KANGAROOS, CITIES AND SPACE: A FIRST APPROACH TO THE AUSTRALIAN URBAN SYSTEM

Daniel ARRIBAS-BEL[*]
Fernando SANZ GRACIA[**]
Domingo P. XIMENEZ-DE-EMBUN[**]

Abstract - *Australia shapes a unique urban system. This paper examines the Australian urban system using data for urban centers and localities in 1996 and 2001. A summary and a basic descriptive analysis of the database are provided, followed by an examination of whether the system follows Zipf's and Gibrat's laws. None of them are found to hold. An Exploratory Spatial Data Analysis (ESDA) as well as a confirmatory analysis are carried out by using some of the most recent developments in spatial econometrics (Heteroskedastic Consistent GM Estimation) to analyze the spatial dimension of city size and growth, finding no influence for the former but a significant one for the latter.*

Key-words: AUSTRALIAN URBAN SYSTEM, ZIPF'S LAW, GIBRAT'S LAW, ESDA (EXPLORATORY SPATIAL DATA ANALYSIS), SPATIAL GIBRAT.

JEL Classification: J11, R00, R12

_____
[*] GeoDa Center for Geospatial Analysis and Computation, School of Geographical Sciences and Urban Planning, Arizona State University (USA) ; darribas@asu.edu.
[**] Department of Economic Analysis, Universidad de Zaragoza (SPAIN) ; fsanz@unizar.es, dpxe@unizar.es.

1. MOTIVATION

The case of Australia is a very interesting and particular one: with an extension of 7,741,220 Km^2, which represents 5.2% of the total world area, (the sixth largest country in the World), it hosts 22,466,224 people (Official Australian Population Clock, September 20th., 2010), only about 0.32% of the total world population, which sets it in the 53th. position. Taking both measures together, this implies the Australian population density is one of 2.6 people per Km^2, which makes it ranks as the seventh lowest density in the World. Australia is also one of the richest and most developed countries: by looking at the per capita GDP (nominal), it is within the first fifteen. Finally, its condition of an island as well as its special and very unique geography have shaped the distribution of population across space in a way that most people live by the coast (specially in the eastern one), leaving in the inner land an incredibly large empty space that may be called demographic desert.

On the other hand, there is a large branch in the urban economic literature analyzing both theoretically and empirically the distribution of the population within an urban system as well as its evolution over time, most of it using Zipf's and Gibrat's Law as tools to describe it. On the theoretical side, Cordoba (2008), Duranton (2007), and Gabaix and Ioannides (2004) are good examples; on the empirical side, although the main target of the studies is USA (Beeson *et al*., 2001; Black and Henderson, 2003; Overman and Ioannides, 2001), several other countries have been chosen, such as China (Anderson and Ge, 2005), India (Sharma, 2003), Malaysia (Soo, 2007), Japan (Davis and Weinstein, 2002; Davis and Weinstein, 2008), France (Eaton and Eckstein, 1997), Austria (Nitsch, 2003), Germany (Bosker *et al*., 2007; Bosker *et al*., 2008; Brakman *et al*., 2004), Spain (Lanaspa *et al*., 2003; Lanaspa *et al*., 2004) and even some cross-country analysis (Rosen and Resnick, 1980; Soo, 2005). However, very little papers have looked at Australia in a detailed fashion, despite its special characteristics, already noted by Rosen and Resnick (1980).

The present paper is thought to fill that (almost) empty space: it examines the whole Australian urban system in 1996 and 2001 in a detailed way displaying many features which set Australia far apart from other countries. While most of the results obtained for other countries show a Zipf coefficient around one, regressions in this paper show Australia has a much lower one, around 0.7, which means a more uneven population distribution among the cities of the system. When analyzing the relation between growth and city size, we also find Australia does not meet Gibrat's law. In addition, special emphasis is set in the spatial dimension of both variables (size and growth) to test to which extent we can speak of spatial association in the urban system. Related to this, an exploratory as well as a confirmatory spatial analysis are carried out being the main conclusion that while sizes do not show any kind of autocorrelation, urban growth does appear to be spatially related.

It is important to make a point regarding the methods used in this paper. While those employed in the non-spatial part are fairly common (i.e. adaptive kernel estimation or OLS regression), the ones that try to account for space

combine visualization tools (i.e. the *Clustercart*) with some of the most recent advances in spatial econometrics to try to offer correct estimates in presence of spatial autocorrelation. This kind of framework, although very useful to include spatial effects into the analysis, has rarely been used when looking at city size distribution.

The remainder of the paper is organized as follows. Section 2 describes the data set, why it has been chosen this way and gives some very basic statistics to get a first feeling. Section 3 analyses the Zipf's relation in Australia, while Section 4 looks at the link between the city size distribution and urban growth, testing whether Gibrat's law holds for the sample or not. An exploratory spatial data analysis (ESDA) to urban population and growth is applied in Section 5 while the confirmatory analysis may be found in Section 6. Section 7 closes the paper by adding some conclusions and pointing to further steps to be taken.

2. DATASET

Since the main purpose of this paper is to analyze the Australian urban system, the spatial unit used will be the Australian Bureau of Statistics' (ABS) "Urban Center and Locality" (UC/L from now on), which groups Collection Districts (CD's from now on, they are the smallest ABS's spatial unit in the Australian Standard Geographical Classification) together to form defined areas according to population size criteria by using census counts. In broad terms, an Urban Center is considered to be a population cluster of 1.000 or more people while a Locality is a population cluster of between 200 and 999 people (thus it does not cover the entire Australia). Each UC/L has a clearly defined boundary and comprises one or more whole CDs[1]. The data set used for this paper contains census counts from the 2001 Census of Population and Housing and 1996 Census data based on 2001 Census geography.

Both the choice of UC/L as unit and the adjustment between 2001 and 1996 boundaries imply a drop in the final dataset size considered which makes the sample smaller than the total Australian population. We are using 86.42% and 88.39% of the total population in 1996 and 2001, respectively.

In the rest of the paper we use the absolute population of city i (S_i) divided by the mean (*relative size* or $\frac{S_i}{S^{av}}$) or the total population (*share* or $\frac{S_i}{\sum_i S_i}$). This allows to relate each nucleus behavior to that of the whole distribution. Finally, there is an additional reason to use relative measures, and it is that, as Gabaix and Ioannides (2004) put it, "talking about steady-state distributions requires a normalization of this type".

---

[1] For more information, refer to:
• Australian Standard Geographical Classification. July 2007,
• Australian Bureau of Statistics website (www.abs.gov.au).

Finally, we conclude this section with a non-parametric analysis. Figure 1 shows the estimation of the log of relative population density function for both 1996 (dotted line) and 2001 (straight line) by means of an adaptive kernel *à-la* Silverman (1986)[2].

We may observe two main features in the figure: one dealing with the general shape of both kernels and the other regarding the evolution from 1996 to 2001. The first one is that, in both years, most of the probabilistic mass is placed on the left of zero; provided it is log of the relative population, zero implies the city whose population is average and then we can observe most of the Australian cities have a size far below the average one. In relation to the evolution, as one might expect from a short period of time, the main conclusion is there are not big differences. However, if some, one might notice the peak has moved rightwards. This, together with the fact the average city has grown, comes from the general growth of Australian population (6.1%): as more people exist, it is logical to think cities will be bigger.

Figure 1. Kernel density plot
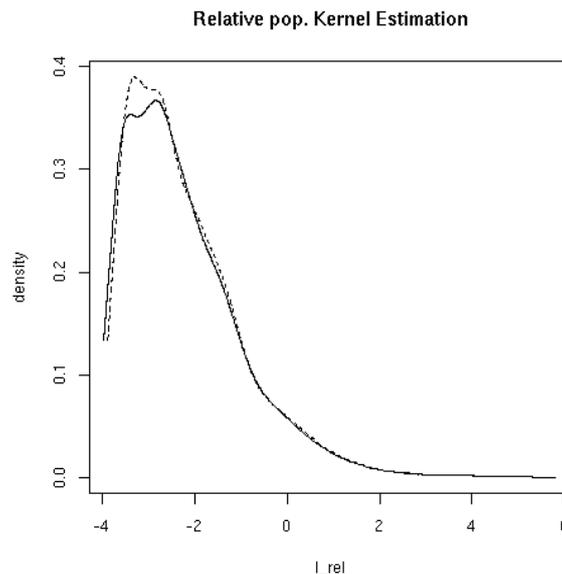
Dotted line corresponds to 1996 and the straight one to 2001.

3. ZIPF IN AUSTRALIA?

3.1. Zipf's Law

A common procedure widely used in the literature to rapidly characterize an urban system is to test how well the sample fits a power law. The theoretical basis of this practice comes from the statistical definition of Zipf's law, which

---

[2] The kernel was estimated with R's package "quantreg", freely available in the CRAN repositories (http://CRAN.R-pro ject.org).

relates to the notion of Pareto distribution. An urban system is said to follow a power law if :

$$P(Size > S) = \frac{a}{S^{\alpha}} \tag{1}$$

After some easy algebra, from eq. 1 and the consideration of the empirical distribution, we have :

$$lnR_i = A - \alpha lnS_i \tag{2}$$

where A is a constant and R_i is the rank of city i (1 for the largest one, 2 for the second one, etc.), which is the common specification to test empirically Zipf's law. In (2), α can be understood as a measure of the degree of evenness in the system: extremely, if $\alpha = \infty$ the graph is a vertical line around a size and every city has that size; opposite, if $\alpha = 0$ the degree of unevenness is maximum. We call the "rank-size rule" when α is around 1 and, in such case, we consider Zipf's law holds, because the power law is just an approximation of the real Zipf's expression. As Gabaix and Ioannides (2004) put it: "even if Zipf's law holds perfectly, the rank-size rule would hold only approximately". In this situation, the second largest city is half the size of the first one, the third largest one is one third the first one, and so on.

In this section, some results on the power law for Australia are offered. Since (2) is invariant to increasing monotone transformations in Si , there is no difference between any of the three measures (absolute, relative and shares of the total) and hence only relative sizes will be extensively shown.

3.2. Basic Zipf

Figures 2(a) and 2(b) show the Zipf plots for both 1996 and 2001 for which expression (2) has been run and Table 1 displays the regression output for both years. As it can be observed, the standard error has been corrected following Gabaix and Ioannides (2004).

The parameter α, indicating the way the population is distributed across the cities in the system, shows always as significant and around 0.74, which implies a distribution very unequal and sets Australia far from meeting Zipf's law. Furthermore, there has been a decrease from 1996 to 2001, which would mean more inequality in the distribution. However, we cannot take this result as very sure since the time horizon is not long enough and urban evolution is a phenomenon which evolves basically in the long run. Also, this decrease in α might be due to the fact that we are taking only those settlements above 200 people. Since the population is increasing over time, the minimum value will be always 200 (thought there do exist smaller settlements, which are not included in the sample) but the maximum may increase. This, everything else hold constant, may cause this increase in the degree of unevenness.

Figure 2. Zipf Plots

(a) UC/L-2001

(b) UC/L-1996

(c) UC-2001

(d) UC-1996
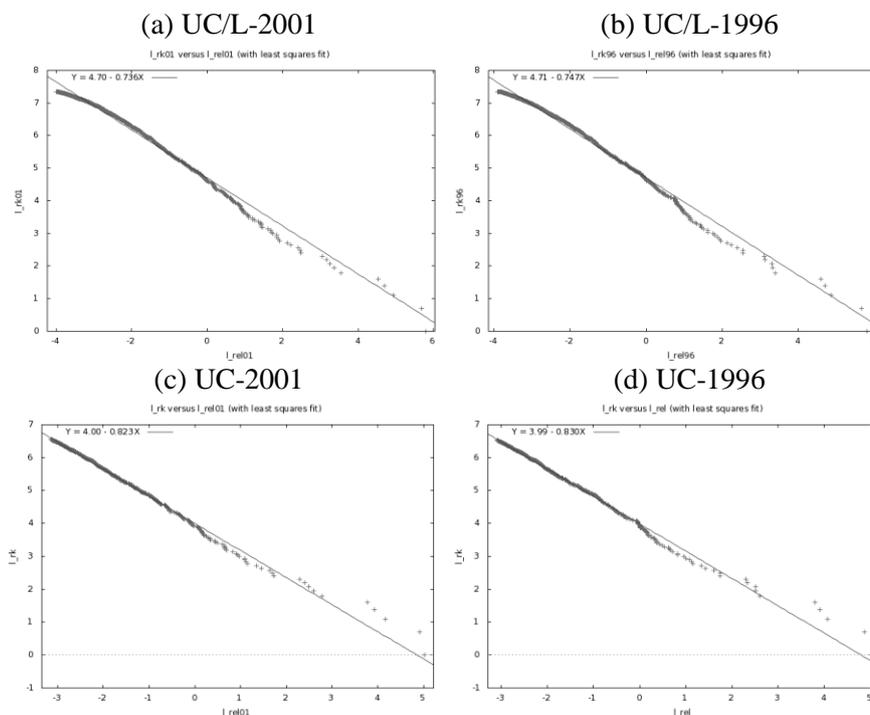
Horizontal axis represents the logarithm of the relative population and the vertical one the logarithm *of the rank.*

Table 1. Log of Australian relative population. OLS results

	Coefficient	Std. Error	t-stat.	p-value	Adj. S.E.	t-Stat.
			2001			
A	4.70264	0.00600962	782.5182	0.0000	-	-
α	0.736492	0.00230305	319.7901	0.0000	0.036	-20.458
	Unadj. R^2	0.985003		Adj. $\overline{R}^2$	0.984994	
			1996			
	Coefficient	Std. Error	t-stat.	p-value	Adj. S.E.	t-Stat.
A	4.70875	0.00554564	849.0899	0.0000	-	-
α	0.746700	0.00215998	345.6980	0.0000	0.036	-20.742
	Unadj. R^2	0.987139		Adj. $\overline{R}^2$	0.987131	

3.3. (Yet) more Zipf

By looking at the two first plots in Figure 2, we can also observe the actual distribution does not exactly fit a straight line but there are several deviations. Specially, there are downwards curves at the upper and lower extremes. These deviations from a straight line usually appear when not only the upper tail but the whole urban system is taken (when there is no cut-off), as Eeckhout (2004) states. Indeed, if we shorted the data-set so that only the biggest cities were considered (upper tail), the graph would look more like a straight line. Following Eeckhout (2004), this occurs because the underlying distribution is log-normal and not Pareto as it used to be assumed. Comparing Figure 2 (a)-(b) to Figure 2 (c)-(d) allows the reader to notice such phenomenon.

In addition, another feature of shortening the sample is that the line becomes steeper, that is, Zipf's parameter (α) increases. We can confirm this if we consider only the Urban Centers (settlements above 1000 people) instead of Urban Centers and Localities (above 200 people). By doing such experiment, we observe how α increases from around 0.74 up to about 0.83 (still far from Zipf's rule).

Furthermore, inspired by Ellis and Andrews (2001), the Australian urban system is divided into seven sub-regions[3] and Zipf's analysis is performed again to try to verify their argument. Their idea is that, due to the fact Australia has a relatively small population spread over a large area, "transport costs and political institutions may have induced multiple centers of economic activity", leading to a nationwide urban system made up of several state rank-size relations where the largest city is a primate (much larger than the rest) and the rest meet Zipf's Law. However, rank-size regressions were performed for each sub-region, finding roughly the same results as in the general case. The largest α coefficient was 0.75 (Southern Australia in 1996), which is still far from the unity. This leads to the conclusion there is not such a regionalization of Zipf, but rather a mirroring of the general picture.

Finally, we zoom out to the international context in Table 2. So far, we have described Australia as a *different* urban system; however, we have said nothing about other economies. Here our purpose is to confirm our suspects that it really shapes a different case. In order to compare results obtained for several different countries when applying a Zipf's regression, we pick examples spread around the world with some apparent similarities, such as area (Canada), population or GDP (Netherlands). As we can observe, Australia's coefficient scores as the lowest one. Since they do not take the same number of cities, nor the same cutoff, one should not directly compare results, but yet this can be taken

---

[3] Apparently, Ellis and Andrews divide it by States (which formally would make up 11 divisions, accounting for both States and Territories, according to Edwards (2001)). However, in this paper Australia has been divided only into seven sub-groups because of three reasons: the geoeconomic reasonability of the seven divisions, the small-sized the data sets would get otherwise and the fact the Australian Bureau of Statistics handles Urban Centers and Localities this way when offering the data.

as a sign that Australian population is distributed across the urban system very unevenly, especially when compared to other countries in the world.

Table 2. International comparison

Country	$|\alpha|$ coef.	Year	n	Reference
Algeria	1.351**	1998	62	Soo (2005)
Australia	0.8234***	2001	703	
Brazil	1.1341**	2000	411	Soo (2005)
Canada	1.2445**	1996	93	Soo (2005)
China	1.3**	1999	2651	Anderson and Ge (2005)
India	1.1876**	1991	309	Soo (2005)
Japan	1.3169**	1995	221	Soo (2005)
Malaysia	0.856***	2000	171	Soo (2005)
Netherlands	1.4729**	1999	97	Soo (2005)
USA	1.3781**	2000	667	Soo (2005)

** Significant at 5% ; *** Significant at 1%.

Data for other countries than Australia are taken from Junius (1999) and relate to 1990. Australia's index has been calculated for 1996 using the Urban Centers only.

This table tries to show the most comparable results, hence data from the UC sample are displayed for Australia and data from Soo (2005) are taken for USA. In the latter case, if we considered data from Gonzalez-Val *et al*. (2010) or Eeckhout (2004) instead, the coefficient happens to be much lower due to the fact these works use the whole distribution, which implies a much larger n.

4. DOES SIZE MATTER FOR SPEED? THE GIBRAT'S LAW

4.1. Gibrat's Law

So far, we have only analyzed the static relationship between size and rank and compared it over different points in time. Though the sample here is not the most suitable one for these purposes (only two years are certainly not enough to draw strong conclusions), it is also interesting to examine how an urban system has dynamically changed, if so. Traditionally, there are two ways in the literature to analyze dynamical processes in cities: the parametric and the non-parametric approach. The former consists of linear regressions *à-la* β-convergence, as in growth and development theory, while the later uses Markov's transition matrices or density kernels. Here we focus on the first one.

One question we might wonder about is whether the growth of a city depends on its initial size or it is independent of it. The situation of no relation between the growth rate of the city and its initial size is called of proportionate growth and if that is the case, Gibrat's Law is said to hold. The conceptual Gibrat's expression to estimate is as follows:

$$\Delta S_i = c + \beta S_i \qquad (3)$$

where ΔS_i represents the growth rate of the city i. If Gibrat's law does not hold we can consider two possibilities: either there is a positive or a negative relationship between being big or not and growing fast or not. If such relation was positive, there would be a premium for bigger cities to attract people; on the contrary if smaller cities grew faster than bigger ones, the tendency would be to convergence among all of them. Finally, if there was proportionate growth, there would be no apparent relation between size and growth.

Gibrat's analysis tells us information about the evolution and direction of the urban system, and there are several implications for each scenario regarding economic or landscape-planning policy which make this kind of analysis of special interest to real world.

4.2. Gibrat and Australia: not quite good friends

In order to test Gibrat's Law, we consider two different specifications that we apply to the absolute and the relative populations, giving rise to the following four equations that we run:

$$\ln \frac{S_{01}}{S_{96}} = c + \beta \ln S_{96} + \varepsilon \qquad (4)$$

$$\ln \frac{S_{01}}{S_{96}} = c + \beta \ln \frac{(S_{01} + S_{96})}{2} + \varepsilon \qquad (5)$$

$$\ln \frac{\frac{S_{01}}{S_{01}^{av}}}{\frac{S_{96}}{S_{96}^{av}}} = c + \beta \ln \frac{S_{96}}{S_{96}^{av}} + \varepsilon \qquad (6)$$

$$\ln \frac{\frac{S_{01}}{S_{01}^{av}}}{\frac{S_{96}}{S_{96}^{av}}} = c + \beta \ln \frac{(\frac{S_{01}}{S_{01}^{av}} + \frac{S_{96}}{S_{96}^{av}})}{2} + \varepsilon \qquad (7)$$

where the subindices refer to the year and ε is the innovation term. The reason why we have considered two different definitions of size is for robustness purposes.

Table 3. Gibrat's OLS

Coefficients	Estimate	Std. Error	t-stat.	P-value	
		$2*\ln S_{96}$			
$\hat{c}$	-0.105667	0.023231	-4.549	5.82e-06	***
$\hat{\beta}$	0.021011	0.003242	6.480	1.22e-10	***
		$2*\ln \dfrac{(S_{01}+S_{96})}{2}$			
$\hat{c}$	-0.051841	0.023542	-2.202	0.0278	**
$\hat{\beta}$	0.013398	0.003294	4.067	5e-05	***
		$2*\ln \dfrac{S_{96}}{S_{96}^{av}}$			
$\hat{c}$	-0.029546	0.008554	-3.454	0.000567	***
$\hat{\beta}$	0.004164	0.003331	1.250	0.211519	
		$2*\ln \dfrac{(\frac{S_{01}}{S_{01}^{av}}+\frac{S_{96}}{S_{96}^{av}})}{2}$			
$\hat{c}$	-0.009945	0.008533	-1.165	0.244	
$\hat{\beta}$	0.012964	0.003302	3.926	9.03e-05	***

*** significant at 1% level; ** significant at 5% level; * significant at 10% level.
The estimated equations correspond to those especified in equations 4 to 7.

Results are shown in Table 3. Just by a quick look at the p-values, we can already conclude it is not really clear Gibrat holds for this case. In fact, we may observe that in those three specifications where Gibrat's law does not hold (β being statistically different from zero), the sign for the size parameter is positive. There is one more remark regarding this finding: the signs seem to be in line with the conclusions obtained in the Zipf analysis. In the previous section, we have seen the α coefficient for the size decreased from 0.75 to 0.74, implying a more uneven city system. This could be partly explained because of the faster growth of bigger cities, which widened the difference between them and the smaller ones, leading to greater inequality in terms of populations.

5. WHERE? BRINGING SPACE INTO ACTION. EXPLORATORY ANALYSIS

This paper was started by pointing to the uniqueness of Australia as an urban system, especially due to its particular geography. We have seen many facts confirming the first statement, but no word about the latter one yet. The

picture we have drawn of Australia is as follows. Although the average city size has increased by 8.45%, as many as 619 cities (out of 1559) decreased in population from 1996 to 2001. Moreover, if we looked at relative sizes, it was 1089 (out of 1559) cities that experienced a negative growth rate. This leaves us with a system becoming more uneven, with a few larger cities growing so as to push the average up, and many more cities declining in population.

With those numbers in mind, now we would like to be able to detect *where* such changes have happened. In this section we provide tools to visualize the spatial dimension of these phenomena by means of what is called Exploratory Spatial Data Analysis (ESDA) and look for patterns that we try to confirm later in Section 6 when we carry out the confirmatory analysis. But, before that comes, let us briefly explain the analytical framework that will help us go through the task.

5.1. Methodology

A very useful concept to step forward in this direction is that of *spatial dependence*. Following Anselin (1988, p.11), "spatial dependence can be considered to be the existence of a functional relationship between what happens at one point in space and what happens elsewhere". Translating that into our topic, if Australian geography played any role in explaining urban outcomes (in terms of either size or growth), we should be able to detect any type of spatial dependence. We can express the idea of spatial dependence in our case by means of a functional form:

$$S_i = f(S_j) \quad \forall i \neq j \tag{8}$$

or
$$gr_i = h(gr_j) \quad \forall i \neq j \tag{9}$$

where *gr* is the growth rate of the population of a city. One common way to introduce space into the formal analysis and account for the functional relation in (8) is by means of the *spatial weight matrix* (*W*). It is an *n* by *n* matrix and is usually constructed considering relations of either physical contiguity or distance, although it can also be designed to express more complex spatial linkages such as economic or cultural distance, for instance. Every element w_{ij} of *W* reflects the spatial connection (or absence of it) between the observations *i* and *j*. To construct a spatial weight matrix based on contiguity, we need the space to be divided into polygons, not spattered with points. Since we are dealing with cities (which are considered to be points in a map), the first step is to convert the points into polygons. For that purpose, we define a Thiessen/Voronoi lattice.

The matrices used here are binary in the sense that neighbors are weighted with 1 and the remainder receives a weight of 0. In this paper, we have used the *queen* criteria. Due to the great concentration of cities in the coast, especially in the East, the distance based weight matrix is not suitable for the Australian case as the average number of neighbors for everyone to have at least one is 623 (while the average in the *queen* case is 5.93).

Once we have obtained *W*, the next concept to introduce is that of *spatial lag*. Analytically, it is expressed as follows:

$$sl(y) = Wy \qquad (10)$$

where *sl* stands for *spatial lag* and *y* is a variable. As Moreno and Vayá (2000, p. 27) put it: "the spatial lag consists of a weighted average of the values in the neighbor regions, taking the weights as fixed and given in an exogenous way". This lag can be understood as the analog for spatial econometrics of the time series' observation of the period *t* - 1 and in the same way the spatial dependence would correspond to the serial autocorrelation.

There are several tests to explore the presence of spatial dependence. Here we use the Moran's I (Moran, 1948) and Moran's scatter plot, two of the most common ways to test for the presence of spatial dependence. Moran's I can be seen as a measure of the correlation between each observation x_i and the rest of regions to which it is spatially linked.

There is one more analytical issue in relation to the global Moran's I: if the variable to be used is a rate, there is a variance instability problem (unequal precision) due to the use of rates as estimates for an underlying 'risk'. In order to correct for this problem, one can smooth the ratio by using several transformations proposed in the literature. Here we use the one following the Empirical Bayes (EB) principle, suggested by Assuncao and Reis (1999). Once the rate is transformed, the Moran's statistic can be applied as usual.

It is interesting to note that Moran's I is a global statistic and, as such, sum up in only one number the degree of spatial correlation among all the observations in the sample. However, it is not able to distinguish those situations in which such spatial correlation is homogeneously spread among the sample from those in which it is clustered in only a few observations. For that purpose, it is necessary to use local indicators of spatial association (LISA), which allow to *decompose* a global statistic of spatial correlation (such as the global Moran's I) into sub-indices for each observation, being very useful to identify clusters. In order to do that, the local test is computed for every observation in the sample, instead of computing a global measure of autocorrelation. Since the sample size may be large, the most common way to show the results is by means of a map in which different colors display different types of outcome.

Here we use the local version of Moran's I, proposed in Anselin (1995), whose expression is the following:

$$I_i = \frac{z_i}{\sum_i z_i^2 / N} \sum_{j \in J_i} w_{ij} z_j \qquad (11)$$

where z_i is the standardized Moran's I for observation *i* and J_i is the group *i*'s neighboring observations.

It is possible to assume the standardized I_i is distributed as a standard normal. Once standardized, positive (negative) values may be regarded as clusters of similar (dissimilar) values around the observation *i*.

As in the global case, if the variable to be used is a rate, the same instability in the variance is encountered. Then, there is a need to transform the variable in the same way as with the global Moran's I; namely the EB one.

5.2. Results

Once we have set up the theoretical background, we can delve into the Australian dataset and make use of the proposed toolbox[4]. In this section, only those results for absolute log-population and absolute growth rate are shown in order to improve clarity, especially as the conclusions do not differ very much from the other specifications.

Section 4 takes a deep look at the Zipf's relation in the Australian system. Now we wonder whether the variable size displays spatial correlation. In order to do so, we apply Moran's I to the absolute log-population. The effect seems clear: we cannot reject the null of no spatial correlation for both years 1996 and 2001. The result is showed in Table 4, the p-value is clearly larger than 0.01, implying that even at the 1% significance level and the null cannot be rejected. If we observe the graph in Figure 3 (a), we can see how the points are spread across the four quadrants and the fitted line mingles with the horizontal axis. This suggests there is no clear pattern for the size in the spatial configuration of the city system or, in other words, that cities are located in space without following any law regarding their sizes.

Table 4. Moran's I

Variable	Moran's I	Standard Moran's I	P-value
Log-population 1996	-0.001	-0.029	0.489
Log-population 2001	0.013	0.922	0.178
Absolute growth (EB transf.)	0.132	8.919	0.000

Section 4 tests the Gibrat's relation, which links the size of a city with its rate of growth, concluding there is no relation between both variables. Parallel to the Zipf's case, we try to put this phenomenon in space and detect if the way nearby cities evolve has any influence on a city's growth. The result, displayed in Table 4, seems to confirm the other side of the coin: Moran's I shows enough evidence to reject the null of no autocorrelation. This means the statistic is significantly different from zero. In fact, the sign is positive, implying the growth of a city and that of its nearby partners are positively correlated and, therefore, the growth tends to cluster in space giving room to the idea of loser and winner

---

[4] The spatial analysis was carried out with the open source package STARS (Rey et al., 2006), freely available at the REGAL's website : http://regionalanalysislab.org/index.php/Main/STARS and the free package GeoDa by Luc Anselin, available at the GeoDa Center for Geospatial Analysis and Computation's website (http://geodacenter.asu.edu).

areas. This relation can be observed graphically in Figure 3 (b): many points are located in the up-right and down-left quadrants, and the fitted line is clearly upwards.

Figure 3. Moran's scatter plots

(a) Absolute Log-Population: 2001 Growth   (b) EB transf. Absolute Population
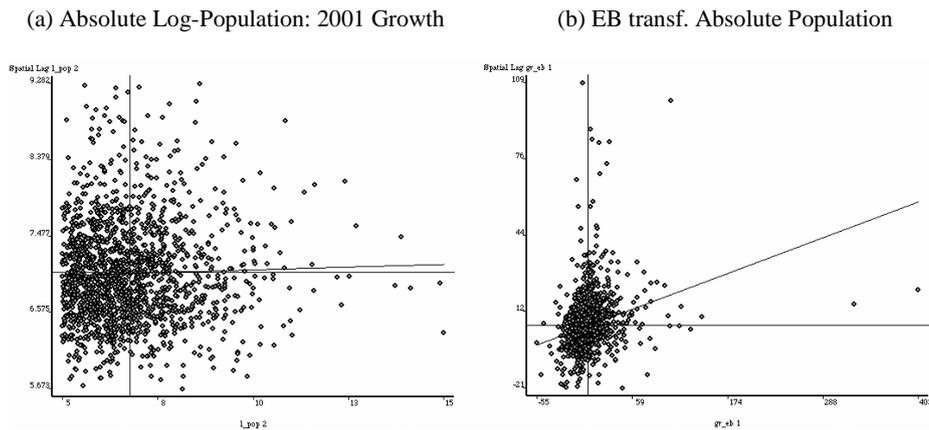
Horizontal axes represent the variable of interest and vertical ones its spatial lag.

Since the first spatial approach does not support the idea of the sizes being spatially correlated, it does not make much sense to try to look for the existence of cities with similar sizes grouped nearby, that is of actual clusters in the variable size. However, we have found the growth to be spatially dependent, which means there is some degree of spatial association for cities with similar growth rates. Therefore, it is interesting to take one step forward and try to look for the *hot* and *cold spots* of the Australian urban growth. We use the local indicators explained above (LISA) to do so. As said before, the usual way to present LISA results is by means of a map in which different colors imply different outcomes.

Due again to the spatial characteristics of the Australian urban system, we propose to use what we come to call a *ClusterCart*, a new alternative way of visualizing LISA indices that proves helpful in this sort of geographical layouts. Basically, the *ClusterCart* is the result of embedding the cluster results from the LISA statistics into a standard cartogram. We explain this idea more in detail below.

A cartogram is a map in which some thematic mapping variable is substituted for land area. As an example, Figure 4 (a) shows a standard cartogram of the Australian cities for a dummy variable of zeros. The polygons have turned to circles of the same size (due to the fact the variable each polygon represents is just a zero for each observation) which barely overlay each other. This produces an abstract representation which distorts the original shape of Australia but which, and this is why it is useful here, allows to see all the observations at

a glance[5]. It also shows a great illustration of why a clustergram is useful in this case: the vast majority of points are located on the east side of the island and they all stand very close to each other. As said before, this feature cannot be recognized by looking at a Voronoi map as the polygons on the east are too small to be noticed. However, since the clustergram gives the same size (still allowing to thematically color it) to every city, this fact can rapidly be discovered and overcome.

Figure 4. *ClusterCart* and its composition

(a) Australian Cartogram   (b) Bicolor LISA for growth
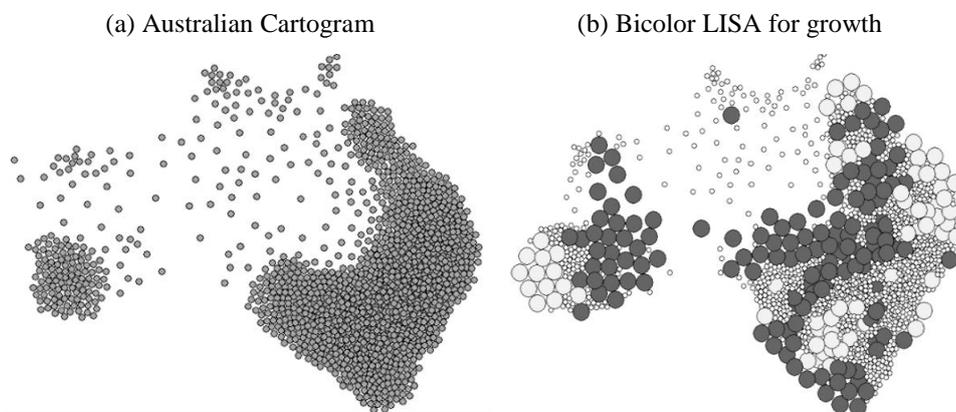

Figure 4 (b) shows the *ClusterCart*. To build it up we have created a cartogram using the cluster results so the light (dark) circles represent the cities with a high-high (low-low) outcome in the LISA and the transparent ones are the rest of the cities.

By looking at Figure 4 (b), we can extract some insights about the urban dynamics in Australia. The first one is that a vast part of the dark circles are not by the coast, except for some of them located in the South[6]. This suggests the Australian population is moving outwards, there is a "push-out" effect that makes cities in the inland decline their population. The follow-up obvious question is: "if population in Australia is growing over time and the inland is decreasing*, where is growth taking place?*" The answer can be found if we observe the lighter circles: basically, coastal and well-watered places are those displaying positive growth clusters, which is in line with the ideas stated in Hugo (2002). But, if we observe them more in detail, we can also detect most of the light clusters locate around some of the largest cities in the country: the one

---

[5] The reason why the map becomes distorted is because now the points cannot cover each other and, for that to happen, they need to be slightly moved from their original position in a standard map to leave room so they all fit in.

[6] However, the reader should note those points are located by Tasmanian coast, not the Australian one. This is due to the fact that the *ClusterCart* tends to group all the observations without distinction between one or another island.

in the West is around Perth's area, the one in the South around Melbourne's and the one in the East corresponds to Brisbane. This leads to conclude that it is in big metropolitan areas, rather than uniformly around the country, where the phenomenon of urban growth is taking place. This result also links to the findings in the non-spatial analysis of urban growth: we found that there seems to be a positive correlation between size and growth and, as it turns out when we spatially explore the data, the phenomenon of growth seems to be taking place in those areas where bigger cities are located. In addition, these results point to the power of larger cities to attract people around their *orbits of influence*, people who might or might not live in the inner city but who interact with it (e.g. commute for work). Looking for the reasons underlying such phenomena are beyond the aim of this study, but it certainly represents an interesting road to walk down for future research.

So far, this has only been a first attempt to explore whether space plays any role, but it does not provide any formal insight, nor tells anything about how Zipf's or Gibrat's analyses are affected. For the latter, we need to walk further and step into the confirmatory analysis. That is what next section is about.

6. DOES SPACE ACTUALLY MATTER?
CONFIRMATORY ANALYSIS

One of the main results from the previous section is that while there is no apparent correlation between sizes, there seems to be spatial autocorrelation in the growth rates. If the former is true and there is some sort of spatial dependence, OLS alone is not the best method to estimate coefficients (Anselin, 1988, p. 58-59). In this section, we try to confirm the directions pointed in the previous section and detect if space actually plays any role in the generating process by performing a test on the residual spatial autocorrelation and the robust LM test, to check if there is spatial autocorrelation in the regression analysis. Lastly, for those specifications that suggest there is a spatial autoregressive model, we estimate it by means of the GMM (Generalized Method of Moments) estimator developed by Kelejian and Prucha (1998).

6.1. Testing for space in the regression analysis

The first step we take is to test against the possibility that the residuals of the OLS regression show spatial randomness. We do this by means of the Moran's I test for spatial autocorrelation in the residuals suggested by Cliff and Ord (1981), in which the null is spatial randomness. The results for every specification are shown in the first column of Table 5. As we can observe, the conclusion is clear: while most (three out of four) of the Zipf's specifications show spatial randomness, as we cannot reject the null, we always do it when it comes to the Gibrat analysis, meaning that the residuals are not randomly distributed across space. However, one of the main issues of this test is the absence of a clear alternative hypothesis. It is only useful when the null cannot be rejected, as it happens for the Zipf's case, because then there is no need to go further in the spatial analysis. For this reason, we consider as not relevant to continue analyz-

ing urban size in a spatial context and, from now on, we only focus on urban growth and Gibrat's law[7].

Table 5. Tests on the Gibrat's and Zipf 's Analysis

Model specification	Moran res. test	Rob. LM error test
Zipf		
$\ln R_{01} \sim \ln S_{01}$	DNR	-
$\ln R_{96} \sim \ln S_{96}$	***	-
$\ln R_{01} \sim \ln \dfrac{S_{01}}{S_{01}^{av}}$	DNR	-
$\ln R_{96} \sim \ln \dfrac{S_{96}}{S_{96}^{av}}$	DNR	-
Gibrat		
$\ln \dfrac{S_{01}}{S_{96}} \sim \ln S_{96}$	***	***
$\ln \dfrac{S_{01}}{S_{96}} \sim \ln \dfrac{(S_{01}+S_{96})}{2}$	***	***
$\ln \dfrac{S_{01}/S_{01}^{av}}{S_{96}/S_{96}^{av}} \sim \ln \dfrac{S_{96}}{S_{96}^{av}}$	***	***
$\ln \dfrac{S_{01}/S_{01}^{av}}{S_{96}/S_{96}^{av}} \sim \ln \dfrac{\left(S_{01}/S_{01}^{av}+S_{96}/S_{96}^{av}\right)}{2}$	***	***

*** significant at 1% level; ** significant at 5% level; * significant at 10% level; DNR Do Not Reject the hypothesis of not significant.
In the model specification, Moran Res. Test refers to the Moran's residual test for spatial autocorrelation; and Rob. LM Error Test to the LM robust test for an error model.

In order to get a better insight about the underlying process, we perform the Lagrange Multiplier statistic for spatial dependence in the error. We use the variant which tests for the presence of spatial dependence in the error term and it is robust to the presence of a missing spatially lagged dependent variable[8]. The test considers the following structure for the error term:

$$u = \rho W u + \varepsilon \qquad (12)$$

---

[7] Although there is one Zipf's specification rejecting the null, for the sake of focus on one issue (urban growth), we do not consider it any further in the analysis.
[8] The computations were performed with the R package 'spdep', freely available at the CRAN website (http://cran.r-project.org) and PySpace, also open-source available at http://geodacenter.asu.edu/software. For more information on the latter one, see Rey and Anselin (2007).

and the hypothesis considered as null is $\rho = 0$. Thus, if we reject it, we may say there is spatial autocorrelation in the error term.

Although it would also be desirable to test for the presence of a missing spatially lagged dependent variable as well, there is one reason why we have not. The usual procedure to estimate a model with a spatial lag is by means of the two steps least squares (2SLS) procedure, which uses instrumental variables (IV's) to correct for the endogeneity created by the dependent variable spatially lagged on the right hand side. Kelejian and Prucha (1998) set the conditions for this method to work and, as assumption 7b in their paper implies (p. 105), it requires at least one of the explicative variables (excluding the constant term) be significant. Since the main goal of Gibrat's analysis is to check whether size, as the only explicative variable, explains any of the growth, we cannot ensure *ex-ante* that is significant and thus that the 2SLS works. For this reason, we discard the introduction of a spatial lag in the equation.

The results of the LM test are shown in the second column of Table 5. As we can observe, the conclusion is clear as well: there seems to be an error model behind the scenes generating the data we actually see. This points to be the case for both specifications (the one considering the initial size and the one using the average of both periods) as well as for absolute and relative populations. Last, due to the choice of the robust variant of the test, the conclusions about the existence of the error model remain the same even in presence of a lag model as well, which is important here, given the particularities outlined above. Accounting for the spatial error structure and analyzing whether it has any implication in the final conclusions we may draw about Gibrat's law are the aims of the next subsection.

6.2. Introducing a spatial error model in the Gibrat analysis

We now turn to the actual modeling of the spatial effects the ESDA suggested and some tests performed before seem to point out. The equations to be estimated are the same as in equations 4 to 7 but now we substitute the innovation term by the one specified in equation 12.

The procedure we use to estimate the model is that proposed in Kelejian and Prucha (2010). They introduce a new class of consistent GM (Generalized Moments) estimators for the autoregressive disturbance process that allows for heteroskedastic innovations and, unlike the version in Kelejian and Prucha (1998), permits statistical inference in the autoregressive term. Although in tests not reported we have not found evidence of heteroskedasticity, the possibility to examine the statistical significance of the error parameter has led us to adopt this method. The spatial counterpart of Table 3 is presented in Table 6. There are in particular two key points to comment about it.

The first one relates to the robustness of the significance of the spatial parameter. The estimate of ρ proves to be statistically different from zero in every specification, which can be interpreted as an argument to state there is actually a spatial error model for Gibrat's law and that space plays a role in urban growth

for Australia. The error term may be seen as a black box which captures to some extent the relevant variables that were not included in the regression as well as other measurement errors. In our estimates, the sign of such parameter is always positive, implying a process of positive spatial autocorrelation, which means similar values of the residuals tend to locate close to each other giving rise to a clustering pattern.

Table 6. GM Consistent Estimator with Heteroskedastik Disturbances

Variable	Coefficient	Std. error	t-Statistic	p-value	
$\ln S_{96}$					
$\hat{c}$	-0.07160020	0.02125529	-3.36858218	0.000774	***
$\hat{\beta}$	0.01627260	0.00306525	5.30874311	0.000000	***
$\hat{\rho}$	0.34144198	0.03663266	9.32069863	0.000000	***
$\ln \dfrac{(S_{01} + S_{96})}{2}$					
$\hat{c}$	-0.02075664	0.02148717	-0.96600127	0.334194	
$\hat{\beta}$	0.00907776	0.00308730	2.94035387	0.003327	***
$\hat{\rho}$	0.35340532	0.03617795	9.76852859	0.000000	***
$ln \dfrac{S_{96}}{S_{96}^{av}}$					
$\hat{c}$	-0.03688607	0.00815158	-4.52501874	0.000006	***
$\hat{\beta}$	0.00062126	0.00319667	0.19434565	0.845931	
$\hat{\rho}$	0.36342201	0.03580290	10.15063088	0.000000	***
$\ln \dfrac{(\frac{S_{01}}{S_{01}^{av}} + \frac{S_{96}}{S_{96}^{av}})}{2}$					
$\hat{c}$	-0.01876152	0.00853607	-2.19791072	0.028102	**
$\hat{\beta}$	0.00879073	0.00309128	2.84372391	0.004517	***
$\hat{\rho}$	0.35351934	0.03619188	9.76791865	0.000000	***

*** significant at 1% level; ** significant at 5% level; * significant at 10% level.
In the model specifications, the names refer to equations 4 to 7 and the term $\hat{\rho}$ to that in equation 12.

The second one is that, although space proves to be significant via the error term, the main conclusions we withdrew in the OLS section do not change. Initial on averaged urban size is still statistically significant and, hereby, Gibrat's law does not hold when we account for spatial effects in the error term. In fact, the coefficients look very similar from one specification to the other one, being the parameter for the size also greater than zero, which means a positive relationship between size and urban growth; again, bigger cities grow faster. This is important because it is not always the case: the estimation by OLS of a

model with spatially correlated errors may lead to conclude some parameters are significant when in fact they are not. However here, we have found how even when we consider such effects, the size of the city seems to be related to growth. We take these results as an extra argument for the deviation from Gibrat in Australia in the period analyzed.

7. CONCLUSIONS AND FUTURE STEPS

The present paper examines in both exploratory as well as confirmatory way the Australian urban system for the years 1996 and 2001. To do so, it uses the largest data set available so that the whole distribution (starting at 200 people) is covered and three different measures of the city size are used, namely the absolute population, the relative one and the size as a share of the total population. Also, for the confirmatory part, it applies some of the most recent developments in spatial econometrics. Australia is a very unique example of low population density, and its very special geography has shaped the distribution in a way that makes it very appealing for the urban researcher.

We first characterize the data set, and it already shows that, despite the short period chosen, some noticeable changes can be perceived. Then Zipf's analysis is carried out in order to test if the '*rank – size*' parameter is around one, but the evidence points to a much lower value (around 0.74), which implies a very uneven distribution of the population over the system and confirms what we had already sketched about Australia being a very unique case. A low parameter is related to a situation in which the sizes of the larger cities differ a great deal from the smaller ones. Moreover, we can observe how, from one year to another one, such coefficient has even decreased, deepening the inequality across cities.

After having realized Zipf does not hold for Australia, we examine the dynamic processes behind the city system to examine the relation between growth and size by means of Gibrat's law. We use different specifications to conclude that we cannot talk about *proportionate growth* in Australia for the period of study. City size coefficient always appears to be significant for population growth. In addition, there is a positive relation between both variables, so larger cities tend to grow faster than smaller ones. If we couple results from these two sections, we can visualize a system where, from 1996 to 2001, population has concentrated more in bigger cities, which increase their size more and more while the gap between them and the smaller ones tends to widen.

One of the main goals of this work is to explore the role space plays in urban population distribution and how the particular geography of Australia influences the outcome of its cities. The following step then has been to bring the spatial dimension into the analysis. To do so, we begin with an exploratory spatial data analysis (ESDA) procedure in which we try to determine whether there is any degree of global spatial dependence. The main conclusion is that although urban sizes are not spatially correlated, growth rates do show association in space. We then try to locate clusters of high and low growth by means of LISA indicators and an alternative visualization tool we call *ClusterCart*. This

step shows declining cities are located mainly in the inland while the growing centers tend to cluster around large cities by the coast.

We follow the exploratory by the confirmatory analysis. To do so we perform some tests and, when suggested, we estimate spatial models. First we analyze the residuals of the OLS regressions and that allows us to rule out the presence of any sort of spatial dependence in the Zipf's analysis. However, we reject the null of spatial randomness in the Gibrat's specification. We then perform the LM test for the existence of a spatial error, finding it relevant. The next step we take is to include an autoregresive structure in the error term. We estimate it with the GMM procedure suggested by Kelejian and Prucha (2010) and find the spatial parameter significant for every specification. However, this does not change the main conclusions regarding Gibrat's law.

The main picture we can draw after this study has a non-spatial and a spatial side, and we can find some similarities between both views. On the non-spatial world, Zipf's analysis sheds a very uneven distribution of people across the urban system, with city sizes more diverse than in other countries in the world. Moreover, Gibrat's approach sketches a link between larger size and faster growth. Once we go down on surface to the spatial world, we also find a very unequal distribution of cities across the Australian geography and a very unbalanced but space-led distribution of growth among cities.

In order to conclude the paper, here we suggest two directions which could be followed to expand the study of Australia. The first one would be to take the data set further back in time and the second one to dig into the causes which give rise to this outcome. Although they certainly give useful information, two points in time with a five-year lag in-between are certainly not enough to study long-term processes such as the evolution of a city system. That is why this paper should be seen rather as a static picture. Covering more years would bring the whole movie and would surely shed more light about the dynamical processes underlying the result pictured here. On the other hand, this study is rather descriptive in the sense that it focuses on characterizing Australia and on withdrawing systematic patterns in the way Australian cities are configured but it falls short in explaining why such trends and distributions are so. Mining possible explanatory variables to get deeper into the causes would surely be of great interest.

REFERENCES

Anderson, G. and Ge, Y. (2005). The size distribution of Chinese cities. *Regional Science and Urban Economics*, 35 (6): 756–776.

Anselin, L. (1988). *Spatial econometrics*. Kluwer Academic Publishers Dordrecht.

Anselin, L. (1995). Local indicators of spatial association-LISA. *Geographical Analysis*, 27 (2): 93–115.

Assuncao, R.M. and Reis, E.A. (1999). A new proposal to adjust Moran's I for population density. *Statistics in Medicine*, 18 (16): 2147-2161.

Beeson, P., DeJong, D., and Troesken, W. (2001). Population growth in US counties, 1840–1990. *Regional Science and Urban Economics*, 31 (6): 669–699.

Black, D. and Henderson, V. (2003). Urban evolution in the USA. *Journal of Economic Geography*, 3 (4): 343–372.

Bosker, M., Brakman, S., Garretsen, H., and Schramm, M. (2007). Looking for multiple equilibria when geography matters: German city growth and the WWII shock. *Journal of Urban Economics*, 61 (1): 152–169.

Bosker, M., Brakman, S., Garretsen, H., and Schramm, M. (2008). A century of shocks: The evolution of the German city size distribution 1925–1999. *Regional Science and Urban Economics*, 38 (4): 330-347.

Brakman, S., Garretsen, H., and Schramm, M. (2004). The Strategic Bombing of Cities in Germany in World War II and it Impact on City Growth. *Journal of Economic Geography*, 4 (1): 1–18.

Cliff, A. and Ord, J. (1981). *Spatial processes: models & applications*. Pion Ltd.

Cordoba, J. (2008). On the distribution of city sizes. *Journal of Urban Economics*, 63 (1): 177–197.

Davis, D. and Weinstein, D. (2002). Bones, Bombs, and Break Points: The Geography of Economic Activity. *American Economic Review*, 92 (5): 1269–1289.

Davis, D. and Weinstein, D. (2008). A Search for Multiple Equilibria in Urban Industrial Structure. *Journal of Regional Science*, 48 (1): 29-65.

Duranton, G. (2007). Urban Evolutions: The Fast, the Slow, and the Still. *American Economic Review*, 97 (1): 197–221.

Eaton, J. and Eckstein, Z. (1997). Cities and growth: Theory and evidence from France and Japan. *Regional Science and Urban Economics*, 27 (4-5): 443–474.

Edwards, R. W. (2001). Australian standard geographical classification. Technical report, Australian Bureau of Statistics.

Eeckhout, J. (2004). Gibrat's Law for (All) Cities. *American Economic Review*, 94 (5): 1429–1451.

Ellis, L. and Andrews, D. (2001). City sizes, housing costs, and wealth. *Reserve Bank of Australia Research Discussion Paper*, No 2001-08.

Gabaix, X. and Ioannides, Y. (2004). The Evolution of City Size Distributions. *Handbook of Regional and Urban Economics*, 4: 2341–2378.

Gonzalez-Val, R., Lanaspa, L., and Sanz, F. (2010). Gibrat's law for cities revisited. *MPRA paper* 10411, University Library of Munich, Germany.

Hugo, G. (2002). Changing Patterns of Population Distribution in Australia. Populations of New Zealand and Australia at the Millennium: A Joint Special Issue of the *Journal of Population Research and the New Zealand Population Review*, 1–22.

Junius, K. (1999). Primacy and Economic Development: Bell Shaped or Parallel Growth of Cities? *Journal of Economic Development*, 24 (1): 1–22.

Kelejian, H. and Prucha, I. (1998). A generalized spatial two-stage least squares procedure for estimating a spatial autoregressive model with autoregressive disturbances. *The Journal of Real Estate Finance and Economics*, 17 (1): 99–121.

Kelejian, H. and Prucha, I. (2010). Specification and estimation of spatial auto-regressive models with autoregressive and heteroskedastic disturbances. *Journal of Econometrics*, 157 (1): 53-67.

Lanaspa, L., Perdiguero, A., and Sanz, F. (2004). La distribucion del tamaño de las ciudades en España, 1900-1999. *Revista de Economía Aplicada*, 34: 5–16.

Lanaspa, L., Pueyo, F., and Sanz, F. (2003). The evolution of spanish urban structure during the twentieth century. *Urban Studies*, 40: 567–580.

Moran, P. (1948). The interpretation of statistical maps. *Journal of the Royal Statistical Society. Series B (Methodological)*, 10 (2): 243–251.

Moreno, R. and Vayá, E. (2000). *Técnicas econométricas para el tratamiento de datos especiales: la econometría espacial*. Barcelona (España). Ediciones de la Universidad de Barcelona.

Nitsch, V. (2003). Does history matter for urban primacy? The case of Vienna. *Regional Science and Urban Economics*, 33 (4): 401–418.

Overman, H. and Ioannides, Y. (2001). Cross-Sectional Evolution of the US City Size Distribution. *Journal of Urban Economics*, 49 (3): 543–566.

Rey, S. and Anselin, L. (2007). Pysal, a Python library of spatial analytical methods. *The Review of Regional Studies*, 37 (1): 5–27.

Rey, S. and Janikas, M. (2006). STARS: Space-Time Analysis of Regional Systems. *Geographical Analysis*, 38 (1): 67–86.

Rosen, K. and Resnick, M. (1980). The size distribution of cities: an examina- tion of the Pareto law and primacy. *Journal of Urban Economics*, 8 (2): 165–186.

Sharma, S. (2003). Persistence and stability in city growth. *Journal of Urban Economics*, 53 (2): 300–320.

Silverman, B. (1986). *Density estimation for statistics and data analysis*. Chapman & Hall/CRC.

Soo, K. (2005). Zipf 's Law for cities: a cross-country investigation. *Regional Science and Urban Economics*, 35 (3): 239–263.

Soo, K. (2007). Zipf 's Law and Urban Growth in Malaysia. *Urban Studies*, 44 (1): 1-14.

KANGOUROUS, VILLES ET ESPACE : UNE APPROCHE EXPLORATOIRE DU SYSTÈME URBAIN AUSTRALIEN

Résumé - *Cet article tente de dégager les caractéristiques du système urbain australien en 1996 et 2001. La première partie propose une approche descriptive du système urbain australien. La seconde partie examine la validité des lois de Zipf et de Gibrat pour les villes. Aucune de ces lois n'est validée. En s'appuyant sur les développements récents en économétrie spatiale (Heteroskedastic Consistent GM Estimation), une analyse exploratoire des données spatiales est utilisée pour étudier la répartition spatiale de la taille et de la croissance des villes.*

Mots-clés : SYSTÈME URBAIN AUSTRALIEN, LOI DE ZIPF, LOI DE GIBRAT, ANALYSE EXPLORATOIRE DES DONNÉES SPATIALES.