New Evidence on Gibrat's Law for Cities

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

22 November 2010

New Evidence on Gibrat's Law for Cities

Rafael González-Val[a]

Luis Lanaspa[b]

Fernando Sanz[b]

[a] Universidad de Barcelona & Instituto de Economía de Barcelona

[b] 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 on the complete distribution of cities (without size restrictions or a truncation point) in three countries (the US, Spain and Italy) for the entire 20th century. For this we use different techniques. First, panel data unit root tests tend to confirm the validity of Gibrat's Law in the upper tail distribution and, second, we find mixed evidence in favour of Gibrat's Law in the long term (in general, size affects the variance of the growth process but not its mean) when using nonparametric methods that relate the growth rate to city size. Moreover, the lognormal distribution works as a good description of city size distribution across the whole century when no truncation point is considered.

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

JEL: R00, C14.

Corresponding 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? By contrast, 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 well approximated with a lognormal, and that the explanation lay in the growth process of firms tending to be multiplicative and independent of their sizes. This proposition became known as Gibrat's Law and prompted a deluge of work exploring the validity of this Law on the distribution of firms (see the surveys of Sutton (1997) and Santarelli et al. (2006)). Gibrat's Law states 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 of the variable; in the words of Gibrat (1931): "*The Law of proportionate effect will therefore imply that the logarithms of the variable will be distributed following the (normal distribution).*"

In the field of urban economics, Gibrat's Law, especially since the 1990s, has given rise to numerous empirical studies testing its validity for city size distributions, arriving at a majority consensus, although not absolute at all, that it tends to hold 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 well-known empirical regularity 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 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 to 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, because they normally accept the fulfilment of Gibrat's Law but at the same time affirm that the distribution followed by city size (at least the upper tail) is a Pareto distribution, very different to the lognormal. Eeckhout (2004) was able to reconcile both results by demonstrating (as Parr and Suzuki (1973) affirmed 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) and Giesen et al. (2010) are the only studies to consider the entire city size distribution. But these are short term analyses[2], and the phenomenon under study (Gibrat's Law) is a long term result.

---

[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) took data from the United States censuses 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 considered no truncation point, but only for the 2000 US census data. Giesen et al. (2010), based on the pioneering work of Reed (2002), fit the double Pareto lognormal (DPLN) and lognormal distributions to data from eight countries (Germany, France, the Czech Republic, Hungary, Italy, Switzerland, Brazil and the US). The data for the US were the same as used by Eeckhout (2004, 2009) and Levy (2009). It should be noted that the DPLN also builds on Gibrat's Law, particularly a generalized version of the Law (Reed, 2002).

The aim of this work is to test empirically the validity of Gibrat's Law in the growth of cities, using data on the complete distribution of cities (without size restrictions or a truncation point) in three countries (the US, Spain and Italy) for the entire 20th century. The following section offers a brief overview of the literature on Gibrat's Law for cities, and the results obtained. Section 3 presents the databases, with special attention to the US census. From the results we deduce that panel data unit root tests tend to confirm the validity of Gibrat's Law in the upper tail distribution (Section 4.1), and we find mixed evidence in favour of Gibrat's Law in the long term (in general, size affects the variance of the growth process but not its mean) when using nonparametric methods that relate growth rate with city size (section 4.2). The validity of the Law in the short term is weaker. In Section 5, we test if the lognormal distribution is a good description of city size distributions across the entire century. The American case is different as the number of cities increases significantly. Finally, in Section 6 we study the behaviour of the new entrants. 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 studies on urban economics that we know of. The countries considered, statistical and econometric techniques used and sample sizes are heterogeneous, and the results are quite mixed, with the acceptance of the Law the predominating outcome, although by a scarce margin.

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. 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 seven most populous Canadian cities), although the main contribution of their work is to propose the use of a data panel methodology and unit root tests in the analysis of urban growth. This is also the methodology that 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 inhabitants 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 Metropolitan Statistical Areas (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).

Other works exist that reject 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

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

during the period 1836–1990 there seems to be a fairly strong correlation between city size and growth rate, a correlation that 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 seems 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) and Giesen et al. (2010), none of these studies considers the entire distributions of cities[4], because all of them impose a truncation point, whether explicitly by taking cities above a minimum population threshold or implicitly by working with MSAs[5]. This is usually because of practical reasons regarding data availability. For this reason, most studies focus on analysing the most populous cities, the upper tail distribution. There are two reasonable justifications for this approach. First, the largest cities represent most of the population of a country. Second, the growth rate of the biggest cities has less variance than that of the smallest ones (scale effect).

However, any test 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

[4] Michaels et al. (2010) use data from Minor Civil Divisions to track the evolution of populations across both rural and urban areas in the United States from 1880 to 2000.

[5] In the US, to qualify as an 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.

deduction can lead to biased conclusions, because it must not be forgotten that what is being analysed is the behaviour of a few cities, which in addition to 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 that cannot be representative of all urban centres.

3. Databases

We use city population data from three countries: the US, Spain and Italy[6]. We have taken the data corresponding to the census of each decade of the 20th century[7]. Table 2 presents the number of cities for each decade and the descriptive statistics.

The data for the US are the same as those used by González-Val (2010). 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 20th 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.

Three details should be noted. First, that all the cities in Alaska, Hawaii and Puerto Rico for each decade are excluded because 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

[6] We use data from "legal" cities. However, there are problems of international comparability because the administrative definition of city changes from one country to another. However, the concepts of municipality used in Spain and Italy are similar.
[7] No census exists in Italy for 1941 because of its participation in the Second World War, so we have taken the data for 1936.

be heterogeneous in geographical terms and thereby could not be compared. Second, for the same reason we also exclude all the unincorporated places (concentrations of population that form no part of any incorporated place but which are locally identified with a name), which began to be accounted after 1950 (they were renamed Census Designated Places (CDPs) in 1980). However, these settlements did exist earlier, so that their inclusion would again present a problem of inconsistency in the sample. In addition, 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 had grown to 5,366 places (11.27%).

Third, the percentage of the total US population that our sample of incorporated places represents can appear low compared with other studies using MSAs. However, it is similar to that of other works using cities.[8] The population excluded from the sample is what the US Census Bureau calls population not in place. Incorporated places and CDPs do not cover the whole territory of the US. Some territory is excluded from any recognised place. For example, more than 74 million people (26.64% of the total US population) lived in a territory that, at least officially, was not in a place in 2000.[9] Most of this amount (61.58% in 2000) is rural population.

Although these people living outside incorporated places are excluded from our sample, they are included in some MSAs because these are multi-county units and this population is counted as inhabitants of the counties. MSAs cover huge geographic areas and include a large proportion of the population living in rural areas. This explains why the percentage of total population represented by MSAs is higher than our sample of incorporated places. However, despite the sample of incorporated places covering a

---

[8] For example, see Kim (2000) and Kim and Margo (2004), where city is defined as an area having a population of greater than 2,500 inhabitants.

[9] Census 2000 data on the population in places and not in places can be found in Table 9 of PHC-3 (US Summary, part 1), available online at: http://www.census.gov/prod/cen2000/index.html.

lower percentage of the total US population, the population of incorporated places is almost entirely urban (94.18% in 2000) compared with 88.35% of urban population in the MSAs.

For Spain and Italy, the geographical unit of reference is the municipality, and the data come 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[10], INE, www.ine.es. The de facto resident population has been taken for each city. It can be observed that our samples of all cities represent almost the total populations of both countries.

Figure 1 displays the mean growth rates for each decade, calculated from gross growth rates, defined as $g_{it} = \frac{S_{it} - S_{it-1}}{S_{it-1}}$, where S_{it} is the population of the city i in the year t. In the US, it can be observed that the first decades of the century saw strong growth rates for city sizes. However, this period of growth ended in 1920–1930. Between 1940 and 1980, the high growth rates seem to recover and then fall in the last two decades. The two periods of lowest growth, 1930–1940 and 1980–1990, are very close to two profound economic crises (the Great Depression and the second oil supply shock in 1979). Spain and Italy present lower growth rates, even with some periods of negative growth of cities (on average). Since overall population did not fall in these countries in any decade, the declines in growth rates are related to composition effects: small municipalities tend to shrink while the few large ones grow, in such a way that the

[10] The official INE census have been improved in an alternative database, created by Azagra et al. (2006), reconstructing the population census for the 20th 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.

unweighted average decreases. In Italy, this period of negative growth rates (1951–1971) coincides with the post-war period after the Second World War, while in Spain the growth rates of cities are strongly negative during the military dictatorship period.

The US is an extremely interesting country in which to analyse the evolution of urban structure because it is a relatively young country whose inhabitants are characterised by high mobility. By contrast, the European countries have a much older urban structure and their inhabitants present a greater resistance to movement; specifically, Cheshire and Magrini (2006) estimate mobility in the US is 15 times higher than that in Europe.

Considering these two types of countries provides information about different urban behaviours. Although Spain and Italy have a consolidated urban structure and new cities are rarely created (urban growth is produced by population increases 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.

4. Testing for Gibrat's Law

4.1. Parametric analysis: panel unit root testing

Clark and Stabler (1991) suggest 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, 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. Second, the presence of cross-sectional dependence across the cities in the panel can give rise to estimations that are not very robust. It has been well established in the literature that panel unit root and stationarity tests 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 based on the CADF statistic (cross-sectional ADF (see below) statistic).

To eliminate 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:

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

where a_i is the individual city-specific average growth rate. We will test for the presence of a unit root in the natural logarithm of city relative size $(y_{it} = \ln s_{it})$. City relative size (s_{it}) is defined as $s_{it} = \frac{S_{it}}{\bar{S}_t} = \frac{S_{it}}{\frac{1}{N}\sum_{i=1}^{N} S_{it}}$; from a long term temporal perspective of steady state distributions it is necessary to use a relative measure of size (Gabaix and Ioannides, 2004). Null hypothesis assumes that all series are nonstationary, 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 3 shows the results of Pesaran's 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 for all the decades. 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, since when the sample size is more than the 200 largest cities the unit root is rejected, indicating a positive relationship between relative size and growth rate even for the largest cities. This result could be a consequence of the political regime, a military dictatorship in most decades of the century. In this context, but certainly only for the capital city, Ades and Glaeser (1995) find that this city will tend to be more dominant the more political instability there is in a country and the more authoritarian is its regime.

4.2. Nonparametric analysis: kernel regression conditional on city size

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

$$g_i = m(s_i) + \varepsilon_i \, , \qquad (2)$$

where g_i is the growth rate $(\ln s_{it} - \ln s_{it-1})$ normalised (subtracting the contemporary mean and dividing by the standard deviation in the relevant decade) and s_i is the logarithm of the *i*th city's 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 estimate $\hat{m}(s)$ the Nadaraya-Watson method is used, exactly as it appears in Härdle (1990), based on the following expression[11]:

$$\hat{m}(s) = \frac{n^{-1}\sum_{i=1}^{n} K_h(s - s_i) g_i}{n^{-1}\sum_{i=1}^{n} K_h(s - s_i)} \, , \qquad (3)$$

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 to permit comparisons between countries.

---

[11] The calculation was performed 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.

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:

$$\hat{\sigma}^2(s) = \frac{n^{-1}\sum_{i=1}^{n} K_h(s-s_i)(g_i - \hat{m}(s))^2}{n^{-1}\sum_{i=1}^{n} K_h(s-s_i)}. \qquad (4)$$

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

Following Gabaix and 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."* As growth rates are normalised, if Gibrat's Law in mean was strictly fulfilled, the nonparametric estimate would be a straight line on the zero value. Values different to zero involve deviations from the mean. Moreover, the estimated variance of the growth rate would also be a straight line on the value one, which would mean that the variance does not depend on the size of the city. To be able to test these hypotheses, we construct bootstrapped 95% confidence bands (calculated from 500 random samples with replacements).

We offer a first approach to the behaviour of city growth from a short term perspective, i.e., considering each decade individually. Figures 2, 3 and 4 show the nonparametric estimates for the US, Spain and Italy, respectively, corresponding to three representative decades (the behaviour in the rest of the adjacent periods is similar).

Two different behaviours can be observed; while in US the estimate of growth is very close to the zero value (this value falls within the confidence bands for most of the distributions, supporting Gibrat's Law even in the short term), in Spain and Italy a different pattern of growth can be seen. Starting from the beginning of the century until the mid-century, city growth exhibits clear divergent behaviour in both European countries, although Gibrat's Law can only be rejected for some values at the upper tail distribution. However, in the second half of the century growth changes gradually to an inverted U-shaped pattern.

There is a negative relationship between the estimated variance of growth and city size in the three countries for most of the decades (this is especially true for the US, where Gibrat's Law can be rejected at the upper tail), although in Spain and Italy the behaviour of the variance is irregular, particularly in the first decades of the century.

Moreover, to analyse the entire 20th century we build a pool with all the growth rates between two consecutive periods. This enables us to carry out long term analysis. Figure 5 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 the growth rates being significantly different for any city size cannot be rejected. For Spain and Italy, the estimated mean grows with the sample size, although it is significantly different to zero only for the largest cities[12]. 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

[12] In the case of Spain, this divergent behaviour could be the explanation for the rejection of the unit root null hypothesis obtained in the previous section.

the population[13]. However, the estimations by decade indicate that this tendency would have reversed in the second half of the century. Therefore, we find evidence in favour of Gibrat's Law for the US throughout the 20th century. Also for Spain and Italy, although for the Mediterranean countries the largest cities present some divergent behaviour.

Figure 5 also shows the nonparametric estimates of the variance of growth rates 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 variances.

Our results, obtained with a 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 (Figure 5) show that while the mean of growth (Gibrat's Law for means) seems to be independent of size in the three countries in the long term (although in Spain and Italy the largest cities present some divergent behaviour), the variance of growth (Gibrat's Law for variances) does depend negatively on size: the smallest cities present clearly higher variances in all three countries (although in Spain and Italy the behaviour is more erratic and the biggest cities also have high variances). In the short term (Figures 2, 3 and 4) the evidence regarding Gibrat's Law is weaker, as corresponds to a Law which is thought to hold mainly in the long term (Gabaix and Ioannides, 2004).

[13] This result could be related with the "safe harbour effect" of Glaeser and Shapiro (2002), which is a centripetal force that tends to agglomerate the populations in large cities when there is an armed conflict.

This points to Gibrat's Law holding weakly (growth is proportional in means but not in variance). Gabaix (1999) contemplates the 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 generalisation of Gibrat's Law that allows size to affect the variance of the growth process but not its mean.

5. What about city size distribution?

Proportionate growth implies a lognormal distribution, and this is a statistical relationship (Gibrat, 1931; Kalecki, 1945). However, as Eeckhout (2004) shows, city size distribution follows a lognormal only when we consider all cities without any size restriction. Our results show that the growth process leads to a lognormal distribution with standard deviation that increases 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 nonparametric 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 large sample sizes, although the deviations from lognormality are arbitrarily small.

Table 4 shows the results of the test. The conclusion is that the null hypothesis of lognormality is not rejected 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 not rejected 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 6 shows the empirical density functions estimated by adaptive Gaussian kernels for 1900, 1950 (the last year in which lognormality is rejected at 5%) and 2000. The motive for this systematic rejection seems to be an excessive concentration of density in the central values, higher than would correspond to the theoretical lognormal distribution (dotted line). Starting in 1900 with a very leptokurtic distribution and with a great deal of density concentrated in the mean value, from 1930 (not shown), when the growth of the urban population slows, the distribution loses kurtosis and concentration decreases, not rejecting 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 for the entire 20th century and in the US for most decades, depending on the significance level).

6. Entrant cities

We must distinguish between the American and European cases, because 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. Moreover, 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 addition, if there is a drift (when there is average city growth), the distribution from which new entrants are drawn is unlikely to be stationary if the result obtained is proportionate growth.

Therefore, Figure 7 shows the nonparametric estimates of the growth rate and 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 (the first period of our sample in which new cities appear). Bootstrapped 95% confidence bands are also presented. The estimations show how the cities entering the sample from 1910 usually had growth rates that 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. These differences in variance indicate that part of the increased variance at the bottom of the size distribution can be explained by the cities that entered the distribution throughout the 20th century.

In addition, Figure 8, 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 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. Briefly, 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, 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 centres and its evolution is fundamentally stochastic without any fixed pattern of behaviour. Moreover, although this issue 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 20^{th} 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 employs different methods (parametric and nonparametric).

There are three basic conclusions, the first two being more important. First, 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 200 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 nonparametric 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 US and, to a lesser extent, for Spain and Italy. In these two European countries, there is a positive relationship between city size and growth, although this divergent behaviour is only significant for the largest cities. For variances the predominant behaviour, in turn, agrees with the Law, except for the largest and smallest cities, depending on the country. In the short term, as could be anticipated, the evidence regarding the validity of the Law is more mixed.

Second, the lognormal distribution works as a good description of city size distribution across the entire century when no truncation point is considered. Wilcoxon's rank-sum test shows that, except for the US in the first half of the century, the lognormal distribution is systematically never rejected.

Finally, the case of the US differs in that the number of cities doubles over the 20th 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. These 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 that entered the distribution throughout the 20th century.

Acknowledgements

The authors would like to thank the Spanish Ministerio de Educación y Ciencia (ECO2009-09332 project), 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), at the XXXV Reunión de Estudios Regionales (Valencia, 2009), at the 9th Annual Conference of the European Economics and Finance Society (Athens, 2010) and at the 50th Congress of the European Regional Science Association (Jönköping, 2010), with all the comments made by participants being highly appreciated.

References

[1] Ades, A. F. and Glaeser, E. L. (1995). Trade and circuses: explaining urban giants. The Quarterly Journal of Economics, 110: 195–227.

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

[3] 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.

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

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

[6] 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.

[7] 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.

[8] Cheshire, P. C. and Magrini, S. (2006). Population growth in European cities: weather matters - but only nationally. Regional Studies, 40(1): 23–37.

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

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

[11] 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.

[12] Duranton, G. (2006). Some Foundations for Zipf´s Law: product proliferation and local spillovers. Regional Science and Urban Economics, 36: 542–563.

[13] Duranton, G. (2007). Urban evolutions: the fast, the slow, and the still. American Economic Review, 97(1): 197–221.

[14] 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.

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

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

[17] Gabaix, X. (1999). Zipf’s Law for Cities: An Explanation. Quarterly Journal of Economics, 114(3): 739–767.

[18] 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.

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

[20] Giesen, K., Zimmermann A. and Suedekum J. (2010). The size distribution across all cities – double Pareto lognormal strikes. Journal of Urban Economics, 68: 129–137.

[21] 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.

[22] 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. DOI 10.1111/j.1467-9787.2010.00685.x

[23] Guérin-Pace, F. (1995). Rank-size distribution and the process of urban growth. Urban Studies, 32(3): 551–562.

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

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

[26] Ioannides, Y. M. and Overman, H. G. (2003). Zipf's Law for cities: an empirical examination. Regional Science and Urban Economics, 33: 127–137.

[27] Kalecki, M. (1945). On the Gibrat Distribution. Econometrica, 13(2): 161–170.

[28] Kim, S. (2000). Urban Development in the United States. Southern Economic Journal, 66: 855–880.

[29] Kim, S. and Margo R. A. (2004). Historical Perspectives on U.S. Economic Geography. Handbook of Urban and Regional Economics, vol. 4, J. V. Henderson

and J. F. Thisse, eds. Amsterdam: Elsevier Science, North-Holland, Chapter 66, pp. 2982–3019.

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

[31] Michaels, G., Rauch F. and Redding S. J. (2010). Urbanization and Structural Transformation. Unpublished manuscript, London School of Economics.

[32] Parr, J. B. and Suzuki, K. (1973). Settlement populations and the lognormal distribution. Urban Studies, 10: 335–352.

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

[34] 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.

[35] Reed, W. (2002). On the rank-size distribution for human settlements. Journal of Regional Science, 42: 1–17.

[36] Resende, M. (2004). Gibrat’s Law and the growth of cities in Brazil: a panel data investigation. Urban Studies, 41(8): 1537–1549.

[37] 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.

[38] Soo, K. T. (2007). Zipf's Law and urban growth in Malaysia. Urban Studies, 44(1): 1–14.

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

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

Tables

Table 1. Empirical Studies on City Growth and Gibrat's Law. A Survey

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	Seven 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. Number of Cities and Descriptive Statistics by Year and Country

US							
Year	Cities	Mean	Standard deviation	Minimum	Maximum	Country Population (CP)	Percentage of CP in our sample
1900	10,596	3,376.04	42,323.90	7	3,437,202	76,212,168	46.9
1910	14,135	3,560.92	49,351.24	4	4,766,883	92,228,496	54.6
1920	15,481	4,014.81	56,781.65	3	5,620,048	106,021,537	58.6
1930	16,475	4,642.02	67,853.65	1	6,930,446	123,202,624	62.1
1940	16,729	4,975.67	71,299.37	1	7,454,995	132,164,569	63.0
1950	17,113	5,613.42	76,064.40	1	7,891,957	151,325,798	63.5
1960	18,051	6,408.75	74,737.62	1	7,781,984	179,323,175	64.5
1970	18,488	7,094.29	75,319.59	3	7,894,862	203,302,031	64.5
1980	18,923	7,395.64	69,167.91	2	7,071,639	226,542,199	61.8
1990	19,120	7,977.63	71,873.91	2	7,322,564	248,709,873	61.3
2000	19,296	8,968.44	78,014.75	1	8,008,278	281,421,906	61.5
SPAIN							
Year	Cities	Mean	Standard deviation	Minimum	Maximum	Country Population (CP)	Percentage of CP in our sample
1900	7,800	2,282.40	10,177.75	78	539,835	18,616,630	95.6
1910	7,806	2,452.01	11,217.02	92	599,807	19,990,669	95.7
1920	7,812	2,621.92	13,501.02	82	750,896	21,388,551	95.8
1930	7,875	2,892.18	17,513.90	79	1,005,565	23,677,095	96.2
1940	7,896	3,180.65	20,099.96	11	1,088,647	26,014,278	96.5
1950	7,901	3,479.86	26,033.29	64	1,618,435	28,117,873	97.8
1960	7,910	3,801.71	33,652.11	51	2,259,931	30,582,936	98.3
1970	7,956	4,240.98	43,971.93	10	3,146,071	33,956,047	99.4
1981	8,034	4,701.40	45,995.35	5	3,188,297	37,742,561	100.0
1991	8,077	4,882.27	45,219.85	2	3,084,673	39,433,942	100.0
2001	8,077	5,039.37	43,079.46	7	2,938,723	40,847,371	99.6
ITALY							
Year	Cities	Mean	Standard deviation	Minimum	Maximum	Country Population (CP)	Percentage of CP in our sample
1901	7,711	4,274.84	14,424.61	56	621,213	32,963,000	100.0
1911	7,711	4,648.11	17,392.98	58	751,211	35,842,000	100.0
1921	8,100	4,863.80	20,031.61	58	859,629	39,397,000	100.0
1931	8,100	5,067.10	22,559.85	93	960,660	41,043,000	100.0
1936	8,100	5,234.38	25,274.48	116	1,150,338	42,398,000	100.0
1951	8,100	5,866.12	31,137.52	74	1,651,393	47,516,000	100.0
1961	8,100	6,249.82	39,130.55	90	2,187,682	50,624,000	100.0
1971	8,100	6,683.52	45,581.66	51	2,781,385	54,137,000	100.0
1981	8,100	6,982.33	45,329.33	32	2,839,638	56,557,000	100.0
1991	8,100	7,009.63	42,450.26	31	2,775,250	56,778,000	100.0
2001	8,100	7,021.20	39,325.47	33	2,546,804	56,996,000	99.8

Table 3. 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: standardised Ztbar statistic, $Z[\bar{t}]$
Variable: Relative size (in natural logarithms)
Sample size: (N, 11)

Table 4. Wilcoxon Rank-sum Test of Lognormality by Year and Country

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

Figures

Figure 1. Decennial Average Growth Rates by Country

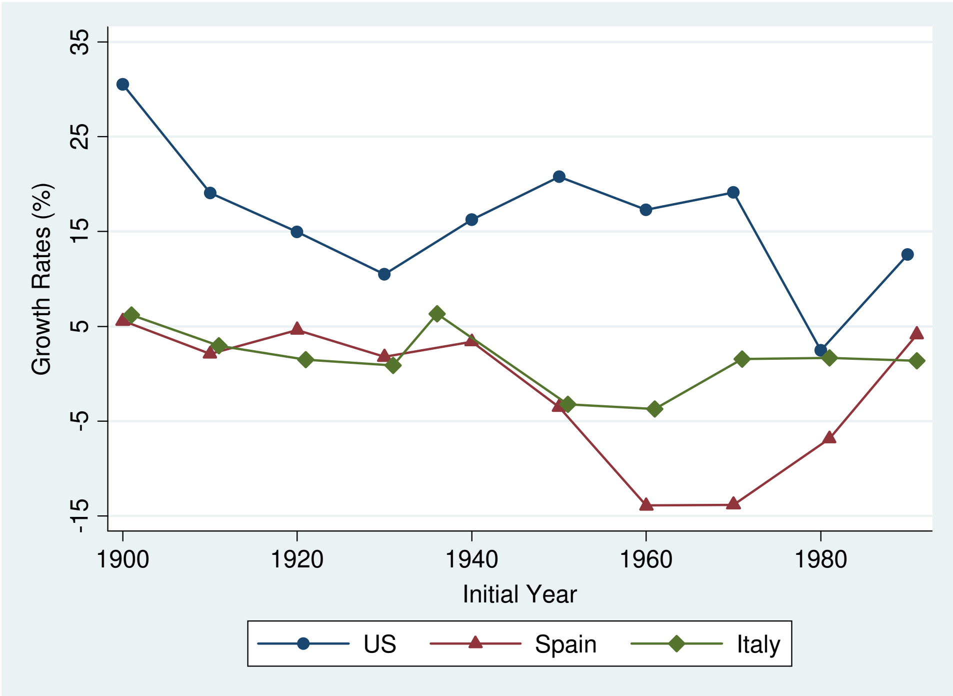

Figure 2. Nonparametric Estimates (bandwidth 0.5) of the Growth Rate and its Variance for the US by Decade

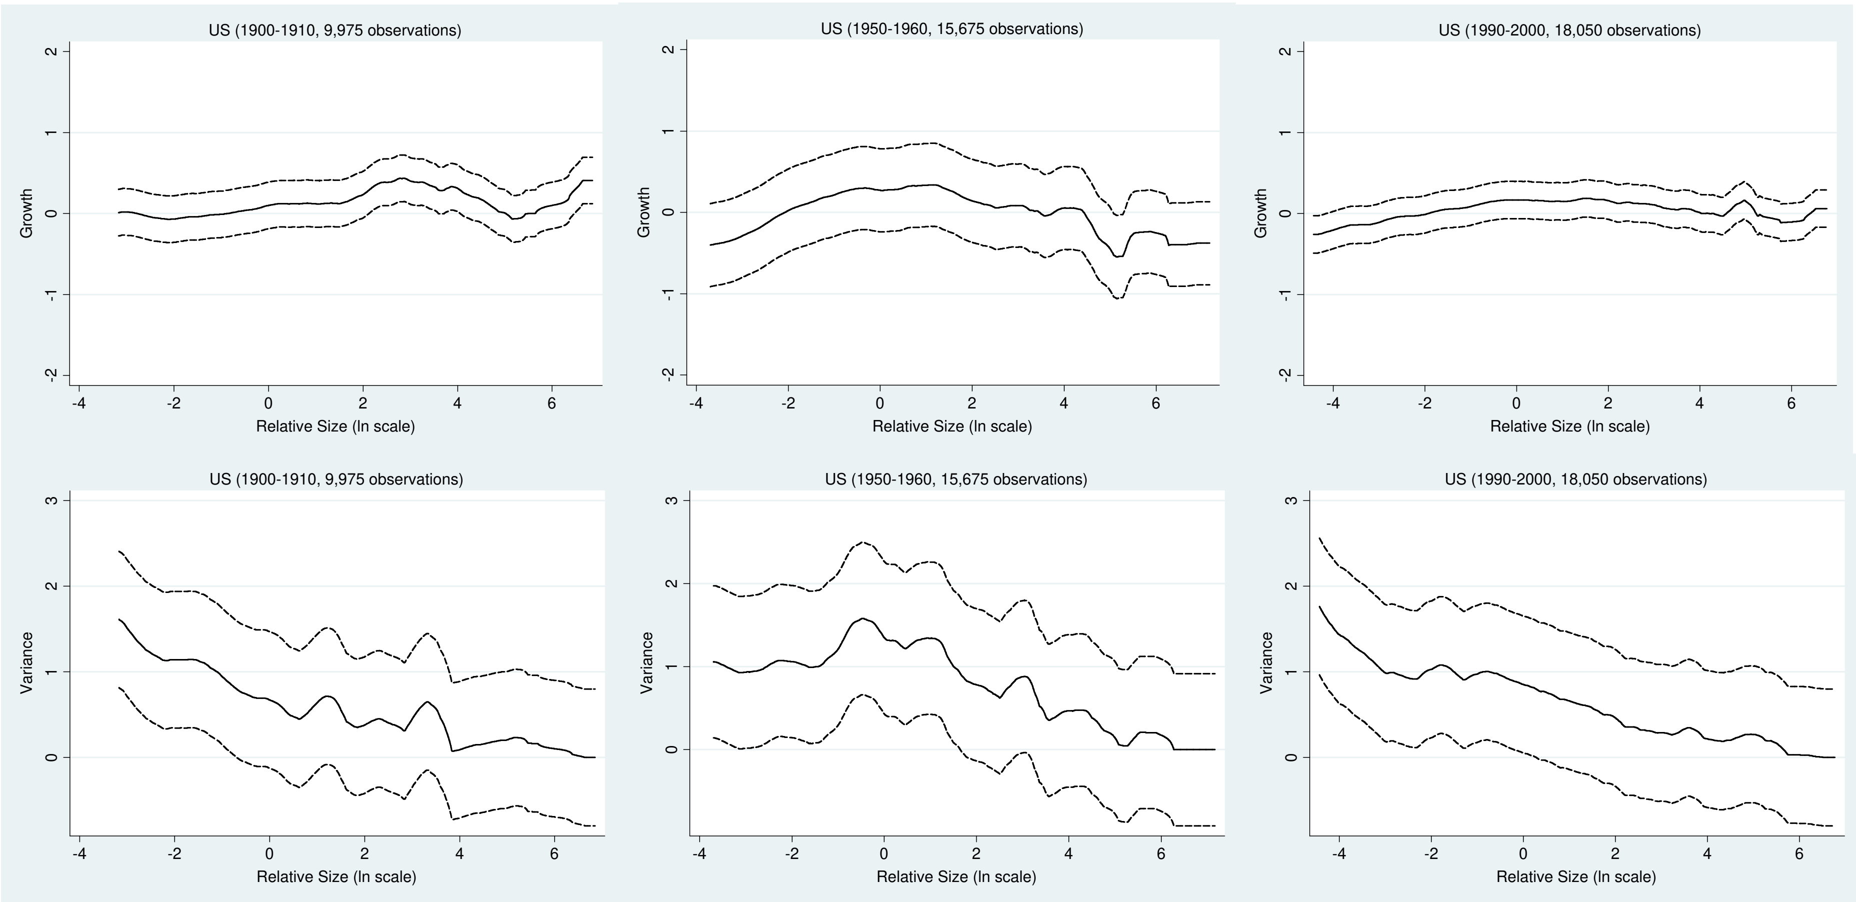

Figure 3. Nonparametric Estimates (bandwidth 0.5) of the Growth Rate and its Variance for Spain by Decade

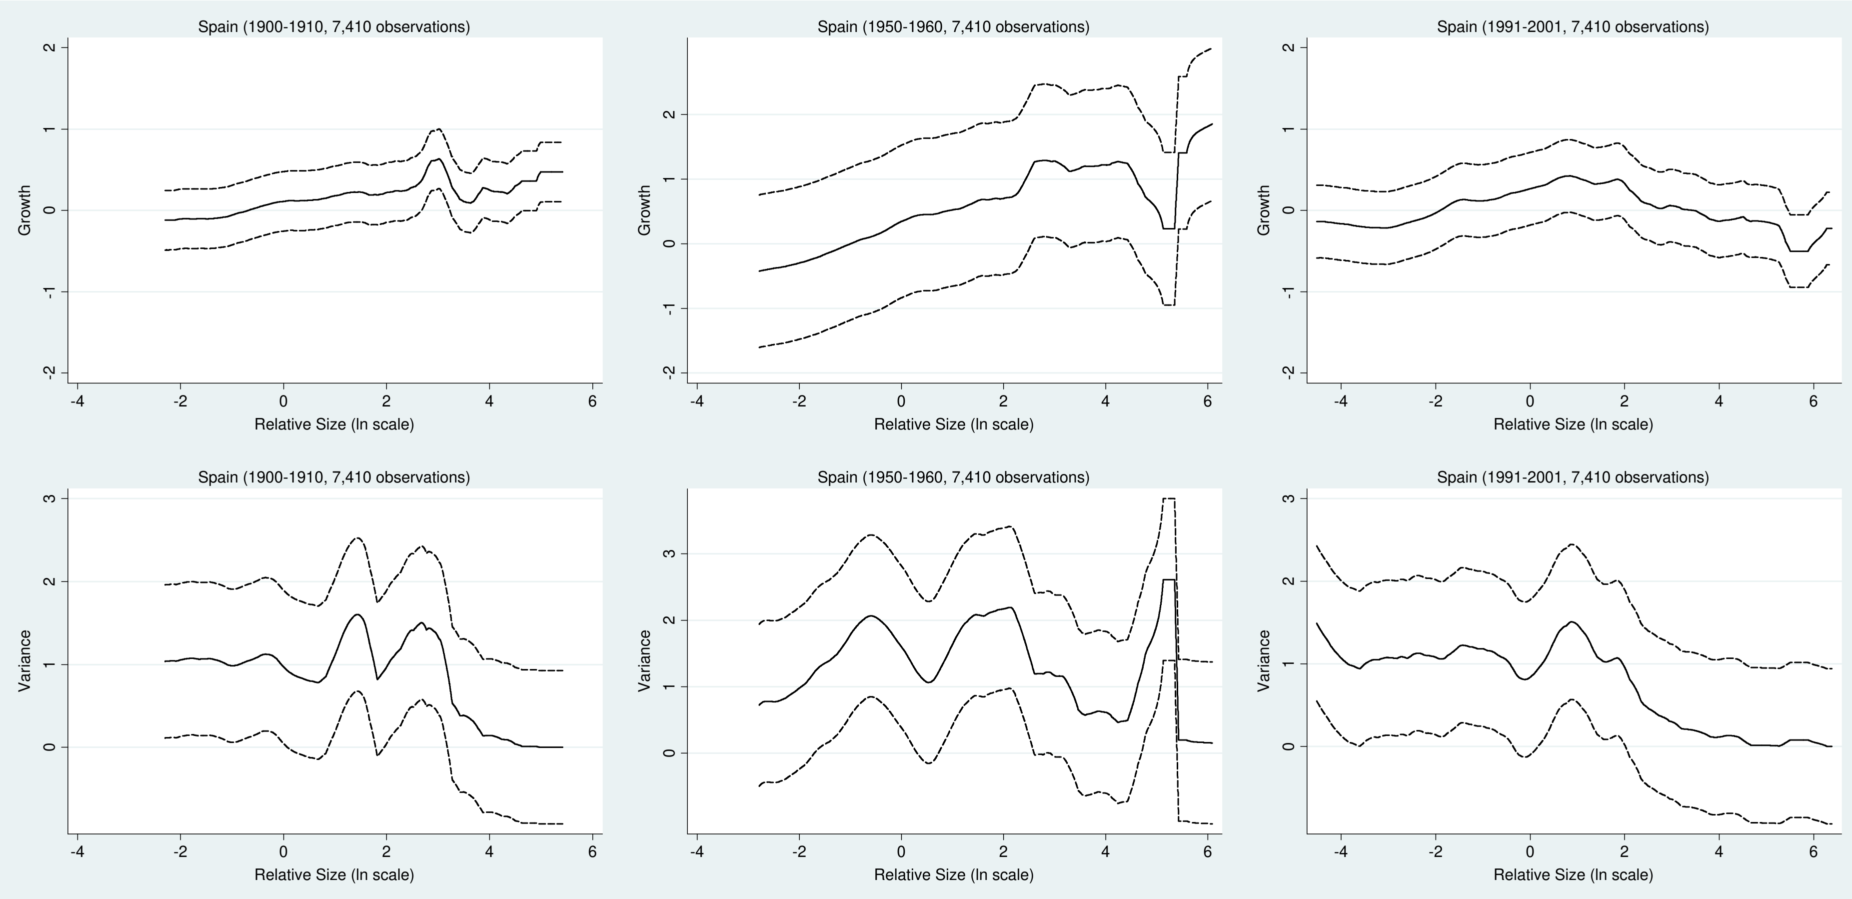

Figure 4. Nonparametric Estimates (bandwidth 0.5) of the Growth Rate and its Variance for Italy by Decade

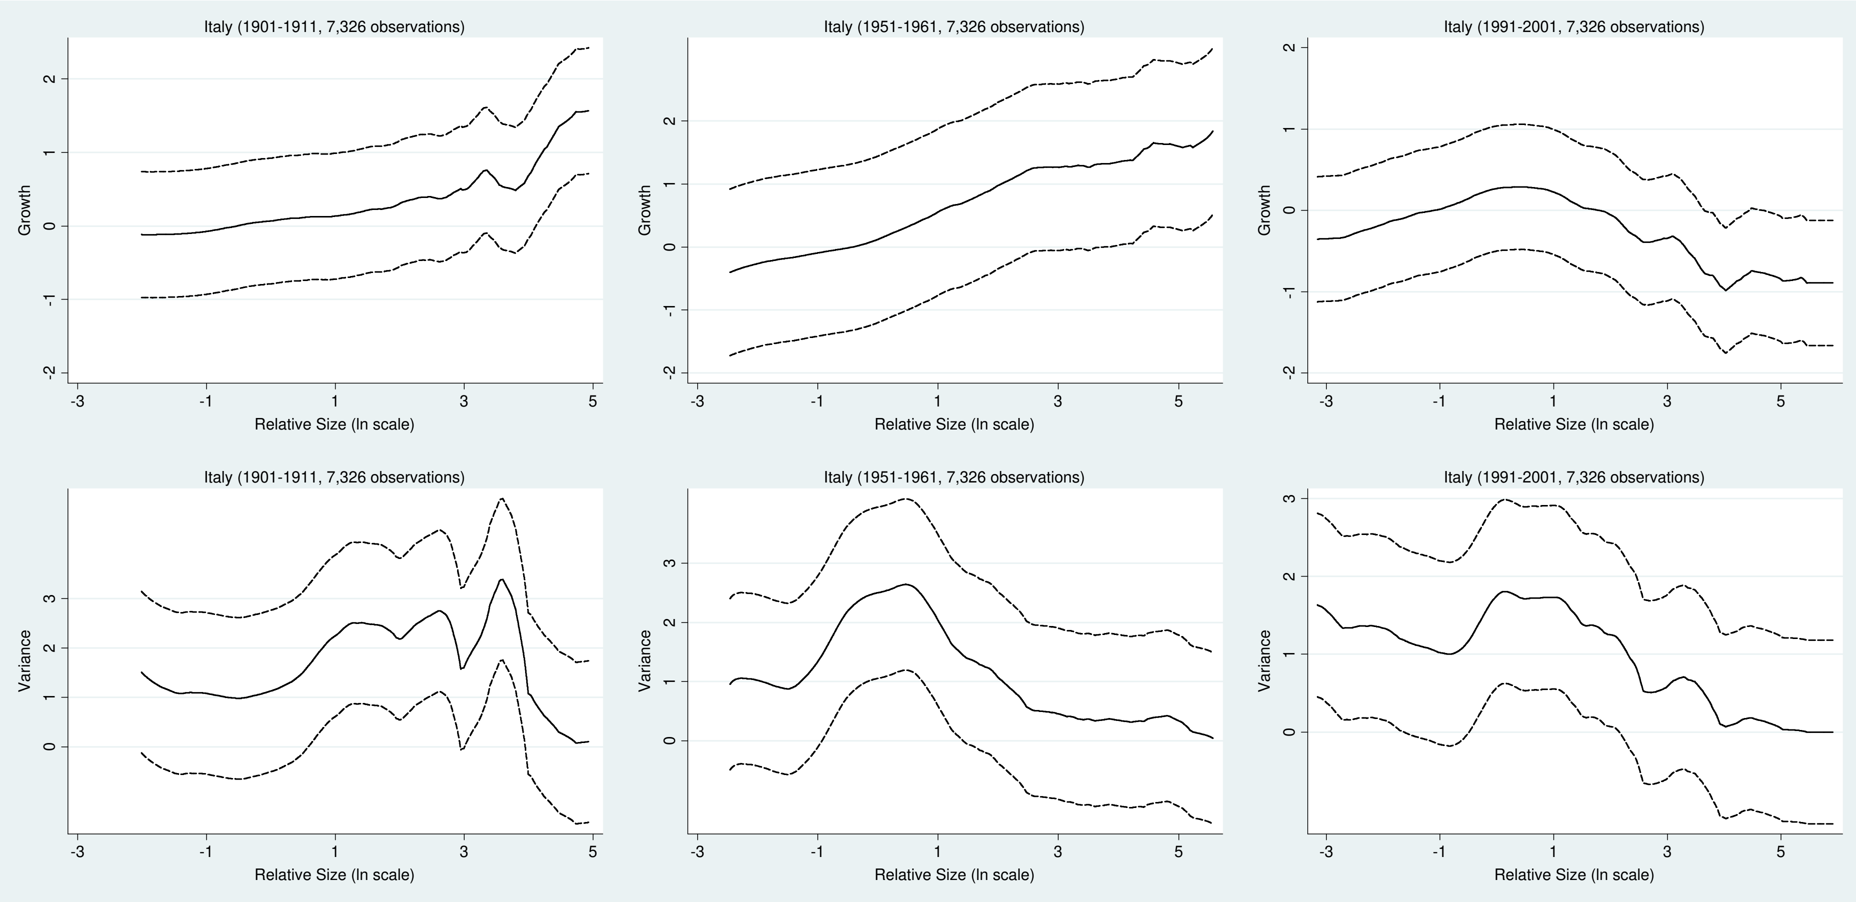

Figure 5. Nonparametric Estimates (bandwidth 0.5) of the Growth Rate and its Variance. All the 20th Century

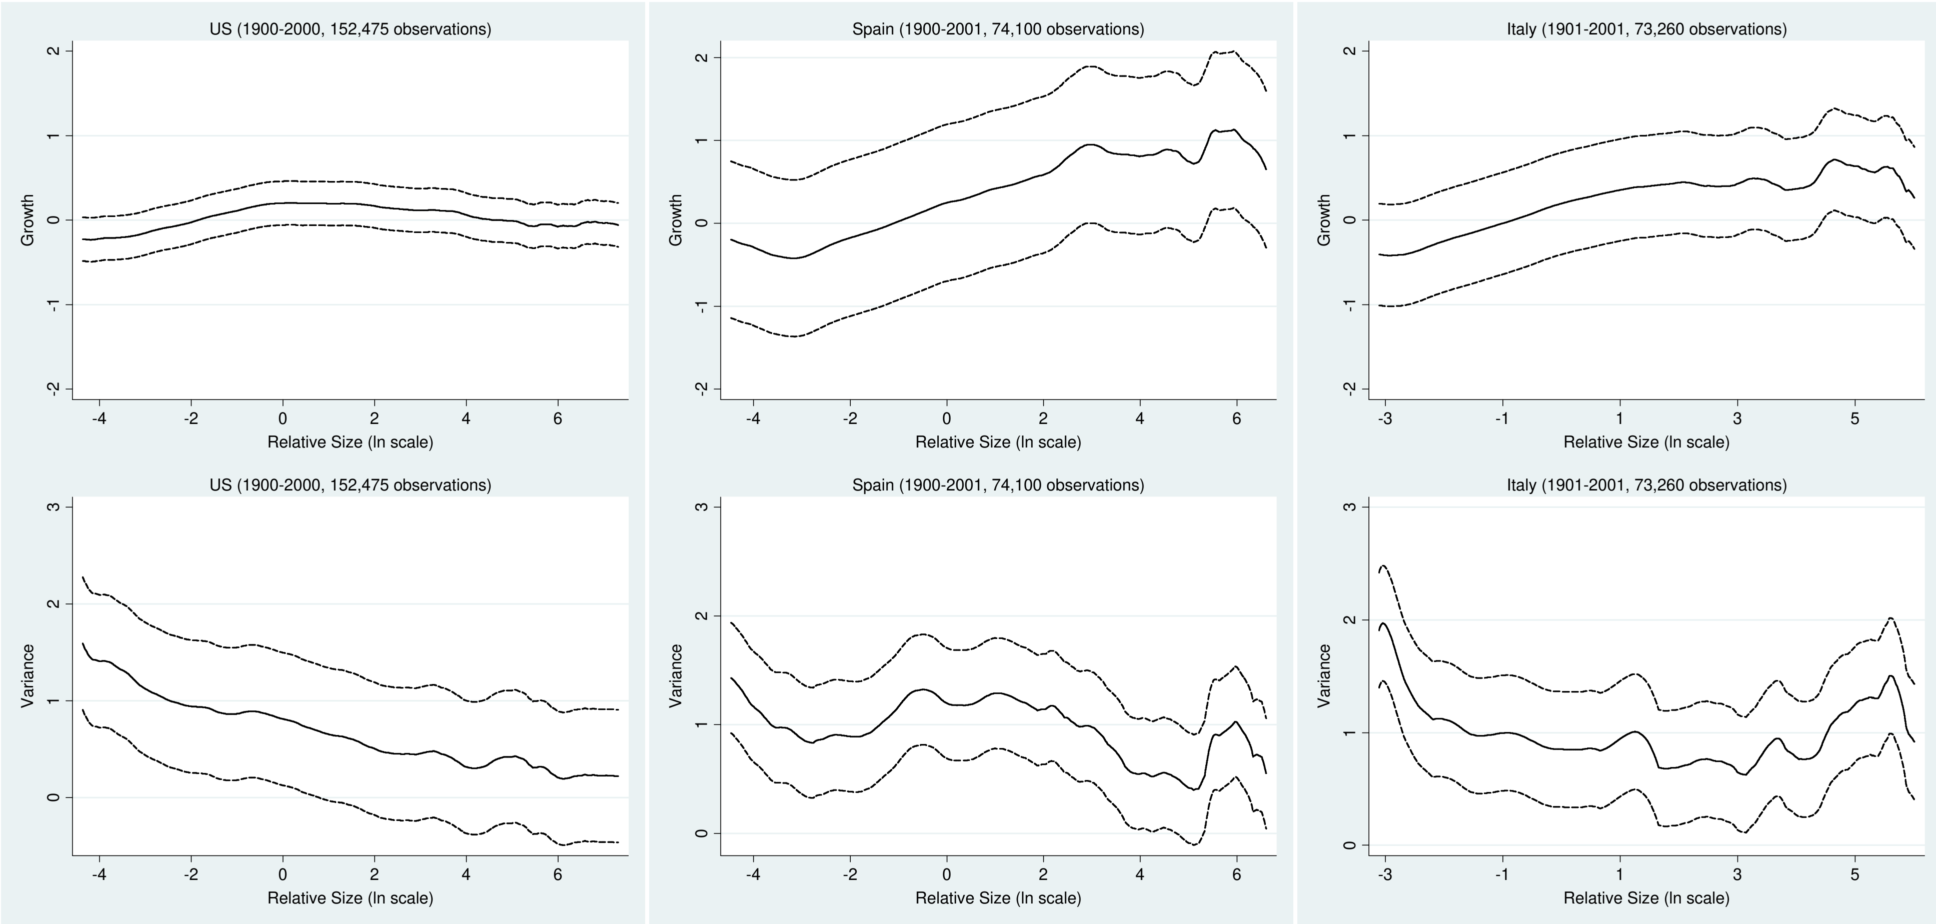

Figure 6. Comparison of the Estimated Density Function (ln scale) and the Theoretical Lognormal (dotted line) for the US in 1900, 1950 and 2000

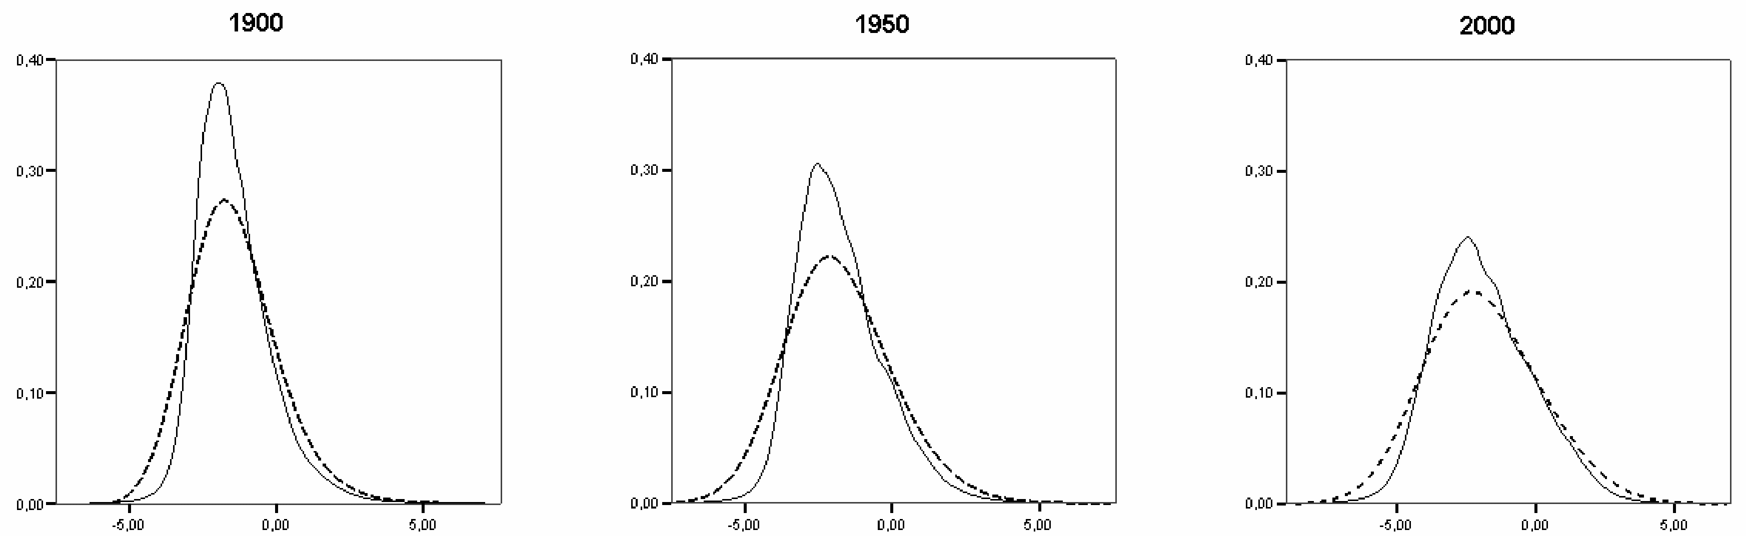

Figure 7. New Entrants Nonparametric Estimates (bandwidth 0.5) of the Growth Rate and its Variance, (US, 1910–2000), 59,865 observations

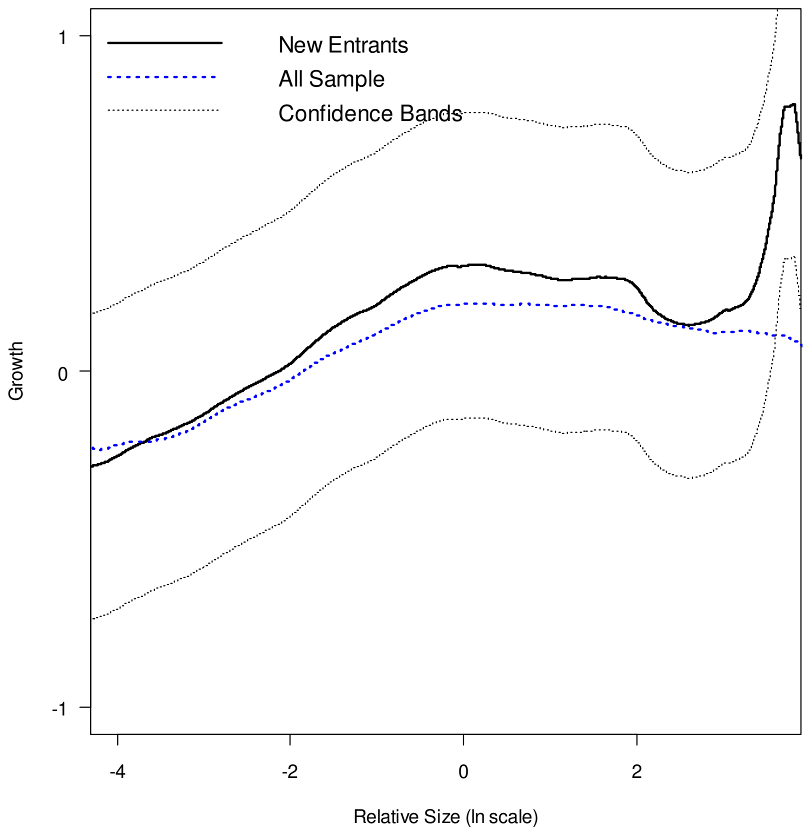

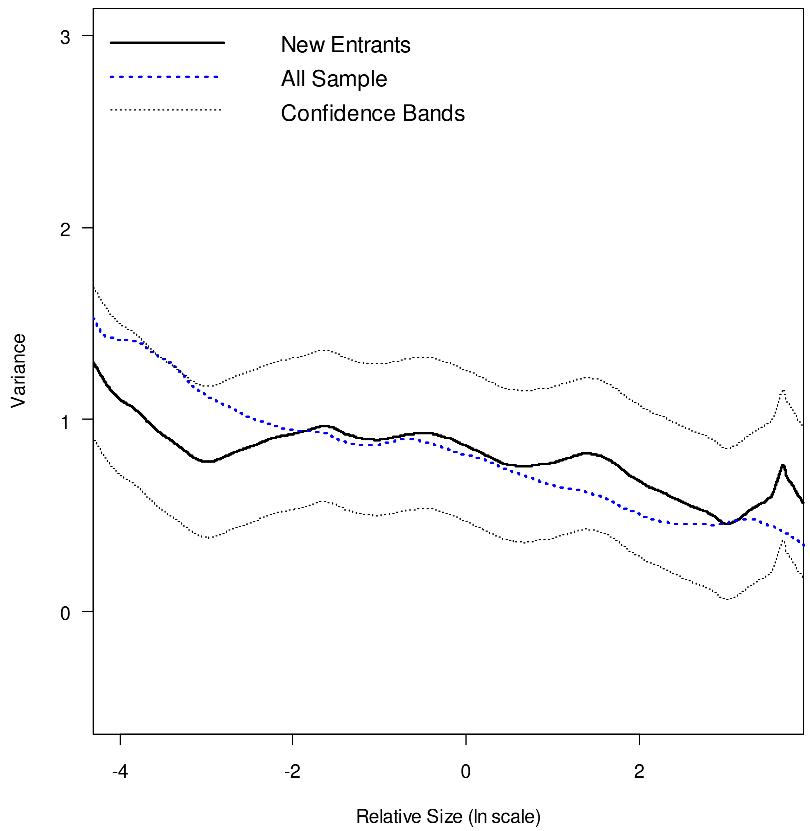

Figure 8. Empirical Density Functions of the New Entrants for the US in 1910 and 2000