no controls for the relative position of the area within the whole economy. For instance, wage equations derived from fully-specified economic geography models, such as Redding and Venables (2004) and Hanson (2005), account for spatial proximity via structural demand and supply access variables. It is beyond the scope of this paper to replicate such a sophisticated and difficult to implement approach. Here we only include, as well as density, a Harris (1954) market potential variable based on the employment accessible from any given area, divided by the distance necessary to reach them:⁸

$$\text{Market Potential} = \sum_{a' \neq a} \frac{Y_{a'}}{Dist_{a,a'}}, \quad (2)$$

where $Y_{a'}$ is employment in area a and $Dist_{a,a'}$, the distance between areas a and a' . The results for gross and net wages are listed in Tables 8 and 9 respectively.

Once market potential is accounted for, the impact of density on gross wage is attenuated. Regardless of the zoning system and the wage (gross vs net), the density coefficients are around

⁸The literature shows that this atheoretic market potential often has the same explanatory power as structural market potential.

Table 8: The spatial determinants of gross wages

Zoning system	Dependent Variable: Log of gross wage (pooled across years)					
	(EA)	(SR)	(DE)	(MR)	(RE)	(LR)
Density	0.055 ^a (0.001)	0.065 ^a (0.002)	0.059 ^a (0.002)	0.05 ^a (0.002)	0.09 ^a (0.003)	0.098 ^a (0.006)
Market Potential	0.1 ^a (0.004)	0.099 ^a (0.008)	0.062 ^a (0.005)	0.079 ^a (0.008)	0.024 ^b (0.011)	-0.009 (0.02)
Time dummies	yes	yes	yes	yes	yes	yes
Obs.	6138	6118	1692	1638	378	396
R ²	0.52	0.26	0.75	0.41	0.77	0.55

Notes: (i) (EA) = employment areas; (SR) = small rectangles; (DE) = départements; (MR) = medium rectangles; (RE) = régions; (LR) = large rectangles. (ii) All variables in logarithms. (iii) Standard-errors in brackets. (iv) ^a, ^b, ^c : Significant at the 1%, 5% and 10% levels respectively.

Table 9: The spatial determinants of net wages

Zoning system	Dependent Variable: Log of net wage (pooled across years)					
	(EA)	(SR)	(DE)	(MR)	(RE)	(LR)
Density	0.027 ^a (0.001)	0.026 ^a (0.001)	0.021 ^a (0.002)	0.023 ^a (0.002)	0.048 ^a (0.003)	0.052 ^a (0.004)
Market Potential	0.037 ^a (0.004)	0.043 ^a (0.007)	0.036 ^a (0.006)	0.044 ^a (0.007)	0.023 ^b (0.01)	-0.0002 (0.012)
Time dummies	yes	yes	yes	yes	yes	yes
Obs.	6138	6118	1692	1638	378	396
R ²	0.23	0.10	0.35	0.26	0.62	0.57

Notes: (i) (EA) = employment areas; (SR) = small rectangles; (DE) = départements; (MR) = medium rectangles; (RE) = régions; (LR) = large rectangles. (ii) All variables in logarithms. (iii) Standard-errors in brackets. (iv) ^a, ^b, ^c : Significant at the 1%, 5% and 10% levels respectively.

30% lower. The impact of market potential is slightly stronger in the medium rectangles than for their administrative counterparts, départements. This is consistent with the intuition that cross-boundary discrepancies should be more salient for grid units that were not designed to minimize them in the first place. Market potential is no longer significant at the regional level, which is consistent with French régions being large enough to depend mainly on themselves (or possibly on foreign markets, which are not considered here) rather than on each other. Even so, this variable depends on distance which makes it definitely more sensitive to the MAUP than, for instance, density, as we noted in section 3.4.

Table 9 exhibits striking similarities. Shape has virtually no effect, and size alters only slightly the results. In addition to the aforementioned larger impact of density and the smaller impact of market potential for large zoning systems, most differences are insignificant and are much lower than those from a change in specification. For instance, density increases productivity by only 2.7% at the small-unit level, once skills and market potential are controlled for, while the baseline estimates were over 7%. This difference supports the findings in the literature (see Combes, Mayer, and Thisse (2008)) and confirms our conclusion that MAUP is of secondary concern compared to modeling issues.⁹

Figure 9 maps out the density and market potential coefficients obtained from the three partly random zoning systems. For a given size, the dispersion of estimates is much lower than that induced by a shift of specification, which confirms the absence of shape effects. The only significant difference due to size regards density in the largest units. Even so, this almost vanishes in the best specification (net wages), as do the differences in the impact of market potential. These conclusions clearly echo the findings of Amrhein and Flowerdew (1992) and suggest that a good specification is actually an efficient way to circumvent the MAUP, even when variables do depend on distance.

Our previous conclusion regarding the sensitivity of concentration measures is confirmed by the analysis of productivity: specification is more important than the MAUP and, within the MAUP, size matters more than shape, and more so when the model is mis-specified and variables are distance-dependent.

6 Gravity equations

To test whether distance plays a systematic role in aggravating the distortions due to MAUP, we turn to the estimation of gravity equations.

⁹One important concern is not tackled here. In the analysis of how agglomeration enhances performance, we inevitably face the major difficulty that causality could run both ways: then worker's location is actually determined by wages. However, we leave this issue to one side here, as it has already been extensively discussed in the literature, and is orthogonal to the MAUP.

Figure 9: The size- and shape- dependency of wage determinants

Note: (REA): Random employment areas, (RDE): Random départements, (RE): Random régions.

6.1 Basic gravity

The gravity model has been widely used to investigate the determinants of trade. A basic specification explains the trade flow $F_{aa'}$, originating from area a and shipped to area a' , by various proxies for the proximity between a and a' . These include the distance between a and a' , $Dist_{aa'}$, and a dummy variable stating whether the areas are contiguous, $Contig_{aa'}$. Finally, the famous “border effect” is captured by a dummy variable for within-area flows, $Within_{a=a'}$. As a first step, we estimate the following two-way fixed-effect specification:

$$\ln(F_{aa'}) = \theta_a + \theta_{a'} - \rho \ln(Dist_{aa'}) + \phi Contig_{aa'} + \psi Within_{a=a'} + \epsilon_{aa'}, \quad (3)$$

where θ_a and $\theta_{a'}$ are destination and origin fixed effects respectively, and $\epsilon_{aa'}$ is an error term. This fixed-effect approach has the attractive property of being structurally compatible with many trade models (based on comparative advantage as well as imperfect competition).¹⁰

Table 10 reports the related estimates under both the administrative and grid zoning systems. The distance elasticity is around 20% larger for grid than for administrative zoning systems, at a given scale. Contiguity is less affected by shape. Again, size effects are slightly more salient, especially when moving from the EA-SR to either the DE-MR or RE-LR zoning systems. The magnitude of the distance effect (in absolute value) increases with size (for the administrative and grid zoning systems), as does that of contiguity. The border effect is always lower for grid zoning systems, which is further evidence of the economic consistency of administrative units. Measurement errors stemming from the use of grid zoning systems are less random for trade and proximity variables than for wage and density. This makes sense as trade and proximity are inherently sensitive to distance, which is not the case for the other two variables. Moreover, shape and size systematically bias the measurement of distance, as observed in Section 3.4.

¹⁰See Feenstra (2003).

Table 10: Basic gravity

Zoning system	Dependent Variable: log of positive flows (Year 1996)					
	(EA)	(SR)	(DE)	(MR)	(RE)	(LR)
Distance	-1.002 ^a (0.022)	-1.175 ^a (0.024)	-1.604 ^a (0.061)	-1.899 ^a (0.048)	-1.678 ^a (0.116)	-1.996 ^a (0.09)
Border	1.810 ^a (0.062)	1.289 ^a (0.06)	1.619 ^a (0.113)	0.804 ^a (0.113)	1.676 ^a (0.155)	1.162 ^a (0.252)
Contiguity	1.010 ^a (0.041)	1.128 ^a (0.044)	1.045 ^a (0.063)	1.069 ^a (0.071)	0.768 ^a (0.073)	0.829 ^a (0.086)
Obs.	24849	22189	6600	5069	441	443
R^2	0.518	0.545	0.708	0.756	0.941	0.939

Notes: (i) (EA) = employment areas; (SR) = small rectangles; (DE) = départements; (MR) = medium rectangles; (RE) = régions; (LR) = large rectangles. (ii) All variables in logarithms. (iii) Standard-errors in brackets. (iv) ^a, ^b, ^c : Significant at the 1%, 5% and 10% levels respectively.

Figure 10 illustrates the way in which both size and shape affect the values and standard errors of estimates from partly random zoning systems. Dark dots in the top-left figure stand for distance (and contiguity and border in the top-right and bottom figures respectively). The 95% confidence interval is shown by the surrounding lighter dots. Note that random zoning systems are ranked by increasing estimated values. For all three proximity measures, we find that the dispersion of estimates increases with scale, suggesting more shape-dependency in larger zoning systems. Nonetheless, this dispersion is of lower magnitude than the differences due to moving from one scale to another (from REA to RDE or RRE, regarding distance and border effects). The shape-dependency of larger zoning systems (especially RRE) is due to two joint phenomena. First, coefficient estimation is more likely to suffer from finite-sample bias for larger (and hence less numerous) units. Second, the random process of aggregation is likely to produce more distinct zoning systems when data are aggregated into larger units.

6.2 Augmented Gravity

Barriers to trade do not only concern proximity. Other trade frictions result from costs unrelated to distance (such as trade policy, exchange-rate volatility, delivery times, and inventory or regulation costs), and from more subtle frictions due to the need to acquire information on remote trading partners or to enforce contracts, as emphasized by Rauch (2001). To tackle these, the literature extends the basic gravity model by making trade costs depend not only on spatial proximity but also on cultural and informational proximity. For instance Wagner, Head, and Ries (2002) report that migration between two countries enhances their bilateral trade by around 50%. To evaluate the trade-creating impact of social and business networks within countries,

Figure 10: The size- and shape-dependency of the impact of spatial proximity on trade

Note: (i) The coefficients (b) have to be greater (in absolute value) than 1.96 times the standard error (se) to enter into the 95% confidence interval. (ii) (REA): Random employment areas, (RDE): Random départements, (RE): Random régions.

Figure 11: The size- and shape-dependency of the trade-creating impact of migrants

Note: (i) The coefficients (b) have to be greater (in absolute value) than 1.96 times the standard error (se) to enter into the 95% confidence interval. (ii) (REA): Random employment areas, (RDE): Random départements, (RE): Random régions.

Combes, Lafourcade, and Mayer (2005) estimate:

$$\begin{aligned} \ln(F_{aa'}) = & \theta_a + \theta_{a'} - \rho \ln(Dist_{aa'}) + \phi Contig_{aa'} + \psi Within_{a=a'} \\ & + \alpha \ln(1 + Mig_{aa'}) + \beta \ln(1 + Mig_{a'a}) + \gamma \ln(1 + Plant_{aa'}) + \epsilon_{aa'}, \end{aligned} \quad (4)$$

where $Mig_{aa'}$ is the number of people born in area a' and working in area a , called (relative to area a) immigrants, $Mig_{a'a}$ are analogously emigrants, and $Plant_{aa'}$ is the number of financial connections between plants belonging to the same business group (see Appendix).

Table 11: Augmented Gravity

Zoning system	Dependent Variable: log of positive flows (1996)					
	(EA)	(SR)	(DE)	(MR)	(RE)	(LR)
Distance	-0.628 ^a (0.023)	-0.714 ^a (0.028)	-1.194 ^a (0.063)	-1.285 ^a (0.064)	-1.261 ^a (0.118)	-1.420 ^a (0.129)
Border	1.016 ^a (0.065)	0.691 ^a (0.07)	0.587 ^a (0.129)	0.044 (0.127)	0.192 (0.226)	-0.237 (0.378)
Contiguity	0.321 ^a (0.05)	0.421 ^a (0.051)	0.37 ^a (0.068)	0.335 ^a (0.075)	0.22 ^a (0.081)	0.412 ^a (0.119)
Immigrants	0.221 ^a (0.014)	0.217 ^a (0.014)	0.24 ^a (0.028)	0.245 ^a (0.034)	0.309 ^a (0.074)	0.211 ^b (0.092)
Emigrants	0.232 ^a (0.014)	0.244 ^a (0.015)	0.213 ^a (0.038)	0.274 ^a (0.035)	0.285 ^a (0.108)	0.245 ^c (0.139)
Plant network	0.063 ^a (0.018)	0.033 ^b (0.013)	0.258 ^a (0.08)	0.035 (0.049)	0.204 (0.177)	0.392 ^c (0.204)
Obs.	24561	21606	6600	5059	441	443
R ²	0.541	0.574	0.722	0.772	0.954	0.948

Notes: (i) (EA) = employment areas; (SR) = small rectangles; (DE) = départements; (MR) = medium rectangles; (RE) = régions; (LR) = large rectangles. (ii) All variables in logarithms. (iii) Standard-errors in brackets. (iv) ^a, ^b, ^c: Significant at the 1%, 5% and 10% levels respectively.

It can readily be seen from Table 11 that, controlling for networks reduces the distance elasticity by about one-third, whereas the contiguity effect is three to four times smaller. The border effect is reduced even further, and disappears completely at the RE-LR scales. These effects are far larger than those due to the two determinants of MAUP, size and shape.

It is worth noting that the trade-creating effect of migrants, which does not directly depend on distance, is robust to the shift of zoning system, in terms of both size and shape. By way of contrast, even though the trade-creating impact of business networks increases slightly with the scale of administrative units, this is no longer statistically significant for grid zoning systems. Figure 11 displays the estimated immigrant and emigrant coefficients in the same way as in Figure 10. Both groups of estimates monotonically increase with the level of aggregation.

We therefore continue to find that size matters more than shape. Moreover, the magnitude of this distortion is definitely larger than in our previous exercises. The obvious explanation is

that trade equations involve many distance-dependent explanatory variables. Since the MAUP is fundamentally linked to proximity mis-measurement, it is fairly intuitive that it jeopardizes the estimation of trade equations more than that of wage equations, and that it is more salient for market potential within wage equations, in particular when skill variables are omitted. Specification issues are anyway a more important concern.

7 Conclusions

The overall picture is fairly clear. The use of different specifications to assess spatial concentration, agglomeration economies, and trade determinants produces substantial variation in the estimated coefficients. In most cases, theory provides a clear explanation of such variations, which are much larger than those sparked by the MAUP. Although size might still be important, especially in the context of distance-dependent explanatory variables, it is of second-order compared to specification, while shape is of only third-order concern. On the other hand, when zoning systems are specifically designed to address local questions, as is the case for French employment areas, we definitely argue that they should be used. Those who are left with other administrative units should not worry too much, however. We therefore urge researchers to pay the most attention to choosing the relevant specification for the question they want to tackle.

We do not of course claim that the various specifications used in this paper are actually the best. They are simply those frequently found in the economic geography literature. Many other empirical questions can be considered. We focus on three simple exercises because they are quite different in spirit, and cover a wide range of estimations. This makes us fairly confident that our conclusions are robust to other exercises, even though this remains to be shown. Finally, note that the French economical and institutional design may be, by chance, particularly well-designed to minimize MAUP problems. We therefore encourage researchers to replicate the exercises carried out here in other contexts (countries and periods).

References

- ABOWD, J. M., R. CREECY, AND F. KRAMARZ (2003): "Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data," Cornell University Working Paper.
- ABOWD, J. M., F. KRAMARZ, AND S. ROUX (2006): "Wages, Mobility, and Firm Performance: Advantages and Insights from Using Matched Worker-Firm Data," *The Economic Journal*, 116(512), 245–285.
- AMRHEIN, C. (1995): "Searching for the Elusive Aggregation Effect: Evidence from Statistical Simulations," *Environment and Planning A*, 27, 105–119.
- AMRHEIN, C., AND R. FLOWERDEW (1992): "The Effect of Data Aggregation on a Poisson Regression Model of Canadian Migration," *Environment and Planning A*, 24, 1381–1391.
- AMRHEIN, C. G., AND H. REYNOLDS (1996): "Using Spatial Statistics to Assess Aggregation Effects," *Geographical Systems*, 3, 83–101.
- (1997): "Using the Getis Statistic to Explore Aggregation Effects in Metropolitan Toronto Census Data," *The Canadian Geographer*, 41(2), 137–149.
- ARBIA, G. (1989): *Spatial Data Configuration in Statistical Analysis of Regional Economic and Related Problems*. Kluwer, Dordrecht.
- CICCONE, A., AND R. HALL (1996): "Productivity and the Density of Economic Activity," *American Economic Review*, 86(1), 54–70.
- COMBES, P.-P., G. DURANTON, AND L. GOBILLON (2008): "Spatial Wage Disparities: Sorting Matters!," *Journal of Urban Economics*, 63, 723–742.
- COMBES, P.-P., M. LAFOURCADE, AND T. MAYER (2005): "The Trade-Creating Effects of Business and Social Networks: Evidence from France," *Journal of International Economics*, 66(1), 1–29.
- COMBES, P.-P., T. MAYER, AND J.-F. THISSE (2008): *Economic Geography*. Princeton University Press.
- DURANTON, G., AND H. G. OVERMAN (2005): "Testing for Localization Using Micro-Geographic Data," *Review of Economic Studies*, 72(4), 1077–1106.
- ELLISON, G., AND E. L. GLAESER (1997): "Geographic Concentration in U.S. Manufacturing Industries: A Dartboard Approach," *Journal of Political Economy*, 105(5), 889–927.
- ESPON (2006): "The Modifiable Areas Unit Problem," Discussion paper, European Spatial Planning Observation Network.

- FEENSTRA (2003): *Advanced International Trade: Theory and Evidence*. Princeton University Press.
- FOTHERINGHAM, A. S., AND D. W. S. WONG (1991): "The Modifiable Areal Unit Problem in Multivariate Statistical Analysis," *Environment and Planning A*, 23, 1025–1044.
- GEHLKE, C., AND K. BIEHL (1934): "Certain Effects on Grouping upon the Size of the Correlation Coefficient in Census Tract Material," *Journal of the American Statistical Association*, 29(185), 169–170.
- HANSON, G. (2005): "Market Potential, Increasing Returns, and Geographic Concentration.," *Journal of International Economics*, 67, 1–24.
- HARRIS, C. (1954): "The Market as a Factor in the Localization of Industry in the United States.," *Annals of the Association of American Geographers*, 44, 315–348.
- OPENSHAW, S., AND P. TAYLOR (1979): "A Million of so Correlation Coefficients: Three Experiments on the Modifiable Areal Unit Problem," in *Statistical Applications in the Spatial Sciences*, ed. by N. Wrigley, pp. 127–144. Pion London.
- RAUCH, J. (2001): "Business and Social Networks in International Trade," *Journal of Economic Literature*, 39(4), 1177–1203.
- REDDING, S., AND A. VENABLES (2004): "Economic Geography and International Inequality," *Journal of International Economics*, 62, 53–82.
- WAGNER, D., K. HEAD, AND J. RIES (2002): "Immigration and the Trade of Provinces," *Scottish Journal of Political Economy*, 49(5), 507–25.

Appendix: Data

Economic variables for all zoning systems are obtained by aggregating information over the 36,247 French municipalities (“communes”).

First, over the 1976-1996 period, the composition in terms of establishments (employment size, and number of establishments) and workers (year and place of birth, age, gender, occupation, and wage, among others) is available at the 4-digit industrial level. The data come from the INSEE survey “*Déclaration Annuelle de Données Sociales*” (DADS), which collects matched employer-employee information in France. Our analysis builds on a panel extract covering people born in October of all even-numbered years, excluding civil servants, which is a representative $1/24^{th}$ of the French population. No survey was carried out in 1981, 1983 or 1990, producing a final sample of over 12.3 million plant-individual year observations, which are then re-aggregated by spatial unit, year (18 points), and industry (99 two-digit sectors covering both manufacturing and services).¹¹ As the key parameter of the sampling process is the date of birth, there is no obvious reason to believe that the sample is geographically biased.

For 1996, the above data are matched with information on the trade volumes shipped by road, both within and between municipalities, which we aggregate into different larger zoning systems. The data comes from the French Ministry of Transport, which annually surveys a stratified random sample of trucks.

Regarding social and business networks, we compute migrant stocks based on the number of natives from one area who moved to work in another area.¹² Business networks are captured via the number of financial connections between plants belonging to the same business group. For each business group, we count the number of plants located in each area. We then compute for each pair of areas the sum over all business groups of the product of the two counts. The data source here is the INSEE survey “*Liaisons Financières*” (LIFI), which defines a business group as the set of all firms controlled either directly or indirectly (over 50%) by the same parent firm, which is itself not controlled by any other firm.¹³

Bilateral distances between spatial units are computed as the average of the great-circle distances between their municipalities, weighted by total employment.

¹¹As in Abowd, Kramarz, and Roux (2006), part-timers are retained and outliers (over five standard errors above and below the mean) are dropped. The selection of industries and the removal of sampling errors at the smallest scale follows Combes, Duranton, and Gobillon (2008).

¹²This figure is also calculated using the DADS survey.

¹³See Combes, Lafourcade, and Mayer (2005) for more details on the network variables.