

DISCUSSION PAPER SERIES

IZA DP No. 15995

**Was Robert Gibrat Right? A Test Based on
the Graphical Model Methodology**

Marco Guerzoni
Luigi Riso
Marco Vivarelli

MARCH 2023

DISCUSSION PAPER SERIES

IZA DP No. 15995

Was Robert Gibrat Right? A Test Based on the Graphical Model Methodology

Marco Guerzoni

*DEMS, University of Milano-Bicocca, BETA
and University of Strasbourg*

Marco Vivarelli

*Catholic University Milano, UNU-MERIT
and IZA*

Luigi Riso

Catholic University Milano

MARCH 2023

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Was Robert Gibrat Right? A Test Based on the Graphical Model Methodology*

Using both regression analysis and an unsupervised graphical model approach (never applied before to this issue), we confirm the rejection of the Gibrat's law when our firm-level data are considered over the entire investigated period, while the opposite is true when we allow for market selection. Indeed, the growth behavior of the re-shaped (smaller) population of the survived most efficient firms is in line with the Law of Proportionate Effect; this evidence reconciles early and current literature testing Gibrat's law and may have interesting implications in terms of both applied and theoretical research.

JEL Classification: L11

Keywords: Gibrat's Law, firm survival, market selection, firm growth

Corresponding author:

Marco Vivarelli
Dipartimento di Politica Economica
Università Cattolica del Sacro Cuore
Largo Gemelli 1
20123 Milano
Italy
E-mail: marco.vivarelli@unicatt.it

* Marco Vivarelli acknowledges the support by the Italian Ministero dell'Istruzione, dell'Università e della Ricerca (PRIN-2017, project 201799ZJ5N: "Technological change, industry evolution and employment dynamics"; principal investigator: Marco Vivarelli).

1. Introduction

The standard interpretation of the Law put forward by Gibrat (1931) is that the growth rate of a given firm is independent of its initial size. However, while earlier studies - based on limited samples of well-established and large companies - confirmed the Law, starting from Mansfield (1962) subsequent and recent research has rejected it (see next section). Indeed, the current consensus within the extant empirical literature points out that smaller firms show a higher growth rate than their larger counterparts.

One way to approach this puzzle is to take into account that earlier studies focused on companies which were the outcome of a previous (not investigated) market selection and so represented the industrial “core” within which Gibrat’s Law tended to be confirmed. On the other hand, the current literature - based on more comprehensive and large datasets including newborn and small firms - tests the law investigating a given population of the same firms *over* time and in doing so magnifies the role of smaller and younger firms (which must grow faster in order to reach a minimum efficient size and survive), therefore rejecting the law.

The purpose of this study is to test whether a given population of firms tends to converge to a Gibrat-like behavior *through* time, allowing for market selection and for the correlated exit of the less efficient firms. In this context - and differently from the studies discussed above - the final population is smaller than the initial one and it is made by the sole survived, most efficient companies. In this setting, it may well be the case that Gibrat’s Law is rejected when considering the entire population of firms and the entire period examined (consistently with the current literature); while the law is confirmed when considering the sole tracked population of survived firms (consistently with the earlier literature)¹. In other words, the hypothesis tested in this work is that Gibrat’s Law, although rejectable in general, might actually be accepted when market selection generates a sort of “steady state”, where a much more homogeneous population of survived firms may behave according to the law.

Another important novelty of this study is methodological. While the earlier and current empirical literature have used econometric techniques, in this work we use an innovative unsupervised approach generating graphical models able to elicit the intrinsic structure of the data and represent them as a network (see next section). Standard econometric analysis is however proposed as a preliminary analysis.

¹ This competitive dynamics is well captured by the seminal theoretical model put forward by Jovanovic (1982), based on Bayesian passive learning, and by the models with active learning (Ericson and Pakes, 1995; Pakes and Ericson, 1998).

2. Past and current literature

In 1931 the French engineer and economist Robert Gibrat put forward his (now world-wide well-known) law of proportionate effect stating that the proportional rate of growth of a given company is independent of its absolute size at the beginning of the investigated period (since then called Gibrat's law, or rule of proportionate growth, see Gibrat, 1931)².

After the second world war, the Gibrat's Law of Proportionate Effect was very popular both among economists and statisticians (Santarelli, Klomp and Thurik, 2006). The main reason was that the Law was fully consistent with a log-normal distribution of firm size (or even considered as the data generation process behind such a distribution). In turn, a log-normal distribution of firm size was (and it is nowadays) actually observed in virtually all the economic sectors, where a vast majority of small- and medium-sized firms coexist with few larger counterparts.

Therefore, as stated by Simon and Bonini (1958), if one "...incorporates the Law of proportionate effect in the transition matrix of a stochastic process, [...] then the resulting steady-state distribution of the process will be a highly skewed distribution" (*ibidem*, p.609).

The empirical consistency between Gibrat's Law and the observed size distribution of firms across different industries was also discussed by Steindl (1965) and treated through examples and simulations by Prais (1976, Chapter 2).

However, although in the long-term Gibrat's Law surely generates a log-normal firm size distribution, the latter does not necessarily require firms' proportional rates of growth. Indeed, if we do not limit our attention to the incumbent firms but we extend our focus to the analysis of industrial dynamics – that is entry and exit of companies within a given industry – a log-normal distribution may emerge as well as the consequence of a small group of persisting larger incumbents (core), coexisting with a large fringe of smaller firms, characterized by churning and turbulence (high entry rates, low survival rates, revolving-door firms, see Geroski, 1995). In other words, Gibrat's law is a sufficient, but not necessary condition to generate an observable log-normal distribution of firm size. This argument allows the possibility to falsify the law, without being in contrast with the revealed skewed distribution of firm size (see below).

Indeed, while earlier studies based on subsamples of large and mature firms had tended to confirm the Law (Hart and Prais, 1956; Simon and Bonini, 1958; Hymer and Pashigian, 1962), further research began to challenge its overall validity.

It is important to note that earlier studies were based on limited databases, only comprising large incumbents, namely companies quoted in the London Stock Exchange in Hart and Prais (1956); the largest 500 *Fortune* US corporations in Simon

² Edwin Mansfield, in his seminal paper on the *AER*, describes Gibrat's law with the following words: "the probability of a given proportionate change in size during a specified period is the same for all firms in a given industry - regardless of their size at the beginning of the period" (Mansfield, 1962, p. 1031).

and Bonini (1958); the 1,000 US largest manufacturing firms in the period from 1946-1955 in Hymer and Pashigian (1962). In other words, the earlier consensus about the validity of Gibrat's law was based on empirical tests limited to the core of larger incumbent companies, so neglecting the role of both incumbent SMEs and newborn firms.

The turning point in the literature was the seminal contribution by Mansfield (1962), investigating the U.S. steel, petroleum and tires sectors in different time periods and finding that Gibrat's Law was failing in the majority of cases, with smaller firms growing faster than their larger counterparts. Indeed, when studies take into account the fringe of SMEs, they incorporate their need to reach the minimum efficient size (MES) and so their engagement into an accelerated growth.

Mansfield's outcome has been largely confirmed by subsequent empirical studies, using more comprehensive specifications and also including firm's age and other controlling regressors.

For instance, Hall (1987) studied 1,778 US manufacturing firms which had already reached a certain minimum size (measured in terms of employment) and belonged to two samples spanning the periods 1972-1979 and 1976-1983. Unlike Mansfield (1962), Hall directly regressed growth rates on the logarithm of the initial size and found that the observed negative relationship between size and growth was robust to corrections for both sample attrition and heteroskedasticity³.

Evans (1987a) analyzed 100 4-digit manufacturing industries using firm level data drawn from the US Small Business Data Base (42,339 firms). The novel feature of this study was the introduction of age as a possible factor - in addition to size measured in terms of employment - in explaining departure from Gibrat's Law. A negative relationship between growth and size was found in 89 per cent of the industries examined, while a negative relationship between growth and age was verified in the 76 per cent of the industries. Like the previous study, the estimation procedure controlled for sample selection bias and heteroskedasticity (see also Evans, 1987b).

The work put forward by Dunne, Roberts and Samuelson (1989) also supported the rejection of Gibrat's law found in the previous studies: within each age category, growth rates turned out to decline along employment size classes. Dunne, Roberts and Samuelson obtained these results from data on 219,754 individual plants - rather than firms as in the previous studies - collected in five US censuses of manufactures (1963-67-72-77-82).

Another important contribution to the investigation of Gibrat's Law was put forward by Dunne and Hughes (1994), who tested the Law of Proportionate Effect over the periods 1975-80 and 1980-85 using 2,149 quoted and unquoted UK companies belonging to 19 different manufacturing industries. After controlling for sample attrition and heteroskedasticity, Dunne and Hughes found further confirmation that

³ This type of econometric specification will also be adopted in the present study, see next section, eq. (3).

smaller companies tend to grow faster than their larger counterparts; they also found that younger companies, for a given size, tended to grow faster than elder ones.

By the same token, Hart and Oulton (1996) used data comprising 87,109 UK incumbent companies over the period 1989-93 and tested the Chesher-Mansfield specification (see next section) measuring size in terms of employment, sales and net assets. In all cases, they detected an overall estimated coefficient of less than one: on average, small firms grew more quickly than larger ones; however, they also found a not significant relationship between growth and size when considering the sole larger firms. In other words, Gibrat's law turned out to be rejected in general, but not falsified within the subsample of the core companies (see above).

Audretsch, Santarelli and Vivarelli (1999) used an Italian dataset comprising newborn manufacturing firms tracked since 1987 to 1993 and found that Gibrat's law was indeed rejected in the vast majority of industries, both considering the entire set of firms and the limited set of survived firms.

On the whole, at the end of the '90s, a new consensus was reached, partially in contrast with the one shared in the previous decades: the conclusion was that "*Gibrat's Legacy*" (as named by Sutton, 1997; see also Caves, 1998 and Coad, 2009) was defensible not as a general law, but only as a dynamic rule valid for large and mature firms that had already attained the MES level of output, but not for their smaller counterparts, operating at a sub-optimal scale (Geroski, 1995). In a nutshell, and combining the two consensus reached by the literature, Gibrat's law should be considered rejected when all firms are taken into account, but confirmed when the sole core companies within industries are considered.

The most recent literature has generally supported this overall conclusion. For instance, Calvo (2006) - analyzing 1272 Spanish manufacturing firms over the period 1990–2000 - found smaller firms growing faster than larger ones. By the same token, Oliveira and Fortunato (2006) - using an unbalanced panel of Portuguese manufacturing firms over the period 1990- 2001- found that large and mature firms do have smaller growth rates than small and young firms.

Daunfeldt and Elert (2013) studied Swedish firms within five-digit NACE-industries during the period 1998–2004 and confirmed the rejection of Gibrat's law when considering the entire population of the investigated companies; however, Gibrat's law was more likely to be rejected for industries characterized by a higher MES, while the law was more likely to hold in mature industries, in industries with a high degree of group ownership, and in industries with a high market-concentration.

Tang, A. (2015) studied the Swedish energy market - using an unbalanced longitudinal dataset covering 2,185 firms during the 1997–2011 period - and found out an interesting twofold result: on the one hand, Gibrat's law was rejected, with smaller firms found to grow faster than their larger counterparts; on the other hand, when examining each firm individually, they found that many Swedish energy firms behave

in accordance with Gibrat's law, namely the ones in *steady state*, that is the larger and more mature ones.

Distante, Petrella and Santoro (2018) run quantile regression models using annual data covering US manufacturing firms over six decades (1950–2010) and found that, conditional on survival, small establishments grow faster than their larger counterparts.

Arouri et al. (2020) studied the pattern of growth of Tunisian firms over the period 1996–2010: their key finding was that, consistently with the extant literature, Gibrat's law was overall rejected with smaller firms growing faster than their larger counterparts; however, the negative impact of the initial size was found larger and more significant for young firms rather than for mature larger incumbent firms.

Elston and Weidinger (2023) investigated MENA companies listed in the United Arab Emirates (UAE) stock exchanges and found out that, in most industries, smaller firms grow faster than larger firms with three notable exceptions: energy, telecommunications and industrial manufacturing.⁴

All in all, the extant literature seems to support the general idea that Gibrat's law should be rejected when all firms are taken into account, but can be revived when the core of larger and older incumbents is singled out. As mentioned in the previous section, the purpose (and the novelty) of this paper is to test whether a given population of firms tends to converge to a Gibrat-like behavior through time, allowing for market selection and for the correlated exit of the less efficient firms. In particular, and differently from the previous studies discussed above, we will start from a brand-new population (1720 newborn Italian manufacturing firms) and will track them over 11 years to test whether a convergence to a Gibrat-like behavior emerges over time. Instead of separating *different groups of firms* (that is core vs fringe), we will deal with the *same population of companies* over time, allowing for market selection and so for the exit of the less efficient firms. The purpose being to investigate whether earlier literature can be reconciled with more recent research, that is to test whether the rejection of Gibrat's Law *ex ante* can be coupled with the defense of the Law *ex post* (see the hypotheses proposed in the next section).

To our knowledge, only two previous studies tried to put forward this kind of experiment. In a first work, Lotti, Santarelli and Vivarelli (2003) run quantile regressions using data for 855 Italian manufacturing firms founded in January 1987 and tracked for six years; their main result is that in five industries out of six, Gibrat's Law fails to hold in the years immediately following start-up, whereas it holds, or fails less severely, when firms approach maturity and market selection has done its job. In a later study, the same authors (Lotti, Santarelli and Vivarelli, 2009) focused on the Italian radio, TV, and communication equipment industry over the period 1987-1994, studying the growth patterns of all the incumbent firms which were active in the sector at the beginning of the examined period (3,285 companies). Consistently with the

⁴ Interestingly enough, these industries are characterized by larger MES and the dominant role of core companies, particularly within a sample of listed firms.

former study, results are twofold: on the one hand, Gibrat's Law is rejected over the entire period with smaller firms growing faster than larger ones; on the other hand, a convergence toward the validity of the Law occurs through time, once the annual regressions are run over the sub-population of survived firms.

In what follows, the econometric regressions are similar in nature to what done in Lotti, Santarelli and Vivarelli (2003 and 2009), while the graphical model approach is applied for the first time in testing Gibrat's law, at least to our knowledge.

3. Data, hypotheses and preliminary econometric specification

The analysis is based on AIDA-BvD data, which contains comprehensive information on all Italian firms required to file accounts. Specifically, we acquired a dataset comprising the entire population⁵ of 1720 newborn Italian manufacturing firms (with at least one employee) founded in 2009 and tracked for 11 years, namely until 2020. We selected the following variables of interest: Employees (E); Regional belonging (dummies corresponding to the NUTS-2 classification, R); Sectoral belonging (dummies corresponding to the 2 digits NACE classification, S); Innovativeness (a dummy I that indicates whether a firm is registered as "innovative" according to the Italian decree "innovative firms act 221/2012", see Guerzoni et al., 2020); Profitability, computed as the ratio between "earnings before taxes" and "revenues from sales and services" (P).

By utilizing data on newly established Italian manufacturing firms in 2009, the validity of Gibrat's Law will be examined over the entire period of 2010-2020, as well as year-by-year. Through this set of analyses, we aim to jointly test the following hypotheses, aimed to possibly reconcile the diverging evidence discussed in Section 2:

(H1) Gibrat's Law is rejected over the entire period (a priori hypothesis);

(H2) A convergence towards a Gibrat-like steady state emerges among the population of surviving firms (a posteriori hypothesis).

If both these hypotheses were to be confirmed, this would mitigate the apparently controversial debate surrounding the validity of Gibrat's Law (see Section 2). While the Law may be rejected when examining the overall evolution of a given ex-ante population of companies (a priori hypothesis), it may still accurately describe the patterns of growth for well-established firms within the ex-post sub-population that results from market selection and learning processes (a posteriori hypothesis).

The specification used to test Gibrat's Law econometrically is the following (the same adopted by Evans, 1987a and 1987b, and Lotti *et al.*, 2009):

⁵ Since we are dealing with an entire population, our dataset cannot be affected by sample selection.

$$G_{i,t} = \beta_0 + \beta_1 \log(E_{i,t-1}) + \beta_1 R_i + \beta_1 S_i + \beta_1 \log(P_{i,t-1}) + \beta_1 I_i + \varepsilon_{i,t} \quad (1)$$

Where $G_{i,t} = (E_{i,t} - E_{i,t-1})/E_{i,t-1}$ is the employment growth rate, of the firm i at time t , $E_{i,t-1}$ is the number of employees of the firm i at time $t-1$. Regional dummies, sectoral dummies, profitability (P) and innovativeness (I) act as controls⁶.

Chesher (1979) pointed to the coefficient β_1 in order to test the validity of the Law through the significance of the relevant parameter. In particular, if $\beta_1=0$, Gibrat's Law holds; if $\beta_1 < 0$, smaller firms grow at a higher rate than their larger counterparts, while the opposite case if $\beta_1 > 0$.

We estimate Eq. (1) in each period t on the subsamples of firms still alive in t ; moreover, we estimate the overall employment growth over the entire investigated period, with growth $G_{i,t} = -100\%$ for firms that have exited the market (as in Evans, 1987a and 1987b).

Table 1 presents the summary statistics for each period, displaying the progressive reduction of the population as well as the corresponding growth rate and other variables included in the model for each year. It is noteworthy that only 688 companies out of 1720 survive until the end of the investigated period, pointing to a ten year surviving rate in Italian manufacturing equal to 40% (this is not surprising and in line with the stylized facts pointed out by Geroski, 1995). In order to survive, newborn firms must grow, with an average size moving from 14.10 employees at the beginning of the period to 21.40 at the end of the period. Profitability is hard to be reached by these young firms although more likely in the later years and innovativeness (as defined in this study) is really an exceptional attribute.

⁶ Controlling for age is obviously useless in our context, since we are dealing with a sole cohort, namely companies founded in 2009.

Table 1: Descriptive statistics

Datasets	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	All	
<i>N.</i>	1720	1573	1403	1259	1126	1030	942	881	798	726	688	1720	
G	<i>Mean</i>	0.20	-0.04	0.06	0.00	0.03	0.05	0.01	0.05	0.01	0.03	0.00	-0.31
	<i>St. Dev.</i>	0.70	0.51	0.50	0.37	0.41	0.36	0.30	0.52	0.23	0.31	0.20	1.82
	<i>Min</i>	-0.85	-0.98	-0.99	-0.99	-0.98	-0.93	-0.97	-0.94	-0.96	-0.88	-0.75	-1.00
	<i>Max</i>	18.00	4.00	9.00	5.00	4.40	3.50	2.58	13.00	2.29	5.90	1.34	35.75
E	<i>Mean</i>	14.10	15.50	15.13	15.47	15.64	16.25	17.47	18.60	19.63	20.64	21.40	14.10
	<i>St. Dev.</i>	28.31	24.35	26.57	26.81	28.14	30.08	32.23	35.24	37.69	41.82	46.12	28.31
	<i>Min</i>	1	1	1	1	1	1	1	1	1	1	1	1
	<i>Max</i>	450	436	472	480	480	539	562	587	603	636	679	450
P	<i>Mean</i>	-0.22	-0.02	-0.50	-0.11	-1.34	0.14	-0.34	0.05	-0.07	0.01	21.44	-0.22
	<i>St. Dev.</i>	2.97	0.69	17.87	2.15	32.42	6.66	3.99	2.81	1.68	0.31	562.38	2.97
	<i>Min</i>	-68.34	-19.34	-669.07	-73.20	-1045.74	-11.92	-87.17	-11.78	-44.55	-6.16	-5.40	-68.34
	<i>Max</i>	7.62	10.51	1.00	3.68	1.71	212.52	1.69	81.10	6.09	1.11	14751.00	7.62
I	<i>Yes</i>	5	5	5	5	5	5	5	4	4	4	4	5
	<i>No</i>	1715	1568	1398	1254	1121	1025	937	877	794	722	688	1715

4. A novel graphical model approach

As one of the main contributions of this study, we make use of unsupervised graphical models (GM) to gain a different and more comprehensive perspective on whether a given population of firms tends to converge to a Gibrat-like behavior through time. In fact, GM allow us to jointly describe the overall structure of dependency among variables and in this way we can capture multiple relationships, including non-linear and conditional dependencies, and both direct and indirect influences. In particular - in our context - GM can depict not only whether the proxy for size directly affects growth, but also any other mediating effects.

In more detail, graphical models are a framework that combines network representation and probability theory to specify conditional independence relationships between random variables in a given dataset. These relationships are represented through a graphical representation, specifically a graph $\mathcal{G}(V,L)$, where V is a finite set of nodes corresponding to the variables of interest and L is the set of links in the network, representing the conditional dependence between any pair of variables (Lauritzen, 1996).⁷

One of the central problems in the representation of GM is the estimation of the underlying probability distributions of the variables from a finite sample. Chow and Liu, 1968 proposed an approach for discrete variables that approximates their probability functions via probability distribution of the first-order tree dependence. The connection between nodes of the tree represents the unknown joint probability of the nodes (or associated variables), providing information about their mutual dependence or mutual information. In detail, Chow and Liu, 1968 found out that a probability distribution of a tree dependence approximates the true value probability of a set of discrete random variables composing the tree, if and only if the latter has maximum mutual information. Under the assumption that the cell probabilities of discrete random variables factorize according to an unknown tree τ written as $\mathcal{G}_D = (\Delta, L_\Delta)$, they can be written as:

$$p(d|\tau) = \frac{\prod_{u,v \in E_\Delta} p(d_u, d_v)}{\prod_{v \in \Delta} p(d_v)^{d_v - 1}} \quad (2)$$

⁷ In order to explain better the following analysis, we introduce a mixed dataset, X , composed of n observation and p variables. We split the variables into r discrete, $D = (D_1, \dots, D_r)$ and q continuous $C = (C_1, \dots, C_q)$. Denote the i -observation of $X = (D, C)$ as $[(d)_i, (c)_i]$ with d_i and c_i representing the i -observation of the variables $D_i \in D$ and $C_i \in C$, respectively. Given the one-to-one correspondence between variables and nodes, we can write the set of the nodes as $V = \{\Delta, \Gamma\}$. Where Δ and Γ are the nodes corresponding to the variables in D and C , respectively.

Where d_v is the number of links incident to the node v , namely the degree of v . According to Eq.(2), the maximized log-likelihood, up to a constant, turn out to be $\sum_{(u,v) \in L_\Delta} I_{u,v}$, where $I_{u,v}$ is the mutual information between nodes u and v . It is worth noting that the mutual information between two variables is defined as a measure of their closeness (Lewis II, 1959), therefore, is a dimensionless, nonnegative, and symmetric quantity which measures the reduction of uncertainty about a random variable, given the knowledge of another.

Prior to the development of the Chow-Liu algorithm, which has been extensively studied, various other algorithms were created to determine the probabilistic structure and corresponding maximum-likelihood estimator. Specifically, Kruskal (1956) proposed a simple and efficient solution to this problem by starting with a null graph and adding the edge with the highest weight at each step, as long as it does not form a cycle with previously chosen edges. Edwards et al. (2010) extended the Chow-Liu algorithm to be applied to mixed data sets \mathbf{X} , using mutual information between discrete and continuous variables. This algorithm relies on the use of mutual information between a discrete variable, D_u , and continuous variable, C_v . It is characterized by the marginal model which results to be an ANOVA model (Edwards, 2012, section 4.1.7). It is worth noting that, when dealing with mixed variables, the evaluation of the mutual information $I(d_u, c_v)$ between each couple of nodes requires distinguishing between the case when the variance of C_v is distributed homogeneously across the level of the discrete variable D_u , from the case when it is heterogeneously distributed (Edwards, 2012). As pointed out by Edwards et al. (2010), one of the disadvantages of selecting a tree on the basis of maximum likelihood is that it always includes the maximum number of edges, even if the latter are not supported by data. Thus, they suggested the use of one of the following measures to avoid this drawback:

$$I^{BIC} = I(x_i, x_j) - \log(n)k_{x_i, x_j}; \quad I^{AIC} = I(x_i, x_j) - 2k_{x_i, x_j} \quad (3)$$

The degrees of freedom associated with the pair of variables, x_i and x_j , are represented by k_{x_i, x_j} , and are determined based on the nature of the variables involved, continuous or discrete⁸. These measures are used in an algorithm proposed by Edwards et al. (2010) to determine the best spanning tree. The algorithm stops when the graph has reached its maximum number of edges. The algorithm can generate either a tree or a forest, where a forest is a group of trees.

⁸ For discrete random variables, the degrees of freedom are equal to $|D_u| - 1$, where D_u is the number of levels of the discrete random variable. However, for continuous random variables, there is only 1 degree of freedom. Under marginal independence, the statistic $I_{u,v}$ has an asymptotic χ^2 distribution (Edwards et al. 2010).

To test the validity of Gibrat's Law, we employ the extension of the Chow-Liu algorithm (Chow and Liu, 1968) proposed by Edwards et al. (2010) for mixed datasets. This methodology allows us to map the conditional dependence relationships of the variables involved in the Eq. (2) into a graph $\mathcal{G} (V,L)$, where V is a finite set of nodes with direct correspondence to the variables of interest and L is the set of links in the network (Lauritzen, 1996). The links represent the conditional dependence between any pair of variables.⁹ Specifically, the GM employed in this paper belong to the class of multivariate distributions, whose conditional independence properties are encoded in a tree/forest in the following way: the absence of a link between two nodes represents conditional independence between the corresponding variables (Jordan, 2004).

In the context of this study, if there is a direct connection between the node G (employment growth rate at time t) and the node E (number of employees at time $t-1$) or the connection is mediated by another node, Gibrat's Law does not hold. Conversely, if the node G is not connected with node E , Gibrat's Law holds.

Moreover, the GM methodology allows us to understand how the relationships between the variables involved in the model change over time; in particular, we built both a GM for each year and one to test the overall relationships over the entire period; therefore, we exactly mimic the econometric setting put forward in the previous section.

5. Empirical findings

Table 2 presents the output of the regressions, which can be considered as our preliminary baseline. In the last column, which displays the results over the entire investigated period, the key coefficient ($\log E$) is negative and significant, rejecting Gibrat's Law and supporting Hypothesis 1¹⁰.

However, the regressions on single periods tell a different story: the initial size displays a significant and negative impact only in the first seven years, whereas - allowing for market selection - this significance disappears since 2017 (supporting Hypothesis 2).

As far as the controls are concerned, while profitability is boosting growth, innovation does not seem to play a significant role.¹¹

⁹It is important to note that we cannot take the magnitude of these links into account, but only their presence or absence, which defines the structure of the tree itself (Riso and Guerzoni, 2022).

¹⁰ In line with the recent literature (see Section 2).

¹¹ This outcome can be due to the very small incidence of innovative firms within our population.

Table 2: regression analyses; dependent variable: Employment Growth Rate

	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	All
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
<i>Intercept</i>	0.510^{***} (0.077)	0,052 (0.071)	0.269^{***} (0.082)	0,074 (0.057)	0.151^{**} (0.074)	0,064 (0.058)	0.139^{**} (0.068)	0,113 (0.116)	-0,029 (0.041)	-0,046 (0.048)	-0,045 (0.041)	0,013 (0.2)
<i>log (E)</i>	-0.156^{***} (0.029)	-0.041^{**} (0.016)	-0.071^{***} (0.017)	-0.064^{***} (0.015)	-0.059^{***} (0.014)	-0.037^{***} (0.012)	-0.023^{**} (0.01)	-0,043 (0.027)	-0,011 (0.008)	-0,016 (0.011)	0,004 (0.008)	-0.174^{***} (0.063)
<i>I</i>	-0,02 (0.083)	-0,05 (0.112)	-0,002 (0.086)	0,168 (0.107)	0,059 (0.082)	0.095 ^{**} (0.037)	-0,047 (0.065)	0,104 (0.102)	-0,003 (0.059)	0,053 (0.072)	0,02 (0.035)	0,726 (0.642)
<i>log (P)</i>	0,077 (0.073)	0.095[*] (0.051)	0.209^{***} (0.057)	0.297^{***} (0.049)	0.241^{***} (0.049)	0.133^{***} (0.046)	0.138^{***} (0.037)	0.138[*] (0.079)	0.174^{***} (0.04)	0.162^{***} (0.051)	0.071^{**} (0.034)	0.335^{**} (0.144)
<i>Regional dummies</i>												
<i>Sectoral Dummies</i>								YES				
<i>Observations</i>	1720	1573	1403	1259	1126	1030	942	881	798	726	688	1720

Note: robust standard errors reported in brackets; * significant at 10%; ** significant at 5%; *** significant at 1%; significant coefficients in bold.

The graphical models in Fig. 1 corroborate this picture and provide additional insights. Fig.1 depicts the graphical model, with Growth (G) and Employment (E) highlighted in green and blue, respectively.

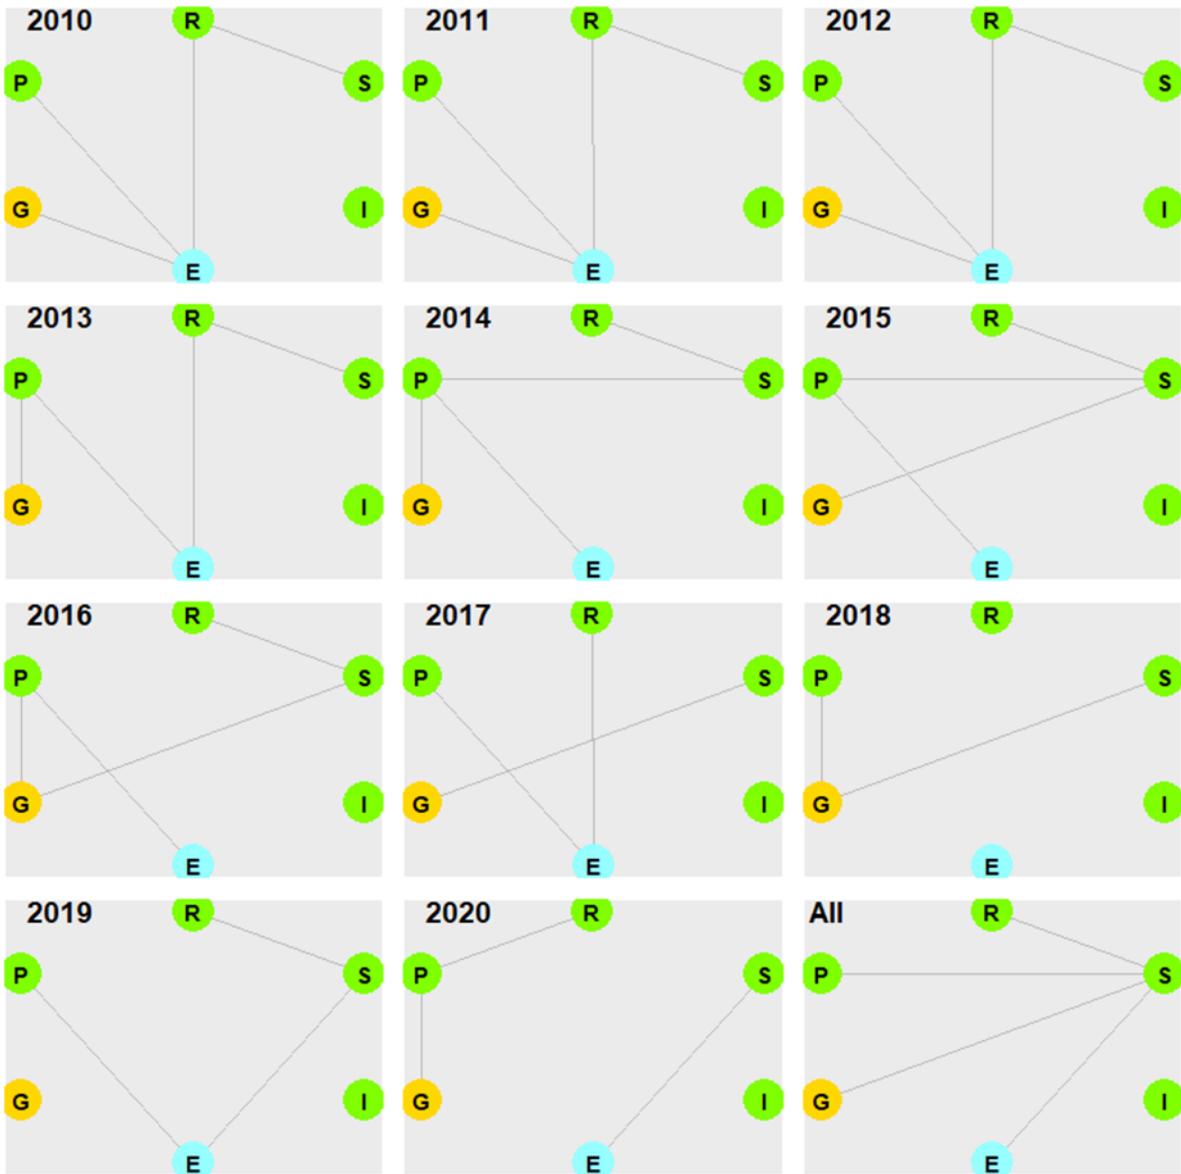
Over the entire period (Figure 1, last panel: All) growth is associated with the initial size, through the mediation of sectoral belonging. This implies that for each given sector, growth and size are not independent. Consequently, Gibrat's law is rejected, since size does have an effect on growth, albeit varying across sectors¹² (this evidence supports Hypothesis 1).

However, shifting to the annual analyses, results present a different picture. In the first three periods, growth is directly linked to employment: even independently from sectoral belonging, size does influence the growth rate of firms. In the subsequent four periods, this relationship persists, but mediated by profitability (and sectoral belonging in 2015). This evidence suggests that Gibrat's Law is initially rejected but becomes stronger overtime. Finally, as it was the case for the annual regressions, since 2017 Growth and Employment are not connected any longer, corroborating the interpretation of Gibrat's Law as a steady state convergence (thus supporting Hypothesis 2).

Remarkably, the results from the graphical models are not only aligned with the regression results, but they also provide a more intuitive and interpretable representation of the data. Specifically, we are able to observe an initial direct correlation between growth and size, then a temporary mediation effect of profitability, while eventually any type of either direct or indirect dependency fades away as market selection proceeds

¹² The mediating role of sectoral belonging might be due to the different "minimum efficient sizes" required by the different sectors (while interesting, this research perspective is beyond the scope of the present work and cannot be tested, given the available data).

Figure 1: Graphical Models



6. Conclusions

As discussed in detail in Section 2, the extant (partially controversial) literature has come to the conclusion that Gibrat's law can be rejected when all firms are taken into account, but can be proved when the core of larger and older incumbents is isolated.

Differently from most of previous studies, in this paper we did not single out different groups of firms (that is core *vs* fringe), but we tracked a brand-new population of companies over time, allowing for market selection and so testing whether a given population of firms tends to converge to a Gibrat-like behavior through time.

Using both standard econometrics and a novel unsupervised approach generating graphical models, this paper showed that the early and current literature testing Gibrat's law can be indeed reconciled, in particular the rejection of Gibrat's Law *ex ante* can be coupled with the defense of the Law *ex post*.

In more detail, while we confirmed the rejection of the law when firms were considered over the entire investigated period, we obtained the opposite when we allowed for market selection and we tracked the sole survived companies. Indeed, the growth behavior of the re-shaped (smaller) population of the survived most efficient firms was in line with the Law of Proportionate Effect.

This twofold evidence may have interesting implications in terms of both applied and theoretical research on the one hand and policy options on the other hand. In particular, policy makers should take into account that employment growth crucially depends on the combination of different factors characterizing industrial dynamics, such new firm formation, firm size, survival rates and market selection.

References

- Arouri, H., Youssef, A.B., Quatraro, F. and M. Vivarelli (2020). “Drivers of growth in Tunisia: young firms vs incumbents”, *Small Business Economics*, 54, 323–340.
- Audretsch, D. B., Santarelli, E. and M. Vivarelli, (1999). “Start-up Size and Industrial Dynamics: Some Evidence from Italian Manufacturing”, *International Journal of Industrial Organization*, 17:965-983.
- Calvo, J.L., (2006). “Testing Gibrat’s Law for Small, Young and Innovating Firms”, *Small Business Economics*, 26:117-123.
- Caves, R. E, (1998). “Industrial Organization and New Findings on the Turnover and Mobility of Firms”, *Journal of Economic Literature*, 36:1947-1982.
- Chesher, A., (1979). “Testing the law of proportionate effect”, *Journal of Industrial Economics*, 27:403–411.
- Chow, C. and C. Liu, (1968). “Approximating discrete probability distributions with dependence trees”. *IEEE transactions on Information Theory*, 14:462–467.
- Coad, A. (2009). “The growth of firms: a survey of theories and empirical evidence”. Cheltenham, UK and Northampton, MA, USA: *Edward Elgar*.
- Daunfeldt, S.O. and N. Elert, (2013). “When is Gibrat’s law a law?”, *Small Business Economics*, 41:133–147.
- Distante, R., Petrella, I. and E. Santoro, (2018). “Gibrat’s law and quantile regressions: An application to firm growth”, *Economics Letters*, 164: 5-9.
- Dunne, P. and A. Hughes, (1994). “Age, size, growth and survival: UK companies in the 1980s”. *Journal of Industrial Economics*, 42: 115–140.
- Dunne, T., Roberts M. J. and L. Samuelson, (1989). “The Growth and Failure of US Manufacturing Plants”, *Quarterly Journal of Economics*, 104: 671-698.
- Edwards D., (2012), “Introduction to graphical modelling”, New York: *Springer Science & Business Media*.
- Edwards, D., De Abreu, G. C. and R. Labouriau, (2010). “Selecting high-dimensional mixed graphical models using minimal aic or bic forests”, *BMC bioinformatics*, 11:1-13.
- Elston, J.A. and A. Weidinger, (2023). “An empirical examination of firm growth in the MENA region through the lens of Gibrat’s law”, *Small Business Economics*, 60:121–131.
- Ericson, R. and A. Pakes, (1995). “Markov-Perfect Industry Dynamics: a Framework for Empirical Work”, *Review of Economic Studies*, 62: 53-82.

- Evans, D. S. (1987a). “The relationship between firm growth, size, and age: Estimates for 100 manufacturing industries”, *Journal of Industrial Economics*, 35:567–581.
- Evans, D. S., (1987b). “Tests of alternative theories of firm growth”. *Journal of Political Economy*, 95: 657–674.
- Geroski, P. A., (1995). “What do we know about entry?”. *International Journal of Industrial Organization*, 13, 421-440.
- Gibrat, R., (1931). “Les Inegalites Economiques”, Librairie du Recueil Sirey, Paris.
- Guerzoni, M., Nava, C.R. and M. Nuccio (2021). “Start-ups survival through a crisis. Combining machine learning with econometrics to measure innovation”, *Economics of Innovation and New Technology*, 30: 468-493.
- Hall, B., (1987). “The Relationship Between Firm Size and Firm Growth in the US Manufacturing Sector”, *Journal of Industrial Economics*, 3:583-606.
- Hart, P. E. and N. Oulton, (1996). “Growth and Size of Firms”, *Economic Journal*, 106: 1242-1252.
- Hart, P. E. and S.J. Prais, (1956). “The Analysis of Business Concentration: A Statistical Approach”, *Journal of the Royal Statistical Society*, 119: 150-191.
- Hymer, S, and P. Pashigian, (1962). “Firm Size and the Rate of Growth”, *Journal of Political Economy*, 70:556-569.
- Jordan, M.I., (2004). “Graphical models”, *Statistical science*, 19:140–155.
- Jovanovic, B., (1982). “Selection and Evolution of Industry”, *Econometrica*, 50: 649-670.
- Kruskal, Joseph B. (1956). “On the shortest spanning subtree of a graph and the traveling salesman problem”, *Proceedings of the American Mathematical society* 7.1: 48-50.
- Lauritzen, S. L., (1996). “Graphical Models”, *Oxford Statistical Science Series*. Oxford University Press: Oxford.
- Lewis II. P. M., (1959), “Approximating probability distributions to reduce storage requirements”, *Information and control*, 2: 214-225.
- Lotti, F., E. Santarelli and M. Vivarelli, (2003). “Does Gibrat’s Law Hold Among Young, Small Firms?”, *Journal of Evolutionary Economics*, 13, 213-235.
- Lotti, F., Santarelli, E. and M. Vivarelli, (2009). “Defending Gibrat's Law as a Long-Run Regularity”, *Small Business Economics*, 32:31-44.
- Mansfield, E., (1962). “Entry, Gibrat’s Law, Innovation, and the Growth of Firms”, *American Economic Review*, 52: 1023-1051.

- Oliveira, B. and Fortunato, A. (2006) “Testing Gibrat’s Law: Empirical Evidence from a Panel of Portuguese Manufacturing Firms”, *International Journal of the Economics of Business*, 13, 65-81.
- Pakes, A. and R. Ericson, (1998). “Empirical Implications of Alternative Models of Firm Dynamics”, *Journal of Economic Theory*, 79: 1-45.
- Prais, S. J. (1976). “The evolution of giant firms in Britain”, London: Cambridge University Press.
- Riso, L. and M. Guerzoni, (2022). “Concept drift estimation with graphical models”, *Information Sciences*, 606:786-804.
- Santarelli, E., L. Klomp and A. R. Thurik, 2006, “Gibrat’s Law: An Overview of the Empirical Literature”, in E. Santarelli (Ed.), *Entrepreneurship, Growth, and Innovation: the Dynamics of Firms and Industries*, Springer: New York, 41-73.
- Simon, H. A. and C. P. Bonini, (1958). “The Size Distribution of Business Firms”, *American Economic Review*, 58: 607-617.
- Steindl, J. (1965). “Random processes and the growth of firms: A study of the Pareto Law”, London: Griffin.
- Sutton, J., (1997). “Gibrat’s Legacy”, *Journal of Economic Literature*, 35: 40-59.
- Tang, A. (2015). “Does Gibrat’s law hold for Swedish energy firms?”, *Empirical Economics*, 49:659–674.