Gibrat's Law for Cities Revisited

González-Val, Rafael and Lanaspa, Luis and Sanz, Fernando

20 January 2010

Gibrat's Law for Cities Revisited

Rafael González-Val

Luis Lanaspa

Fernando Sanz

Universidad de Zaragoza

Abstract: The aim of this work is to test empirically the validity of Gibrat's Law in the growth of cities, using data for all the twentieth century of the complete distribution of cities (without any size restrictions) in three countries: the US, Spain and Italy. For this we use different techniques (parametric and non-parametric methods), obtaining mixed evidence, being the time horizon considered the key issue. In the short term, considered decade by decade, we find that growth was divergent in all three countries. Despite this, the distribution of growth in the cities can be approached as a lognormal. In the long term, first panel data unit root tests confirm the validity of Gibrat's Law in the upper tail distribution and, second, we find evidence in favour of a weak Gibrat's Law (size affects the variance of the growth process but not its mean) when using non-parametric methods which relate the growth rate to city size.

Keywords: Gibrat's Law, city size distribution, urban growth

JEL: R00, C14.

Address:

Dpto. de Análisis Económico, Universidad de Zaragoza, Facultad de CC. Económicas y Empresariales, Gran Vía, 2, 50005 Zaragoza (Spain), E-mail: llanas@unizar.es

1. Introduction

The relationship between the growth rate of a quantifiable phenomenon and its initial size is a question with a long history in statistics: do larger entities grow more quickly, or more slowly? On the other hand, perhaps no relationship exists and the rate is independent of size. A fundamental contribution to this debate is that of Gibrat (1931), who observed that the distribution of size (measured by sales or the number of employees) of firms could be approximated well with a lognormal, and that the explanation lay in the growth process of firms tending to be multiplicative and independent of their size. This proposition became known as Gibrat's Law and prompted a deluge of work exploring the validity of this Law for the distribution of firms (see the surveys of Sutton (1997) and Santarelli et al. (2006)). Gibrat's Law establishes that no regular behaviour of any kind can be deduced between growth rate and initial size.

The fulfilment of this empirical proposition also has consequences for the distribution which follows the variable; in the words of Gibrat (1931) himself "*the Law of proportionate effect will therefore imply that the logarithms of the variable will be distributed following the (normal distribution)*". Some years later Kalecki (1945), in a classic article, tested this statistical relationship between lognormality and proportionate growth under certain conditions, consolidating the conceptual binomial Gibrat's Law – lognormal distribution.

In the field of urban economics, Gibrat's Law, especially since the 1990s, has given rise to numerous empirical studies contrasting its validity for city size distributions, arriving at a majority consensus, though not absolute, that it holds in the long term. Gibrat's Law presents the added advantage that, as well as explaining relatively well the growth of cities, it can be related to another empirical regularity well known in urban economics, Zipf's Law, which appears when the so-called Pareto distribution exponent is equal to the unit[1]. The term was coined after a work by Zipf (1949), which observed that the frequency of the words of any language is clearly defined in statistical terms by constant values. This has given rise to theoretical works explaining the fulfilment of Gibrat's Law in the context of external urban local effects and productive shocks, relating them with Zipf's Law and associating them directly to an equilibrium situation. These theoretical works include Gabaix (1999), Duranton (2006, 2007), and Córdoba (2008).

Returning to the empirical side, there is an apparent contradiction in these studies, as they normally accept the fulfilment of Gibrat's Law but at the same time affirm that the distribution followed by city size is a Pareto distribution, very different to the lognormal. Recently, Eeckhout (2004) was able to reconcile both results, by demonstrating (as Parr and Suzuki (1973) affirm in a pioneering work) that, if size restrictions are imposed on the cities, taking only the upper tail, this skews the analysis. Thus, if all cities are taken, it can be found that the true distribution is lognormal, and that the growth of these cities is independent of size. However, to date, Eeckhout (2004) is the only study to consider the entire city size distribution. But this is a short term analysis[2], when the phenomenon under study (Gibrat's Law) is, by definition, a long term result.

The aim of this work is to test empirically the validity of Gibrat's Law in the growth of cities, using data for all the twentieth century of the complete distribution of

---

[1] If city size distribution follows a Pareto distribution the following expression can be deduced: $\ln R = a - b \cdot \ln S$, where R is rank (1 for the biggest city, 2 for the second biggest and so on), S is the size or population and a and b are parameters, this latter being known as the Pareto exponent. Zipf's Law is fulfilled when b equals the unit.

[2] Eeckhout (2004) takes data from the United States census of 1990 and 2000, possibly because they are the only ones to be available online. Levy (2009), in a comment to Eeckhout (2004), and Eeckhout (2009) in the reply, also consider no truncation point, but only for the 2000 US Census data.

cities (without any size restrictions or with no truncation point) in three countries: the US, Spain and Italy. The following section offers a brief overview of the literature on Gibrat's Law and cities and the results obtained. Section 3 presents the databases, with special attention to the US census.

From the results we deduce that, when we consider the complete distribution of cities in the short term (Section 4), a tendency to divergence is seen. However, the empirical evidence (Section 5) shows that this does not impede city size distribution being adequately approximated as a lognormal distribution. Finally, in Section 6 a long term viewpoint is taken. Panel data unit root tests confirm the validity of Gibrat's Law in the upper tail distribution (Section 6.1), and we find evidence in favour of a weak Gibrat's Law (size affects the variance of the growth process but not its mean) when using non-parametric methods which relate growth rate with city size (section 6.2). The work ends with our conclusions.

2. Gibrat's Law for cities. An overview of the literature

In the 1990s numerous studies began to appear which empirically tested the validity of Gibrat's Law. Table 1 shows the classification of all the studies on urban economics that we know of. While the countries considered, statistical and econometric techniques used and sample sizes are heterogeneous, the predominating result is the acceptance of Gibrat's Law.

Thus, both Eaton and Eckstein (1997) and Davis and Weinstein (2002) accept its fulfilment for Japanese cities, although they use different sample sections (40 and 303 cities, respectively), and time horizons. Davis and Weinstein (2002) affirm that long-run city size is robust even to large temporary shocks and, in studying the effect of Allied bombing in the Second World War, deduce that the effect of these temporary shocks disappears completely in less than 20 years.

Brakman et al. (2004) come to the same conclusion when analysing the impact of bombardment on Germany during the Second World War, concluding that, for the sample of 103 cities examined, bombing had a significant but temporary impact on post-war city growth. Nevertheless, nearly the same authors in Bosker et al. (2008) obtain a mixed result with a sample of 62 cities in West Germany: correcting for the impact of WWII, Gibrat's Law is found to hold only for about 25% of the sample.

Meanwhile, both Clark and Stabler (1991) and Resende (2004) also accept the hypothesis of proportionate urban growth for Canada and Brazil respectively. The sample size used by Clark and Stabler (1991) is tiny (the 7 most populous Canadian cities), although the main contribution of their work is to propose the use of data panel methodology and unit root tests in the analysis of urban growth. This is also the methodology which Resende (2004) applies to his sample of 497 Brazilian cities. However, Henderson and Wang (2007) strongly reject Gibrat's Law and a unit root process in their worldwide data set on all metro areas over 100,000 from 1960 to 2000.

For the case of the US, there are also several works accepting statistically the fulfilment of Gibrat’s Law, whether at the level of cities (Eeckhout (2004) is the first to use the entire sample without size restrictions), or with MSAs (Ioannides and Overman (2003), whose results reproduce Gabaix and Ioannides, 2004). Also for the US, however, Black and Henderson (2003) reject Gibrat’s Law for any sample section, although their database of MSAs is different[3] to that used by Ioannides and Overman (2003).

[3] The standard definitions of metropolitan areas were first published in 1949 by what was then called the Bureau of the Budget, predecessor of the current Office of Management and Budget (OMB), with the

Other works exist rejecting the fulfilment of Gibrat's Law. Thus, Guérin-Pace (1995) finds that in France for a wide sample of cities with over 2,000 inhabitants during the period 1836-1990 there appears to be a fairly strong correlation between city size and growth rate, a correlation which is accentuated when the logarithm of the population is considered. This result goes against that obtained by Eaton and Eckstein (1997) when considering only the 39 most populated French cities. Soo (2007) and Petrakos et al. (2000) also reject the fulfilment of Gibrat's Law in Malaysia and Greece, respectively.

For the case of China, Anderson and Ge (2005) obtain a mixed result with a sample of 149 cities of more than 100,000 inhabitants: Gibrat's Law appears to describe the situation well prior to the Economic Reform and One Child Policy period, but later Kalecki's reformulation seems to be more appropriate.

What we wish to emphasize is that, with the exception of Eeckhout (2004), none of these studies considers the entire distribution of cities, as all of them impose a truncation point, whether explicitly, by taking cities above a minimum population threshold or implicitly, by working with MSAs[4]. This is usually due to a practical reason of data availability. For this reason most studies focus on analysing the most populous cities, the upper tail distribution. There are two very reasonable justifications for this approach. First, the largest cities represent most of the population of a country. And second, the growth rate of the biggest cities has less variance than the smallest ones (scale effect).

---

designation Standard Metropolitan Area. This means that if the objective is making a long term analysis it will be necessary to reconstruct the areas for earlier periods, in the absence of a single criterion.

[4] In the US, to qualify as a MSA a city needs to have 50,000 or more inhabitants, or the presence of an urbanised area of at least 50,000 inhabitants, and a total metropolitan population of at least 100,000 (75,000 in New England), according to the OMB definition. In other countries similar criteria are followed, although the minimum population threshold needed to be considered a metropolitan area may change.

However, it should be pointed out that any test done on this type of sample will be local in character, and the behaviour of large cities cannot be extrapolated to the entire distribution. This type of deduction can lead to erroneous conclusions, as it must not be forgotten that what is being analysed is the behaviour of a few cities, which as well as being of a similar size, can present common patterns of growth. Therefore, we might conclude that Gibrat's Law is fulfilled when in fact we have focused our analysis on a club of cities which cannot be representative of all urban centres.

3. The databases

We use city population data from three countries: the US, Spain and Italy. The US is an extremely interesting country in which to analyse the evolution of urban structure, as it is a relatively young country whose inhabitants are characterised by high mobility. On the other hand we have the European countries, with a much older urban structure and inhabitants who present greater resistance to movement; specifically, Cheshire and Magrini (2006) estimate mobility in the US is fifteen times higher than in Europe.

Considering these two types of country gives us information about different urban behaviours, as while Spain and Italy have an already consolidated urban tissue and new cities are rarely created (urban growth is produced by population increase in existing cities), in the US urban growth has a double dimension: as well as increases in city size, the number of cities also increases, with potentially different effects on city size distribution. Thus, the population of cities (incorporated places) goes from representing less than half the total population of the US in 1900 (46.99%) to 61.49% in

2000; at the same time the number of cities increases by 82.11%, from 10,596 in 1900 to 19,296 in 2000.

The data for the US we are using are the same as those used by González-Val (2009). Our base, created from the original documents of the annual census published by the US Census Bureau, www.census.gov, consists of the available data of all incorporated places without any size restriction, for each decade of the twentieth century. The US Census Bureau uses the generic term *incorporated place* to refer to the governmental unit incorporated under state Law as a city, town (except in the states of New England, New York and Wisconsin), borough (except in Alaska and New York), or village, and which has legally established limits, powers and functions.

The number of cities (in brackets) corresponding to each period is: 1900 (10,596 cities), 1910 (14,135), 1920 (15,481), 1930 (16,475), 1940 (16,729), 1950 (17,113), 1960 (18,051), 1970 (18,488), 1980 (18,923), 1990 (19,120), and 2000 (19,296).

Two details should be noted. First, that all the cities corresponding to Alaska, Hawaii, and Puerto Rico for each decade are excluded, as these states were annexed during the 20th century (Alaska and Hawaii in 1959, and the special case of Puerto Rico, which was annexed in 1952 as an associated free state), and data do not exist for all periods. Their inclusion would produce geographical inconsistency in the samples, which would not be homogenous in geographical terms and thus could not be compared. And second, for the same reason we also exclude all the unincorporated places (concentrations of population which do not form part of any incorporated place, but which are locally identified with a name), which began to be accounted after 1950. However, these settlements did exist earlier, so that their inclusion would again present a problem of inconsistency in the sample. Also, their elimination is not quantitatively important; in fact there were 1,430 unincorporated places in 1950, representing 2.36% of the total population of the US, which by 2000 would be 5,366 places and 11.27%.

For Spain and Italy the geographical unit of reference is the municipality and the data comes from the official statistical information services. In Italy this is the Servizio Biblioteca e Servizi all'utenza, of the Direzione Centrale per la Diffusione della Cultura e dell'informazione Statistica, part of the Istituto Nazionale di Statistica, www.istat.it, and for Spain we have taken the census of the Instituto Nacional de Estadística[5], INE, www.ine.es. The de facto resident population has been taken for each city.

We have taken the data corresponding to the census of each decade of the 20th century. For Italy data for the following years have been considered (in brackets, the number of cities for each year): 1901 (7,711), 1911 (7,711), 1921 (8,100), 1931 (8,100), 1936 (8,100), 1951 (8,100), 1961 (8,100), 1971 (8,100), 1981 (8,100), 1991 (8,100), and 2001 (8,100). No census exists in Italy for 1941, due to its participation in the Second World War, so we have taken the data for 1936. For Spain the following years are considered: 1900 (7,800), 1910 (7,806), 1920 (7,812), 1930 (7,875), 1940 (7,896), 1950 (7,901), 1960 (7,910), 1970 (7,956), 1981 (8,034), 1991 (8,077), and 2001 (8,077).

4. Gibrat's Law in the short term: convergent, parallel, vs. divergent city growth processes

In this section we offer a first approach to the behaviour of city growth from a short term perspective, i.e., considering each decade individually. Following Gabaix and

[5] The official INE census have been improved in an alternative database, created by Azagra et al. (2006), reconstructing the population census for the twentieth century using territorially homogeneous criteria. We have repeated the analysis using this database and the results are not significantly different, so we have presented the results deduced from the official data.

Ioannides (2004), Gibrat's Law states that the growth rate of an economic entity (firm, mutual fund, city) of size S has a distribution function with mean and variance that are independent of S. Therefore, if S_{it} is the size of city i at the time t and g is its growth rate, then $S_{it} = S_{it-1}(1+g)$. Taking logarithms and adding that the rate depends on the initial size, we can obtain the following general expression of the growth equation[6,7]:

$$\ln S_{it} - \ln S_{it-1} = \mu + \beta \ln S_{it-1} + u_{it}, \qquad (1)$$

where $\mu = \ln(1+g)$ and u_{it} is a random variable representing the random shocks which the growth rate may suffer, which we shall suppose to be identically and independently distributed for all cities, with $E(u_{it}) = 0$ and $Var(u_{it}) = \sigma^2 \ \forall i,t$. If $\beta = 0$ Gibrat's Law holds and we obtain that growth is independent of the initial size.

In this case, ($\beta = 0$), it is easy to prove that the expected value of the size of city i at the time t depends only on the number of periods which have passed and on size in the first period:

$$E(\ln S_{it}) = \mu \cdot t + \ln S_{i0}, \qquad (2)$$

while the variance would be given by:

$$Var(\ln S_{it}) = t \cdot \sigma^2. \qquad (3)$$

---

[6] The size of a city can be defined, according to the literature, in three ways: in levels (S_{it}), in relative values ($S_{it}/\overline{S}_t$, $\overline{S}_t$ is the mean size) or in shares ($S_{it}/\sum_i S_{it}$). The crucial parameter in (1) is β, which determines whether Gibrat's Law holds. The specification (1) in logs makes the estimation of β robust to the three different definitions of city size.

[7] Taking logarithms we reduce the distortions that may occur in the mean and variance of the growth rate due to changes in the scale of the variable.

Consequently, the mean grows over time, and variance does too. The increased variance over time is consistent with the prediction of a Brownian motion: proportionate growth leads to a lognormal distribution with standard deviation that is increasing in time t.

We adopt the Eaton-Eckstein terminology of convergent, parallel, vs. divergent city growth processes. Remember that if $\beta = 0$ city growth is parallel, as it does not depend on initial size. Thus, if the estimation of β is significantly different to zero we will reject the fulfilment of Gibrat's Law. In the case of being greater than zero, we will have divergent growth, because city growth would depend directly and positively on initial size. A sustained process of divergent growth of this kind would result in an increasingly asymmetrical distribution, with small cities getting further and further away from large ones. And if β is negative urban growth would be convergent, as the growth-size ratio would be negative; a larger initial population would mean less growth and vice versa, so that in the long term distribution would tend to be concentrated around a median value. It is simple to prove that when $\beta \neq 0$ expressions (2) and (3) change, becoming

$$E(\ln S_{it}) = \mu \cdot \frac{(\beta+1)^t - 1}{\beta} + (\beta+1)^t \ln S_{i0}, \qquad (4)$$

$$Var(\ln S_{it}) = \sigma^2 \cdot \frac{(\beta+1)^{2t} - 1}{\beta^2 + 2\beta}, \qquad (5)$$

and it can be demonstrated (see Appendix) that when $t > 1$ and growth is divergent $(\beta > 0)$ variance (5) grows even faster than in (3), while if city growth were convergent $(\beta < 0)$ variance (5) would be less than in (3).

The first result we wish to present is the estimation of equation (1). We will focus on the analysis of the estimation of parameter β, as whether Gibrat's Law is fulfilled or not depends on its significance and its sign. Table 2 shows the results of the OLS estimation of β for each decade in the three countries considering all the cities, without size restrictions. The results of these regressions are usually heteroskedastic, so we have calculated the t-ratios using White's (1980) Heteroskedasticity-Consistent Standard Errors.

The first conclusion we obtain is that when the entire sample of cities is considered, β is always significantly different to zero, for any period and in the three countries. This result is robust as, while the literature usually admits the possibility of occasional deviations from Gibrat's Law in the short term (with some periods in which urban growth may be convergent or divergent), we are rejecting the fulfilment of Gibrat's Law for each decade of the 20th century and for three nations. But the really surprising finding is that despite these three countries having such different urban structures and histories, the estimated parameter is always positive (except in the period 1970-1980 in the US), so that the three exhibit divergent behaviour throughout the 20th century.

The exception to this process of divergence is the estimation obtained for the US in the decade 1970-1980. The fact that this parameter is negative shows that during this decade the most populous cities grew more slowly. However, this result is atypical, and reflects two demographical circumstances in the United States during this period. First, between 1960 and 1990 there was a decline in the growth of the total population of the US, going from a growth rate of 18.5% in 1950-1960 to 9.8% in 1980-1990[8]. Then, that the total population grew by only 11.4% in 1970-1980, the third lowest growth rate in

[8] Source: http://www.census.gov/population/censusdata/table-4.pdf.

the history of the US since the first census was published in the late 18th century. And in this context of low growth of the total population, the percentage of urban population also fell (understood now as the percentage of the population associated with incorporated places), going from 64.51% of the total population in 1970 to 61.78% in 1980, which is by far the biggest fall in the 20th century. The fact that our estimation of β is negative would reflect the cities in the upper half of the distribution being where growth slowed most.

We have obtained that, in the short term, the city growth process was divergent in the three countries. However, this conclusion can change in the long term. But first we will analyse in section 5 the consequences on city size distribution of the divergent tendency we have observed.

5. What about city size distribution? Lognormality is maintained

In the section above, it has been shown that the overall result in the short term when the whole distribution is used is divergence. Also, as $\beta > 0$ the variance will grow more than linearly (equation (5)), so that the growth process could be explosive and it would be expected that city size distribution would be increasingly asymmetrical. But our results show that the growth process lead to a lognormal distribution with standard deviation that is increasing in time t (as a Brownian motion would predict) in the three countries.

We carried out Wilcoxon's lognormality test (rank-sum test), which is a non-parametric test for assessing whether two samples of observations come from the same distribution. The null hypothesis is that the two samples are drawn from a single population, and therefore that their probability distributions are equal, in our case, the lognormal distribution. Wilcoxon's test has the advantage of being appropriate for any sample size. The more frequent normality tests –Kolmogorov-Smirnov, Shapiro-Wilks, D'Agostino-Pearson– are designed for small samples, and so tend to reject the null hypothesis of normality for such large sample sizes, although the deviations from lognormality are arbitrarily small.

Table 3 shows the results of the test. The conclusion is that the null hypothesis of lognormality is accepted at 5% for all periods of the 20th century in Spain and Italy. In the US a temporal evolution can be seen; in the first decades lognormality is rejected and the p-value decreases over time, but from 1930 the p-value begins to grow until lognormal distribution is accepted at 5% from 1960 onwards (the same conclusion is reached by González-Val (2010) through a graphic examination of the adaptive kernels corresponding to the estimated distribution of different decades). In fact, if instead of 5% we take a significance level of 1%, the null hypothesis would only be rejected in 1920 and 1930.

However, the shape of the distribution in the US for the period 1900-1950 is not far from lognormality, either. Figure 1 shows the empirical density functions estimated by adaptive Gaussian kernels for 1900 and for 1950 (the last in which lognormality is rejected). The motive for this systematic rejection appears to be an excessive concentration of density in the central values, higher than would correspond to the theoretical lognormal distribution (in black). Starting in 1900 with a very leptokurtic distribution, with a great deal of density concentrated in the mean value, from 1930 (not shown), when the growth of urban population slows, the distribution loses kurtosis and concentration decreases, accepting lognormality statistically at 5% from 1960.

To sum up, both the test carried out and the visualisation of the estimated empirical density functions seem to corroborate that city size distribution can be approximated correctly as a lognormal (in Spain and Italy during the entire 20th century, and in the US for most decades, depending on the significance level), despite urban growth having been divergent every decade during the entire 20th century for the three countries (with the single exception of the period 1970-1980 in the US).

6. Gibrat's Law in the long term

In this section, we change our temporal perspective to the long term (the entire twentieth century). In order to carry out this analysis, we transform city population $\left(S_{it}\right)$ to city relative size $\left(s_{it}\right)$, defined as $s_{it} = \frac{S_{it}}{\bar{S}_t} = \frac{S_{it}}{\frac{1}{N}\sum_{i=1}^{N} S_{it}}$, as in a long term temporal perspective of steady state distributions it is necessary to use a relative measure of size.

This approach is more interesting, as the phenomenon under study (Gibrat's Law) is, by definition, a long term result. For this we combine parametric methods (the panel dimension of our data has been exploited in order to test for a unit root) with non-parametric ones, enabling us to study the relationship of growth and the variance of growth with city size.

6.1. Parametric analysis: panel unit root testing

Clark and Stabler (1991) suggested that testing for Gibrat's Law is equivalent to testing for the presence of a unit root. This idea has also been emphasized by Gabaix and Ioannides (2004) who expect "*that the next generation of city evolution empirics could draw from the sophisticated econometric literature on unit roots*". In line with this suggestion most studies now apply unit root tests (see Table 1).

Some authors (Black and Henderson, 2003; Henderson and Wang, 2007; Soo, 2007) test the presence of a unit root by proposing a growth equation similar to our equation (1), which they estimate using panel data. Nevertheless, as pointed out by Gabaix and Ioannides (2004) and Bosker et al. (2008), this methodology presents some drawbacks. First, the periodicity of our data is by decades, and we have only 11 temporal observations (decade-by-decade city sizes over a total period of 100 years), when the ideal would be to have at least annual data. And second, the presence of cross-sectional dependence across the cities in the panel can give rise to estimations which are not very robust. It has been well established in the literature that panel unit root and stationarity tests that do not explicitly allow for this feature among individuals. (Banerjee et al. 2005).

For this, we use one of the tests especially created to deal with this question: Pesaran's (2007) test for unit roots in heterogeneous panels with cross-section dependence is calculated on the basis of the CADF statistic (cross-sectional augmented ADF statistic).

To eliminate the cross dependence, the standard Dickey-Fuller (or Augmented Dickey-Fuller, ADF) regressions are augmented with the cross section averages of lagged levels and first-differences of the individual series, such that the influence of the unobservable common factor is asymptotically filtered.

The test of the unit root hypothesis is based on the t-ratio of the OLS estimate of b_i in the following cross-sectional augmented DF (CADF) regression:

$$\Delta y_{it} = a_i + b_i y_{i,t-1} + c_i \bar{y}_{t-1} + d_i \Delta \bar{y}_t + e_{it}. \qquad (6)$$

We will test for the presence of a unit root in the natural logarithm of city relative size $(y_{it} = \ln s_{it})$ taking this into account. Null hypothesis assumes that all series are non-stationary, and Pesaran's CADF is consistent under the alternative that only a fraction of the series is stationary.

However, the problem with Pesaran's test is that it is not designed to deal with such large panels (22,078 cities in the US, 8,077 in Spain and 8,100 in Italy), especially when so few temporal observations are available $(N \to \infty, T = 11)$. For this reason, we must limit our analysis to the largest cities (although the next section does offer a long term analysis of the entire sample).

Table 4 shows the results of Pesaran's (2007) test, both the value of the test statistic and the corresponding p-value, applied to the upper tail distribution until the 500 largest cities in the initial period have been considered. All statistics are based on univariate AR(1) specifications including constant and trend.

The null hypothesis of a unit root is not rejected in the US or Italy for any of the sample sizes considered, providing evidence in favour of the long term validity of Gibrat's Law. Spain's case is different, as when the sample size is more than the 200 largest cities, the unit root is rejected, indicating a relationship between relative size and growth rate even for the largest cities.

6.2. Non-parametric analysis: kernel regression conditional on city size

This section on the nonparametric analysis follows closely the analysis in Ioannides and Overman (2003), and Eeckhout (2004). It consists of taking the following specification:

$$g_i = m(s_i) + \varepsilon_i , \qquad (7)$$

where g_i is the growth rate $(\ln s_{it} - \ln s_{it-1})$ normalised (subtracting the mean and dividing by the standard deviation) and s_i is the logarithm of the ith city relative size. Instead of making suppositions about the functional relationship m, $\hat{m}(s)$ is estimated as a local mean around the point s and is smoothed using a kernel, which is a symmetrical, weighted and continuous function in s.

To analyse all the 20th century we build a pool with all the growth rates between two consecutive periods. This enables us to carry out long term analysis. And the Nadaraya-Watson method is used, exactly as it appears in Härdle (1990), based on the following expression[9]:

$$\hat{m}(s) = \frac{n^{-1}\sum_{i=1}^{n} K_h(s - s_i) g_i}{n^{-1}\sum_{i=1}^{n} K_h(s - s_i)}, \qquad (8)$$

where K_h denotes the dependence of the kernel K (in this case an Epanechnikov) on the bandwidth h. We use the same bandwidth (0.5) in all estimations in order to permit comparisons between countries.

Starting from this calculated mean $\hat{m}(s)$, the variance of the growth rate g_i is also estimated, again applying the Nadaraya-Watson estimator:

$$\hat{\sigma}^2(s) = \frac{n^{-1}\sum_{i=1}^{n} K_h(s - s_i)(g_i - \hat{m}(s))^2}{n^{-1}\sum_{i=1}^{n} K_h(s - s_i)}. \qquad (8)$$

[9] The calculation was done with the KERNREG2 Stata module, developed by Nicholas J. Cox, Isaias H. Salgado-Ugarte, Makoto Shimizu and Toru Taniuchi, and available online at: http://ideas.repec.org/c/boc/bocode/s372601.html.

The estimator is very sensitive, both in mean and in variance, to atypical values. For this reason we decide to eliminate from the sample the 5% smallest cities, as they usually have much higher growth rates in mean and in variance. This is logical; we are discussing cities of under 200 inhabitants, where the smallest increase in population is very large in percentage terms.

Gibrat's Law implies that growth is independent of size in mean and in variance. As growth rates are normalised, if Gibrat's Law in mean were strictly fulfilled, the nonparametric estimate would be a straight line on the zero value. Values different to zero involve deviations from the mean. And the estimated variance of the growth rate would also be a straight line in the value one, which would mean that the variance does not depend on the size of the variable analysed. To be able to test these hypotheses, we have constructed bootstrapped 95-percent confidence bands (calculated from 500 random samples with replacement).

Figure 2 shows the nonparametric estimates of the growth rate of a pool for the entire 20th century for the US (1900-2000, 152,475 observations), Spain (1900-2001, 74,100 observations) and Italy (1901-2001, 73,260 observations). For the US the value zero is always in the confidence bands, so that it cannot be rejected that the growth rates are significantly different for any city size. For Spain and Italy the estimated mean grows with the sample size, although it is significantly different to zero only for the largest cities[10]. One possible explanation is historical: both Spain and Italy suffered wars on their territories during the 20th century, so that for several decades, the largest cities attracted most of the population[11]. Therefore, we find evidence in favour of

[10] In the case of Spain, this divergent behaviour could be the explanation for the rejection obtained in the previous section of the unit root null hypothesis.
[11] This result can be related with the "safe harbour effect" of Glaeser and Shapiro (2002), which is a centripetal force which tends to agglomerate the population in large cities when there is an armed conflict.

Gibrat's Law for the US throughout the 20th century. Also for Spain and Italy, although the largest cities would present some divergent behaviour.

Figure 2 also shows the nonparametric estimates of the variance of growth rate of a pool for the entire 20th century for the US, Spain and Italy. As expected, while for most of the distribution the value one falls within the confidence bands, indicating that there are no significant differences in variance, the tails of the distribution show differentiated behaviours. In the US the variance clearly decreases with the size of the city, while in Spain and Italy the behaviour is more erratic and the biggest cities also have high variance.

Our results, obtained with our sample of all incorporated places without any size restriction, are similar to those obtained by Ioannides and Overman (2003), with their database of the most populous MSAs. To sum up, the nonparametric estimates show that while the mean of growth (Gibrat's Law for means) seems to be independent of size in the three countries (although in Spain and Italy the largest cities would present some divergent behaviour), the variance of growth (Gibrat's Law for variances) does depend negatively on size: the smallest cities present clearly higher variance in all three countries (although in Spain and Italy the behaviour is more erratic and the biggest cities also have high variance).

This points to Gibrat's Law holding weakly (growth is proportional in means but not in variance). Gabaix (1999) contemplates this possibility, that Gibrat's Law might not hold exactly, and examines the case in which cities grow randomly with expected growth rates and standard deviations that depend on their sizes. Therefore, the size of city i at time t varies according to:

$$\frac{dS_t}{S_t} = \mu(S_t)dt + \sigma(S_t)dB_t,$$

where $\mu(S)$ and $\sigma^2(S)$ denote, respectively, the instantaneous mean and variance of the growth rate of a size S city, and B_t is a standard Brownian motion. Córdoba (2008) also introduces a parsimonious generalization of Gibrat's Law that allows size to affect the variance of the growth process but not its mean.

Nevertheless, we must distinguish between the American and European cases, as Gibrat's Law assumes a fixed and invariant number of locations. The number of cities remains almost constant in Spain and Italy, but the same is not true of the US; between the start of the sample and the end, the number of cities doubles. And while a Brownian motion can be adjusted to include new entrants, the distribution from which the entrants are drawn and the magnitude of entrants will affect the distribution. In particular, in the presence of a drift (as in this case where there is average city growth), the distribution from which new entrants are drawn is unlikely to be stationary if one wants to obtain the result that growth is proportionate.

So, Figure 3 shows the nonparametric estimates of the growth rate and of the variance of growth rate of a pool for the entire 20th century for the US (1910-2000, 59,865 observations) considering only the new entrant cities since 1910. Bootstrapped 95-percent confidence bands are also presented. The estimations show how the cities entering the sample from 1910 usually had growth rates which were higher on average and in variance than the average of the entire sample (dotted blue line), although the bands do not permit us to reject their being significantly different. The differences in variance indicate that part of the increased variance at the bottom of the size distribution can be explained by the cities which entered the distribution throughout the twentieth century.

Also, Figure 4, representing the empirical estimated distributions of entrant cities in 1910 and 2000 (normalized by the average size of the cohort of the entire distribution), shows the change in distribution of entrant cities. Starting from a very leptokurtic distribution in 1910 (more leptokurtic than the distribution of the whole sample) concentration decreases until the 2000 distribution is very similar to lognormal.

7. Conclusions

The aim of this work is very simple: to provide additional information on the fulfilment of Gibrat's Law, an empirical regularity which is well known in the literature on Urban Economics. In a nutshell, this Law states that the population growth rate of cities is a process deriving from independent multiplicative shocks, so that two conclusions can be statistically deduced. First, if we take logarithms city size distribution can be well fitted by a lognormal; second, the growth rate is on average independent of the initial size of the urban centers and its evolution is fundamentally stochastic, without any fixed pattern of behaviour. Moreover, although this problem is not dealt with here, if the urban growth process does follow Gibrat's Law this has some implications for the theory, as demonstrated in the excellent survey by Gabaix and Ioannides (2004).

This article contributes in two ways. On the one hand, it uses a database covering three countries (the US, Spain and Italy), with different urban histories, for the entire 20th century. As far as we know, this is the widest-ranging attempt to test the geographical and temporal validity of this Law, focusing on robust results. On the other, it does not use a unique statistical or econometric technique, but looking to find greater robustness in alternative specifications, employs different methods (parametric and non-parametric) and complementary approaches.

There are three basic conclusions, the first two being more important.

First, as shown in Section 2, until now evidence for the fulfilment of Gibrat's Law for cities has been mixed, with a slight predominance of papers defending its fulfilment. This article argues that the time horizon considered is a key issue. In the short term, using decade by decade census data individually, growth in the three countries during the century is divergent: there is a positive correlation between the growth rate and the initial size of the cities. This result, somewhat unexpectedly found consistently in three nations over a hundred years, does not impede that the corresponding contrasts, Wilcoxon's rank sum tests, show that, except for the US in the first half of the century, the lognormal distribution is systematically never rejected. How can these apparently contradictory findings be reconciled? Simply, because as the definition of Gibrat's Law makes clear, this is a long term phenomenon, and the empirical evidence accumulated in these pages shows this to be so, although with some nuances. This fact defines the second conclusion.

Second, therefore, the panel data unit root tests carried out confirm that, in the long term, Gibrat's Law always holds for the upper tail of the distribution for the US and Italy, and only for the two hundred largest cities for Spain. In any case, the use of panel techniques for three countries and eleven census periods is innovative and generates, we believe, important conclusions. Moreover, from the use of non-parametric techniques, also over the long term, such as kernel regressions conditional on city size, we deduce that Gibrat's Law for means is completely fulfilled for the three countries, while for variances the predominant behaviour is, in turn, consonant with the Law, except for the largest and smallest cities, depending on the country.

Finally, the case of the US differs in that the number of cities doubles over the twentieth century. The new entrant cities present higher growth rates in means and in variance than the average for the whole sample, although we cannot reject their being significantly different. The differences are greater in variance, indicating that part of the increased variance at the bottom of the size distribution can be explained by the cities which entered the distribution throughout the twentieth century.

Appendix: Variance and convergent, parallel, vs. divergent city growth processes

We have two expressions:

$$Var(\ln S_{it}) = t \cdot \sigma^2 \qquad (3)$$

$$Var(\ln S_{it}) = \sigma^2 \cdot \frac{(\beta+1)^{2t} - 1}{\beta^2 + 2\beta} \qquad (5)$$

If Gibrat's Law is fulfilled ($\beta = 0$), and applying L'Hôpital's rule we obtain that (5) converges to (3): $\lim_{\beta \to 0} \left(\sigma^2 \cdot \frac{2t(\beta+1)^{2t-1}}{2\beta + 2} \right) = \frac{2t\sigma^2}{2} = t\sigma^2$.

Let's see what happens if $\beta > 0$ or $\beta < 0$:

$$(3) - (5) = t \cdot \sigma^2 - \sigma^2 \cdot \frac{(\beta+1)^{2t} - 1}{\beta^2 + 2\beta} = \frac{\sigma^2}{\beta(\beta+2)}\left[t(\beta^2 + 2\beta) - (\beta+1)^{2t} + 1\right] = \frac{\sigma^2}{\beta(\beta+2)}[f(\beta)]$$

Considering time t as a continuum beginning with zero, the expression between brackets $f(\beta)$ is only defined if $-1 < \beta$. Also, if $\beta > 0$ then $\frac{\sigma^2}{\beta(\beta+2)} > 0$, while if $-1 < \beta < 0$ then $\frac{\sigma^2}{\beta(\beta+2)} < 0$.

Therefore, to find out the total sign of the difference $(3) - (5)$ we must study the behaviour of the function $f(\beta) = t(\beta^2 + 2\beta) - (\beta+1)^{2t} + 1$. The maximum or minimum of this function is obtained by defining the corresponding optimisation problem, whose first order condition is given by:

$$\frac{df(\beta)}{d\beta} = f'(\beta) = 2t\left(\beta + 1 - (\beta + 1)^{2t-1}\right) = 0 ,$$

from which we deduce that at the extreme $1 = (\beta + 1)^{2t-2}$, which means that $f(\beta)$ is maximum or minimum in $\beta = 0$. In order to know if the optimum $\beta = 0$ is maximum or minimum we obtain the second order condition:

$$\frac{d^2 f(\beta)}{d\beta^2} = f''(\beta) = 2t\left(1 - (2t-1)(\beta + 1)^{2t-2}\right),$$

and evaluate the sign in $\beta = 0$: $f''(\beta = 0) = 4t(1-t) < 0$ as long as $t > 1$.

Thus, we already know that the function $f(\beta)$ is concave and reaches its maximum in $\beta = 0$ as long as $t > 1$. Considering that $f(0)$=0, this function always takes negative values except in the maximum.

The final sign of the difference $(3) - (5)$ will be (maintaining the conditions $-1 < \beta$ and $t > 1$):

1. When $\beta > 0$ we have seen that $\dfrac{\sigma^2}{\beta(\beta + 2)} > 0$ is fulfilled and city growth is divergent. The variance of the cities will be bigger than if Gibrat's Law were fulfilled: $(3) < (5)$.

2. When $\beta < 0$ city growth is convergent. The variance of the cities will be less than if Gibrat's Law were fulfilled: $(3) > (5)$.

3. When $\beta = 0$ (3)=(5).

Acknowledgements

The authors would like to thank the Spanish Ministerio de Educación y Ciencia (SEJ2006-04893/ECON project and AP2005-0168 grant from FPU programme), the DGA (ADETRE research group) and FEDER for their financial support. The comments of Arturo Ramos and members of the ADETRE research group contributed to improving the paper. Earlier versions of this paper were presented at the 55th North American Meeting of the Regional Science Association International (New York, 2008), at the XXXIII Symposium of Economic Analysis (Zaragoza, 2008), at the 67th International Atlantic Economic Conference (Rome, 2009), at the 24th Annual Congress of the European Economic Association (Barcelona, 2009), and at the XXXV Reunión de Estudios Regionales (Valencia, 2009), with all the comments made by participants being highly appreciated.

References

[1] Anderson, G. and Ge, Y. (2005). The Size Distribution of Chinese Cities. Regional Science and Urban Economics, 35: 756-776.

[2] Azagra, J., Chorén, P., Goerlich, F. J. and Mas, M. (2006). La localización de la población española sobre el territorio. Un siglo de cambios: un estudio basado en series homogéneas (1900-2001). Fundación BBVA.

[3] Banerjee, A., M. Massimiliano and Osbat, C. (2005). Testing for PPP: should we use panel methods? Empirical Economics, Vol. 30, 77-91.

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

[5] 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: 330–347.

[6] Brakman, S., Garretsen, H. and Schramm, M. (2004). The Strategic Bombing of German Cities during World War II and its Impact on City Growth. Journal of Economic Geography, 4: 201-218.

[7] Cheshire, P. C. and Magrini, S. (2006). Population Growth in European Cities: Weather Matters- but only Nationally. Regional Studies, 40(1): 23-37.

[8] Clark, J. S. and Stabler, J. C. (1991). Gibrat's Law and the Growth of Canadian Cities. Urban Studies, 28(4): 635-639.

[9] Córdoba, J. C. (2008). A generalized Gibrat’s Law for Cities. International Economic Review, Vol. 49(4): 1463-1468.

[10] Davis, D. R. and Weinstein D. E. (2002). Bones, Bombs, and Break Points: The Geography of Economic Activity. American Economic Review, 92(5): 1269-1289.

[11] Duranton, G. (2006). Some Foundations for Zipf´s Law: Product Proliferation and Local Spillovers. Regional Science and Urban Economics, 36: 542-563.

[12] Duranton, G. (2007). Urban Evolutions: The Fast, the Slow, and the Still. American Economic Review, 97(1): 197-221.

[13] 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.

[14] Eeckhout, J. (2004). Gibrat's Law for (All) Cities. American Economic Review, 94(5): 1429-1451.

[15] Eeckhout, J. (2009). Gibrat's Law for (all) Cities: Reply. American Economic Review, 99(4): 1676–1683.

[16] Gabaix, X. (1999). Zipf's Law for Cities: An Explanation. Quaterly Journal of Economics, 114(3): 739-767.

[17] Gabaix, X. and Ioannides, Y. M. (2004). The Evolution of City Size Distributions. In Handbook of Urban and Regional Economics, vol. 4, ed. John V. Henderson and Jean. F. Thisse, 2341-2378. Amsterdam: Elsevier Science, North-Holland.

[18] Gibrat, R. (1931). Les Inégalités Économiques. París: Librairie du Recueil Sirey.

[19] Glaeser, E. L. and Shapiro, J. M. (2002). Cities and Warfare: The Impact of Terrorism on Urban Form. Journal of Urban Economics, 51: 205-224.

[20] González-Val, R. (2010). The Evolution of the US City Size Distribution from a Long-run Perspective (1900-2000) forthcoming in Journal of Regional Science.

[21] Guérin-Pace, F. (1995). Rank-Size Distribution and the Process of Urban Growth. Urban Studies, 32(3): 551-562.

[22] Härdle, W. (1990). Applied Nonparametric Regression. In Econometric Society Monographs. Cambridge, New York and Melbourne: Cambridge University Press.

[23] Henderson, V. and Wang, H. G. (2007). Urbanization and city growth: The role of institutions. Regional Science and Urban Economics, 37(3): 283–313.

[24] Ioannides, Y. M. and Overman, H. G. (2003). Zipf's Law for Cities: an Empirical Examination. Regional Science and Urban Economics, 33: 127-137.

[25] Kalecki, M. (1945). On the Gibrat Distribution. Econometrica, 13(2): 161-70.

[26] Levy, M. (2009). Gibrat's Law for (all) Cities: A Comment. American Economic Review, 99(4): 1672–1675.

[27] Parr, J. B. and Suzuki, K. (1973). Settlement Populations and the Lognormal Distribution. Urban Studies, 10: 335-352.

[28] Pesaran, M. H. (2007). A simple panel unit root test in the presence of cross-section dependence. Journal of Applied Econometrics, 22: 265–312.

[29] Petrakos, G., Mardakis, P. and Caraveli, H. (2000). Recent Developments in the Greek System of Urban Centres. Environment and Planning B: Planning and Design, 27(2): 169-181.

[30] Resende, M. (2004). Gibrat's Law and the Growth of Cities in Brazil: A Panel Data Investigation. Urban Studies, 41(8): 1537-1549.

[31] Santarelli, E., Klomp, L. and Thurik, R. (2006). Gibrat's Law: An Overview of the Empirical Literature. In Entrepreneurship, Growth, and Innovation: the Dynamics of Firms and Industries, ed. Enrico Santarelli, 41-73. New York: Springer.

[32] Soo, K. T. (2007). Zipf's Law and Urban Growth in Malaysia. Urban Studies, 44(1): 1-14.

[33] Sutton, J. (1997). Gibrat's Legacy. Journal of Economic Literature. 35(1): 40-59.

[34] White, H. (1980). A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity. Econometrica, 48: 817-38.

[35] Zipf, G. (1949). Human Behaviour and the Principle of Least Effort, Cambridge, MA: Addison-Wesley.

Tables

Table 1. - Empirical studies on city growth

Study	Country	Period	Truncation point	Sample size	GL	EcIss
Eaton and Eckstein (1997)	France and Japan	1876-1990 (F) 1925-1985 (J)	Cities > 50,000 inhabitants (F) Cities > 250,000 inhabitants (J)	39 (F), 40 (J)	A	non par (tr mat, lz)
Davis and Weinstein (2002)	Japan	1925-1965	Cities > 30,000 inhabitants	303	A	par (purt)
Brakman et al. (2004)	Germany	1946-1963	Cities > 50,000 inhabitants	103	A	par (purt)
Clark and Stabler (1991)	Canada	1975-1984	7 most populous cities	7	A	par (purt)
Resende (2004)	Brazil	1980-2000	Cities > 1,000 inhabitants	497	A	par (purt)
Eeckhout (2004)	US	1990-2000	All cities	19361	A	par (gr reg); non par (ker)
Ioannides and Overman (2003)	US	1900-1990	All MSAs	112 (1900) to 334 (1990)	A	non par (ker)
Gabaix and Ioannides (2004)	US	1900-1990	All MSAs	112 (1900) to 334 (1990)	A	non par (ker)
Black and Henderson (2003)	US	1900-1990	All MSAs	194 (1900) to 282 (1990)	R	par (purt)
Guérin-Pace (1995)	France	1836-1990	Cities > 2,000 inhabitants	675 (1836) to 1782 (1990)	R	par (corr)
Soo (2007)	Malaysia	1957-2000	Urban areas > 10,000 inhabitants	44 (1957) to 171 (2000)	R	par (purt)
Petrakos et al. (2000)	Greece	1981-1991	Urban centres > 5,000 inhabitants	150	R	par (gr reg)
Henderson and Wang (2007)	World	1960-2000	Metro areas > 100,000 inhabitants	1220 (1960) to 1644 (2000)	R	par (purt)
Bosker et al. (2008)	West Germany	1925-1999	Cities > 50,000 inhabitants	62	M	par (purt); non par (ker)
Anderson and Ge (2005)	China	1961-1999	Cities > 100,000 inhabitants	149	M	par (rank reg); non par (tr mat)

Gibrat's Law: GL	EcIss: Econometric Issues	gr reg: growth regressions	corr: coefficient of correlation (Pearson)
A: Accepted	par: parametric methods	ker: kernels	lz: Lorenz curves
R: Rejected	non par: non parametric methods	rank reg: rank regressions	
M: Mixed Results	purt: panel unit root tests	tr mat: transition matrices	

Table 2. - Estimated beta coefficients for the entire sample size

US										
Period	1900-1910	1910-1920	1920-1930	1930-1940	1940-1950	1950-1960	1960-1970	1970-1980	1980-1990	1990-2000
β estimated	0.008	0.022	0.042	0.009	0.048	0.051	0.027	-0.005	0.042	0.015
t-ratio	2.875	10.543	21.568	7.958	30.998	28.674	18.029	-3.934	35.152	12.352
SPAIN										
Period	1900-1910	1910-1920	1920-1930	1930-1940	1940-1950	1950-1960	1960-1970	1970-1981	1981-1991	1991-2001
β estimated	0.010	0.020	0.023	0.028	0.012	0.048	0.115	0.115	0.047	0.013
t-ratio	7.424	14.282	15.512	18.641	6.647	31.734	44.743	49.699	24.524	7.259
ITALY										
Period	1901-1911	1911-1921	1921-1931	1931-1936	1936-1951	1951-1961	1961-1971	1971-1981	1981-1991	1991-2001
β estimated	0.010	0.022	0.019	0.014	0.033	0.042	0.066	0.046	0.025	0.017
t-ratio	5.907	18.507	7.757	10.787	24.031	25.329	33.739	32.319	16.858	16.541

t-ratios calculated using White Heteroskedasticity-Consistent Standard Errors

Table 3. - Wilcoxon rank-sum test of lognormality

US											
Year	1900	1910	1920	1930	1940	1950	1960	1970	1980	1990	2000
p-value	0.0252	0.017	0.0078	0.0088	0.0208	0.0464	0.1281	0.1836	0.2538	0.323	0.4168
SPAIN											
Year	1900	1910	1920	1930	1940	1950	1960	1970	1981	1991	2001
p-value	0.5953	0.6144	0.6233	0.6525	0.4909	0.5792	0.6049	0.522	0.5176	0.622	0.7212
ITALY											
Year	1901	1911	1921	1931	1936	1951	1961	1971	1981	1991	2001
p-value	0.2081	0.2205	0.2352	0.291	0.2864	0.3118	0.2589	0.272	0.382	0.4671	0.5287

Ho: The distribution of cities follows a lognormal

Table 4. –Panel unit root tests, Pesaran's CADF statistic

Cities (N)	US	Spain	Italy
50	-0.488 (0.313)	-0.915 (0.180)	4.995 (0.999)
100	0.753 (0.774)	0.050 (0.520)	5.983 (0.999)
200	1.618 (0.947)	-2.866 (0.002)	-1.097 (0.136)
500	1.034 (0.849)	-12.132 (0.000)	5.832 (0.999)

test-statistic (p-value)
Pesaran's CADF test: standarized Ztbar statistic, $Z[\bar{t}]$
Variable: Relative size (in natural logarithms)
Sample size: (N, 11)

Figures

Figure 1.- Comparison of the Estimated Density Function (ln scale) and the Theoretical Lognormal in black (US)

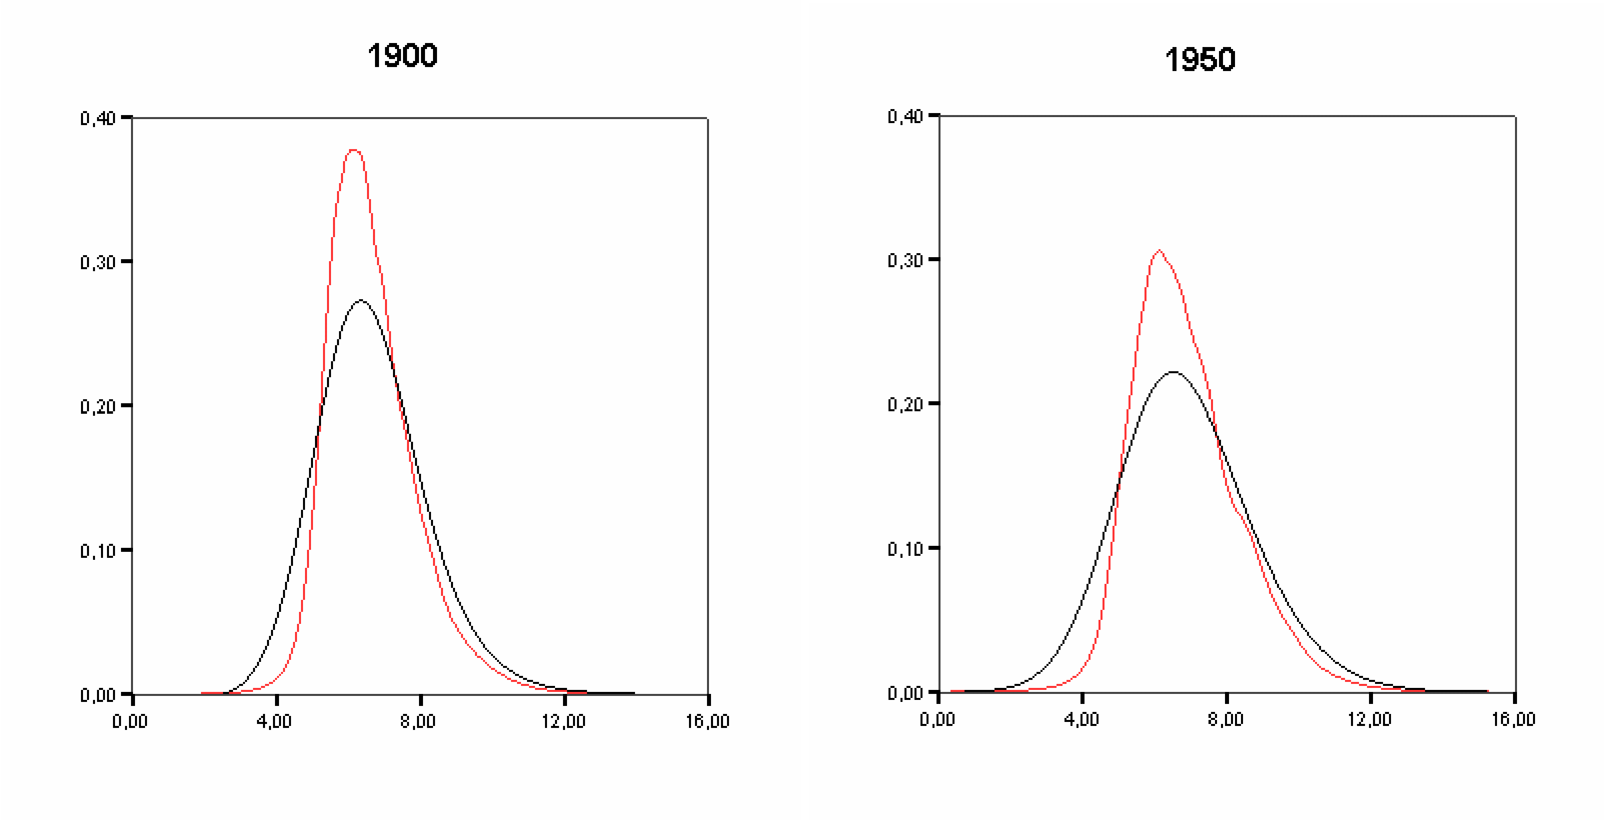

Figure 2.- Nonparametric Estimates (bandwidth 0.5)

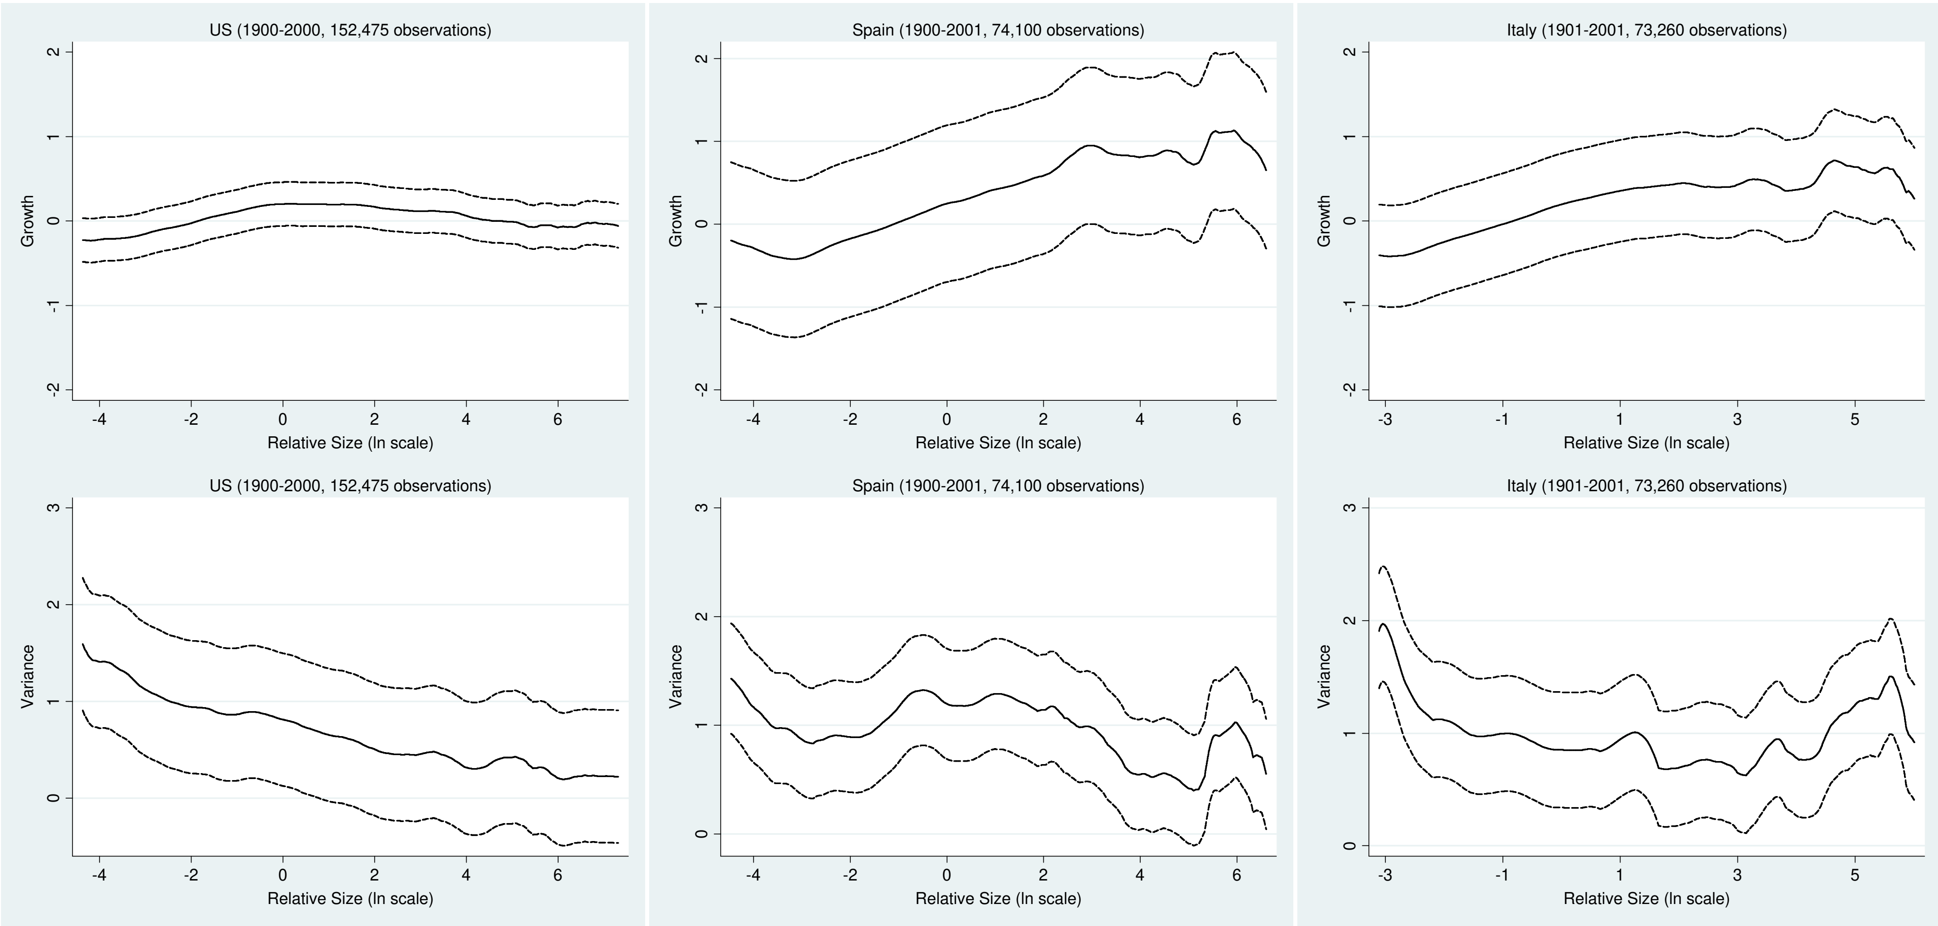

Figure 3.- New Entrants Nonparametric Estimates (bandwidth 0.5), (US, 1910-2000), 59,865 observations

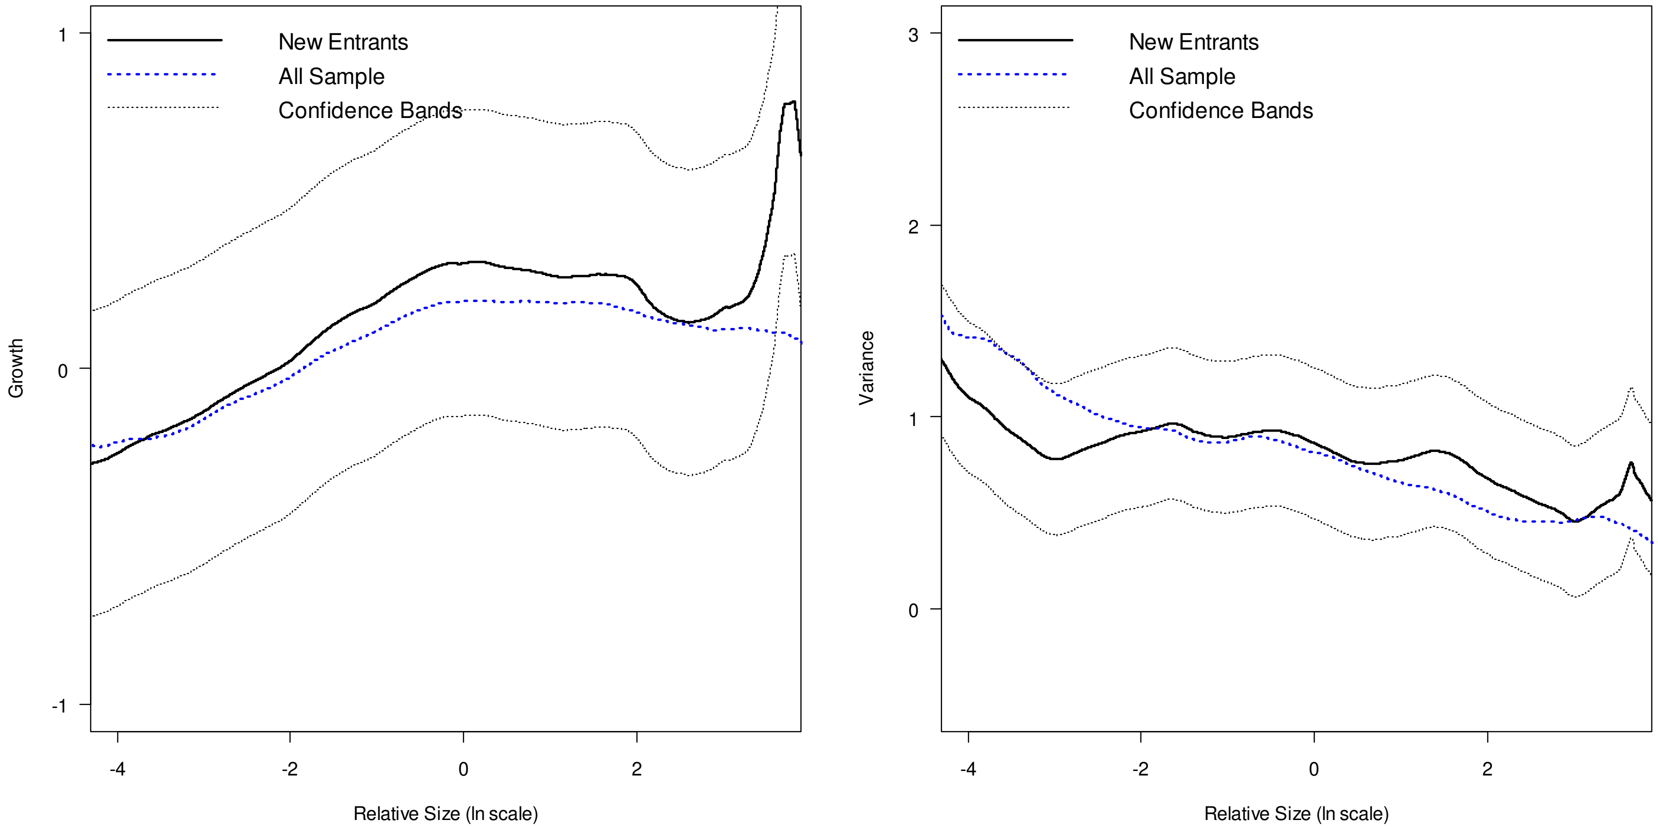

Figure 4.- Empirical Density Functions of the New Entrants