Sector, industry and inter-organizational movement statistics in the Stockholm Region: informing organizational growth models

Hernan Mondani^{1,2}

Published online: 3 July 2018
© The Author(s) 2018

Abstract

Organizational growth processes exhibit interesting statistical regularities, chiefly the heavy-tailed pattern of the size and growth-rate (i.e. yearly change in size) distributions. In spite of its ubiquity, empirical studies of growth are often limited to private activities and specific sectors, and generative models on the other hand are built on simplified assumptions and only aim at reproducing stylized facts. In this study, we use a unique Swedish longitudinal database on employment in the Stockholm Region, to analyze the interplay between organizational growth statistics by ownership sector, industrial activity and inter-organizational employee movements during a period of 14 years. We fit distributions for organizational size and growth rates. We find that the body of the aggregate growth-rate distribution is dominated by public sector growth, while the private sector dominates the tails. Industries with mostly public organizations tend to have a lognormal size distribution, while privately-owned industries are better fitted by a truncated power law. Growth-rate distributions are fitted to an exponential power (Subbotin) distribution. We decompose the change in size into incoming and outgoing employee movements, and find that the distribution of aggregated movements is well approximated by a lognormal distribution. Most organizations that do not grow have however in- and outgoing movements, but these mostly cancel each other out.

Keywords Organizational growth process · Heavy-tailed distributions · Power-law distribution · Exponential power distribution · Sector and industry dynamics · Inter-organizational movements

✉ Hernan Mondani
hernan.mondani@sociology.su.se

¹ Department of Sociology, Stockholm University, Stockholm, Sweden

² Present Address: Stockholm, Sweden

1 Introduction

1.1 Organizational growth statistics: empirical evidence

1.1.1 Size and growth-rate distributions

The size of an organization, in our case given by the number of employees, changes in time. The succession of relative changes in size is one way to describe an organization's *growth process*. The statistical analysis of organizational growth processes has given rise to a number of statistical regularities. The main results are that the distributions for organizational size and growth-rate—i.e. how fast the size changes from one time period to the next—are so-called *heavy-tailed* or *fat-tailed* distributions. The tails of the distribution give the probability for large-size organizations and large growth events of expansion and decline. This probability is higher than a reference distribution, typically the Gaussian distribution, hence the “fat” tails. A Gaussian distribution arises from the aggregation of many independent units, in our case organizations. Heavy-tailed distributions on the other hand hint at the existence of dependencies and interactions between organizations (Amaral et al. 1998) and can therefore be useful in the study of organizational dynamics.

The existence of fat tails in the size and growth-rate distributions has non-trivial implications. The size distribution reflects the known high asymmetry in size observed in the organizational population: there are very few large organizations coexisting with lots of smaller ones, but also with some of intermediate size. The heavy-tailed growth-rate distribution also has important consequences. Most of the time organizations experience small or zero change, meaning their size changes very little or nothing at all. However, because the tails of the distribution that measure extreme events of large growth—positive for expansions and negative for declines—are much more likely to occur than they would be in a “Gaussian” world, huge size variations in a relatively short period of time are to be expected. Moreover, such rare but likely events have been shown to take place in organizations of all sizes, so fluctuations do not have a typical scale.

The study of heavy-tailed statistics has a long tradition in economics, although it is still not part of mainstream economics. It dates back to the studies of income inequality by Pareto during the late 19th century, where he showed that the distribution of wealth is skewed so that roughly 20% of the people concentrate 80% of the wealth [see Mandelbrot (1983, Part 11) for a classic historical review]. Regarding the growth-rate distribution, an additional empirical observation is its ubiquity: the heavy-tailed pattern appears in many different kinds of growth processes. The settings are quite varied, ranging from natural ones like the growth of bird populations (Keitt et al. 2002), to man-made settings like financial systems (Plerou et al. 1999), GDP and the growth of imports and exports (Castaldi and Dosi 2009; Podobnik et al. 2008) and industrial growth (Amaral et al. 1997b; Ishikawa 2006a; Bottazzi and Secchi 2006a), just to name a few examples.

Other quantities have been studied, for example how the variance of the growth-rate distribution depends on organizational size, and the statistics of subunits within organizations. Another example is the autocorrelation between growth rates at different points in time (Bottazzi et al. 2001), although the empirical evidence in the latter case is not as conclusive as for the quantities mentioned previously.

1.1.2 Sector and industry growth

Empirical studies of organizational growth have traditionally focused on manufacturing activities in private firms, either for the whole population or for particular activities (Dosi 2007). Heavy-tailed growth patterns have been reported for example in pharmaceutical (Bottazzi et al. 2001; De Fabritiis et al. 2003; Fu et al. 2005; Matia et al. 2004), furniture, shoes and printing (Bottazzi and Secchi 2003a, 2006a), textiles (Bottazzi and Secchi 2003b), metals and chemicals (Bottazzi and Secchi 2003a) or food (Bottazzi and Secchi 2003a). Apart from these studies, little is known about the statistics of other types of organizations, for example publicly-driven ones. This lack of knowledge limits our understanding of the differences between different types of organizations, and helps instead to reproduce a simplified view of the problem based on few “universal” stylized facts. For instance, it might be the case that privately-owned organizations are driven by a market logic that is not shared by publicly-owned ones, and looking into the similarities and differences in growth statistics will provide information on their differential dynamics. The same goes for different industrial activities.

1.2 Models of organizational growth processes

The use of models to understand underlying processes governing the dynamics of organizational growth processes are almost as old as the empirical findings outlined above. We can distinguish between two lines of research: statistical inference models and generative models. The first line comprises, on the one hand, the statistical estimation of the effect of various relevant variables on organizational growth, and, on the other hand, various techniques for fitting statistical distributions to empirical data. An example of the former comes from economics with the study of the so-called *Gibrat's Law* (Gibrat 1931) sometimes called Law of Proportionate Effect. Under a series of assumptions—listed e.g. in Amaral et al. (1997b)—Gibrat's Law states that the growth rate of an organization is independent of its size (Dosi 2007).

The second line of research, generative models, has to do with developing models that can reproduce one or more of the observed statistical patterns. This line of research dates back at least to the work of Gibrat (1931) and his stochastic model of organizational growth, and includes also the seminal works by Simon (1955). A recent review of models can be found in Mondani et al. (2014), and Luttmmer (2010) provides a review of economic models. Recent generative models are quite dominated by physically inspired concepts and techniques, a field known as econophysics (Mantegna and Stanley 2000; Chakrabarti et al. 2006). Econophysical generative models like (Amaral et al. 1998) reproduce aggregated statistical patterns from stochastic processes at the organization or subunit level. However, econophysics generative models are typically not calibrated with empirical data. The interplay between size, growth-rate and movement distributions should be used to calibrate these models. Even more so since within the economic literature, econophysical generative models have been criticized for lacking proper statistical methodology (Gallegati et al. 2006). So there is a need for empirically-based quantitative descriptions to better inform model design.

1.3 Aim of this study

In the previous paragraphs we have outlined several research needs. On the empirical side, there is a need to complement the existing and abundant findings on statistical regularities in growth patterns with evidence from organizations in publicly driven activities. On the modeling side, model design can benefit from a more detailed and statistically rigorous understanding of how inter-organizational movements generate the observed statistical patterns. In the light of the needs outlined above, this study uses a Swedish longitudinal database on employment in the Stockholm Region to inform organizational growth models. It does so by analyzing the interplay between organizational growth statistics and inter-organizational movements for around 843,000 employees in the region during a period of 14 years. All organizations with registered employees, public and private, as well all types of industrial activities are considered. The remainder of this article is structured as follows. First, we present the data in Sect. 2. Our estimation methods are discussed in Sect. 3. The results are presented in Sect. 4. In a first part, we fit statistical distributions to size (number of employees) and growth-rate probability functions. In order to see how the aggregate patterns result from a superposition of group-level patterns, we categorize the organizations by ownership sector and industrial activity and compute their statistical distributions. In a second part, we look at movement probability distributions to determine how the growth statistical pattern comes about. Finally, conclusions and implications for further research are discussed in Sect. 5.

2 Data

We use the Stockholm database in our analyzes. It is a unique compilation of Swedish governmental registers, providing information on all organizations having people registered as employees in the Stockholm Region, for the period 1990–2003. It is possible to use a legal ownership variable to separate between organizations belonging to the public and private sectors. It is also possible to categorize the industrial activity of each workplace into industry groups, according to the SNI standard.¹ We also have information on the organizational membership for every individual 16 years old or older who is employed by an organization in the Stockholm Region during the period. This allows us to compute the incoming and outgoing employee movements for each organization over time. We measure the size $S_i(t)$ of organization i at year t as the total number of employees. Other measures of size are possible, like for instance company sales in the case of profit organizations. The statistical patterns for growth have been shown to be similar regardless of the measure (Amaral et al. 1997a).

In all our analyzes we restrict ourselves to organizations with size 10 or more employees. Such organizations concentrate the majority of employees in the data, an average of around 694,000 out of 843,000 employees. On the other hand, organizations with at least 10 employees represent slightly more than 10% of the total, 8700 out of 82,000 approximately. There are two reasons for doing this distinction: some related to data and some theoretical. On the data side, the organizational membership reported in the data is for

¹ SNI Näringsgrensindelning, standards SNI69, SNI92 and SNI2002. Statistics Sweden, see http://www.scb.se/sv/_Dokumentation/Publikationer-om-statistiska-metoder/Meddelande-i-samordningsfragor-MIS/SNI-Naringsgrensindelning/.

the largest income source of the individual. It is therefore possible that a person who is registered as employee in more than one organization appears to change membership just because the first income source changed. This effect should be larger for very small organizations. In fact, an average of 45% of the small organizations in our data have only 1 or 2 employees. Moreover, most studies of the kind we are carrying out use an even higher cut-off, so by setting it to 10 we are still able to capture some of the low-size statistical properties.

On the theoretical side, the dynamics of very small organizations differ from those of larger ones. organizations of very small sizes can be start-ups, small family businesses or even individual consulting companies. In our data, over 66% of the organizations with less than 10 employees belong to the activities ‘4—Construction’, ‘5—Commerce and communication’, and ‘6—Financial and company services’. These very small organizations are likely to be more unstable; their growth processes more drastic, and be impacted by processes like innovation in a different way compared with larger organizations (Santarelli 2006; Lotti et al. 2003). We can see this is reflected by the larger volatility of this group measured by the variation coefficient of the total number of employees—i.e. the standard deviation in units of the mean. This value is on average higher among organizations of less than 10 employees (0.12) than among larger ones (0.08). So is the variation coefficient of the total number of organizations, 0.18 for small organizations against 0.14 for larger ones.

3 Methods

In the first part of this study, we analyze aggregate organizational growth statistical patterns and break them down by ownership sector and industry. Regarding the size distribution, a way of presenting the data is through the complementary cumulative distribution function (CCDF) of organizational sizes. This distribution informs on the fraction of organizations in a population of organizations that have size (number of employees) larger or equal to a given size s . Traditionally, the baseline distribution has been the lognormal distribution, but as Axtell (2001) showed, the inclusion of small size ranges and organizations outside the stock exchange makes the *power law distribution* a better fit. The function takes two parameters: the exponent α and the minimum size $S_{\min}$. The exponent is sometimes called the Pareto index in this literature (Ishikawa 2006b), in reference to its original application to inequality by Pareto (Mandelbrot 1960). The minimum size at which the function offers a good fit and it is set to avoid divergence as the size tends to zero and gives an indication of the lower limit of the power-law range. The distribution has the form

$$P_{\text{PL}}(S \geq s) \sim \left(\frac{S_{\min}}{s} \right)^{\alpha}, \quad \alpha > 0. \quad (1)$$

We estimate the parameters $\alpha, S_{\min}$ with the python package `powerlaw` (Alstott et al. 2014). The estimation technique is maximum likelihood (MLE), and we consider the fact that size distributions are discrete, since in our case there can only be an integer number of employees in an organization a given year.

We furthermore test alternative probability density functions (PDF) for organizational size. A minimal first test is against the traditional *lognormal distribution* (LNORM):

$$P_{\text{LNORM}}(S, \mu, \Sigma) \sim \frac{1}{\Sigma \sqrt{2\pi S}} \exp \left(-(\ln S - \mu)^2 / (2\Sigma^2) \right). \quad (2)$$

This distribution can be obtained in the context of the basic Gibrat model, by a simple random walk of the logarithm of organizational size, and is thus a baseline case. It should be noted though that, as in any random walk, the variance of organizational size grows in time in the Gibrat model, so the resulting lognormal distribution is not stable in time. There are ways to make the distribution stable, e.g. by introduction of new firms.

But the population of organizations has limitations in terms of the size of the organizations it can accommodate. So secondly, we consider the *truncated power law distribution* which is a power law with an exponential upper cut-off to model constraints on large sizes by making the less likely [see Newman (2003) for a review in a network context]:

$$P_{\text{TPL}}(S, \alpha, \lambda) \sim S^{-\alpha} \exp(-\lambda S). \quad (3)$$

The test procedure is a likelihood ratio test,² where we report the log-likelihood difference and the p value measuring how significantly the difference deviates from zero. Note that the role of the p value in this context is to provide a criterion for comparison of goodness-of-fit that allows seeing whether there is a significant improvement in fit to the data when choosing one distribution over the other. It should not be interpreted as a significance test like one does in a statistical inference context.

For the growth-rate distribution, if an organization has S_1 employees the current year and S_0 employees the year before, we define the *growth rate* as

$$r_1 = \log_{10} \frac{S_1}{S_0}. \quad (4)$$

The notation follows the usual convention, e.g. (Amaral et al. 1997b). Note that S_0 is not the initial size in the time series. The rate could also be measured in other logarithmic scales, in absolute form, or as a return on sales. The baseline PDF in the case of growth rates is the *Laplace distribution*, also called double-exponential distribution (Amaral et al. 1997b). It is a conditional probability of a growth rate r_1 on one year, given the size S_0 of the organization the year before. With $\bar{r}_1, \sigma_1$ the mean and standard deviation of growth rates, the equation for the PDF reads:

$$P_{\text{LAP}}(r_1 | S_0) = \frac{1}{\sqrt{2}\sigma_1(S_0)} \exp\left(-\frac{\sqrt{2}|r_1 - \bar{r}_1(S_0)|}{\sigma_1(S_0)}\right). \quad (5)$$

In order to give an estimation of deviation from the Laplace, we choose to use the *Subbotin* (or exponential power) function for our estimation of the growth-rate distributions. The Subbotin is a family of distributions that has the Laplace and the Gaussian as particular cases. It consists of a shape parameter b , a scale parameter a , and a position parameter m as follows:

$$P_{\text{SUB}}(r, a, b, m) = \frac{1}{2ab^{1/b}\Gamma(1 + 1/b)} \exp\left(-\frac{1}{b}\left|\frac{r - m}{a}\right|^b\right), \quad (6)$$

where Γ is the gamma function. The Laplace PDF in Eq. (5) can be recovered by setting $b = 1, \sigma_1 = \sqrt{2}a, \bar{r}_1 = m$. As b tends to 2, the Subbotin tends to a Gaussian. Shape parameters lower than 1 imply fatter-than-Laplace tails and thus more likely large growth events.

² Alternative comparison methods such as Bayesian Information Criterion can in principle be used (Clauset et al. 2009).

Since we want to check for differences in positive and negative growth events, we use the asymmetric version of the Subbotin distribution (Bottazzi and Secchi 2011). Here, there exist one shape and one scale parameter for the left and right tails b_l, b_r, a_l, a_r and a position parameter m as above. The MLE estimation of the Subbotin parameters is done with the SUBBOTOLS package (Bottazzi 2004); see Fagiolo et al. (2007) for an application of this package.

4 Results

4.1 Database description

The total number of employees by sector and industry, considering only organizations of 10 or more employees, is shown in Fig. 1. It can be observed that, in terms of total number of employees, the public and private sectors start the 1990's relatively even (the public sector employs a considerable amount of people in the Stockholm Region). As time progresses the private sector grows relatively larger. The public sector is however more stable, with a population of around 290,000 employees (and 520 companies, not shown in the figure) on average, while the population in the private sector varies following the macroeconomic context of the time (an average of 402,000 employees in 8200 organizations).

Looking at industrial activities, we can see that some industries like '1—Agriculture, Forestry and Fishing' and '3—Energy and Waste' are marginal, which is not surprising in a region like Stockholm. Overall, no single industry dominates in number of employees. The larger industries are '5—Commerce and Communication' (155,000 on average), '6—Financial and Commercial Services' (average 131,000) and 'Human Health and social work' (114,000 on average). This is to be expected, since Stockholm as a region concentrates a high percentage of the Swedish population, and thus welfare state services, but also central national services and administrative and financial headquarters.

Table 1 shows the aggregate composition in terms of number of employees by sector and industry as a fraction of the cumulative total number of employees over the whole time period. The public sector concentrates most of the health and education activities, and administration activities have of course a predominant role. The private sector accounts for most of the manufacturing, commerce, communication and financial activities.

4.2 Interplay between sector and industry statistics

4.2.1 Aggregate distributions

We begin by analyzing the size distribution of organizations with 10 or more employees for the aggregate (i.e. total) population of organizations and the whole time period 1990–2003. We plot the size CCDF in log–log scale in the left panel of Fig. 2. The plot also shows an MLE fit to a power-law distribution according to Eq. (1). The corresponding exponent estimate $\hat{\alpha}$ is close to 1, as has been reported for other populations which include small companies (Axtell 2001). We observe that the fit does not hold for the entire size range: the minimum size by MLE fit is 21 employees, and the fit is good for between 2 and 3 orders of magnitude, overestimating the real distribution in the tail, meaning that there are slightly less very large organizations than expected in a power law model.

Fig. 1 Time evolution of employment in Stockholm Region organizations, by sector (left) and industry (right). Considering only organizations with 10 or more employees. Key for industry names: '1—Agriculture, forestry and fishing', '2—Manufacturing and mining', '3—Energy and waste', '4—Construction', '5—Commerce and communication', '6—Financial and company services', '7—Education and research', '8—Human health and social work', '9—Cultural services', '10—Public administration and others'

Table 1 Number of employees per sector and industry, as a percentage of the aggregated total

Industry	Public (%)	Private (%)	Total (%)
0—N/A	7.10	92.90	0.08
1—Agriculture, forestry and fishing	37.30	62.70	0.20
2—Manufacturing and mining	6.65	93.35	12.47
3—Energy and waste	72.96	27.04	0.88
4—Construction	15.31	84.69	3.67
5—Commerce and communication	25.19	74.81	22.40
6—Financial and company services	16.39	83.61	18.81
7—Education and research	83.18	16.82	8.94
8—Human health and social work	81.46	18.54	16.40
9—Cultural services	26.42	73.58	8.22
10—Public administration and others	98.69	1.31	7.93
Total	41.64	58.36	100.00

Considering only organizations with 10 or more employees

Fig. 2 Aggregate distributions for the whole population. Whole time period, organizations with 10 or more employees. Size distribution (left) and growth-rate conditional probability density (right). Power-law MLE fit for reference in the size distribution—Eq. (1)—and Laplace MLE fit for reference in the growth-rate density—Eq. (5)

Full estimations for the power-law parameters for this and the upcoming cases—by sector and by industry—are reported in Table 2. We performed likelihood-ratio tests with our two alternative distributions according to Eqs. (2–3). The lognormal is a worse fit to the data, but the truncated power law offers a statistically significant alternative. The power law fit with an exponent close to 1 has a straightforward interpretation, apparent also from the plot: multiplying the size by a factor of ten decreases the fraction of organizations having that size or more by approximately the same factor. For example, around one in ten organizations have 100 or more employees, while close to one in a hundred have 1000 or more. And this proportionality holds for almost three orders of magnitude, so it is invariant across size scales.

Table 2 Power-law estimations for size distributions, and comparison tests with lognormal (LNORM)—Eq. (2)—and truncated power law (TPL)—Eq. (3)

Distribution	Power-law estimates		Goodness-of-fit test	
	$\hat{\alpha}$ (SE)	$\hat{S}_{\min}$	logLR (p value)	
			LNORM	TPL
Aggregate	1.03 (0.00)	21	14.45 (0.03)	– 10.1 (0.00)
Sector				
Public	0.67 (0.01)	95	– 54.32 (0.00)	– 60.8 (0.00)
Private	1.19 (0.01)	37	– 20.48 (0.00)	– 26.9 (0.00)
Industry				
1—Agriculture, forestry and fishing	1.19 (0.05)	10	– 5.95 (0.02)	– 9.3 (0.00)
2—Manufacturing and mining	0.93 (0.01)	10	– 2.76 (0.28)	– 14.9 (0.00)
3—Energy and waste	1.01 (0.05)	34	– 3.38 (0.10)	– 4.6 (0.00)
4—Construction	1.31 (0.02)	18	– 1.30 (0.34)	– 4.1 (0.00)
5—Commerce and communication	1.16 (0.01)	35	0.63 (0.79)	– 1.1 (0.13)
6—Financial and company services	1.17 (0.02)	72	– 7.27 (0.02)	– 12.9 (0.00)
7—Education and research	0.71 (0.01)	11	– 11.73 (0.00)	– 22.6 (0.00)
8—Human health and social work	0.70 (0.01)	10	1.29 (0.35)	– 7.4 (0.00)
9—Cultural services	1.45 (0.03)	92	– 0.94 (0.36)	– 3.5 (0.01)
10—Public administration and others	1.10 (0.05)	309	– 3.56 (0.05)	68.4 (0.00)

Whole time period. Aggregate, by sector and by industry

The growth-rate conditional PDFs are shown in the right panel of Fig. 2 presents the PDF as defined in Eq. (4). The density functions are binned by initial size. For example, the plot with squared markers must be read as the probability to have a growth rate r_1 the present year, given that the organization had between 10 and 100 employees the year before. Only growth rates for sizes larger than or equal to 10 employees are considered. An MLE fit to a Laplace distribution according to Eq. (5) is also shown in the plot, as the Laplace is usually the baseline distribution. The plot is in semi-log scale, giving the double-exponential Laplace curve the so-called “tent” shape. We see that none of the distributions is Gaussian, and this is the reason for binning the data, since a density of normal distributions with different variance can give a Laplace (West 1987). We observe also that higher initial-size bins have curves with less variance, meaning that larger organizations have smaller growth-rate magnitudes. This observation is well established in the literature, e.g. (Amaral et al. 1997b). It provides a rationale for binning the data by initial size. If Gibrat’s Law held, then one would expect the same behavior at all size ranges. In our case, the change in variance with initial size is not very pronounced, but it is more evident for negative growth rates.

Full estimations for the Subbotin distribution in Eq. (6) are reported in Table 3.³ For the aggregate case discussed here, we see that the shape parameter estimates $\hat{b}_l, \hat{b}_r$ in small organizations (bin 1) are close to 1, meaning close to the Laplace distribution, with the

³ We estimated the original symmetric Subbotin as well. The estimations for the symmetric case are roughly equal to the average of the two estimations for the asymmetric case, except for industries ‘2—Manufacturing and mining’ and ‘4—Construction’.

Fig. 3 Size distributions by sector. Power-law MLE fit for reference. Whole time period

left tail a little fatter. The next two bins, namely, medium and large organizations, have $\hat{b}_l, \hat{b}_r \approx 0.6\text{--}0.8$. The MLE fit here is a Subbotin with tails fatter than the Laplace. The last bin has very few data points, and the error is too large to say anything meaningful. The estimate for the position parameter $\hat{m}$ is zero, so the distributions are essentially centered around zero growth, reflecting the fact that there is very little or no growth/decline at all most of the time. The large rates in the tails, although relevant for the dispersion, have very low probability and do not shift the center to much even for asymmetric densities. The scale parameter estimates $\hat{a}_l, \hat{a}_r$ are related to the variance, and decrease for increasing initial-size bin more clearly so in the right tail, except in the last case due to the few points that go into the estimation.

4.2.2 Distributions by sector

It is difficult to draw meaningful conclusions just by looking at aggregate patterns. For instance, a sum of Gaussian distributions of different variances can shield a power-law distribution in the aggregate (West 1987) and this is the reason we break the growth distributions by size bin. Therefore, we decompose the distributions by ownership sector and industrial activity. Adding the sector dimension to our analyses provides an interesting opportunity, since most studies of growth typically focus on private companies, while in reality the sectors interact with each other, especially in a region like Stockholm where the number of people employed by the public and private sector are comparable.

The size distribution broken down by public/private sector is shown in Fig. 3. We see that the public sector has a “fatter” distribution than the private sector, resembling the aggregate size distribution. The exponent is smaller, so the decay is not as fast as in the aggregate distribution, meaning that large organizations are more frequent in the Stockholm public sector than in the private sector. This is why in Table 2, both the lognormal and the truncated power law offer a statistically significant better fit compared to the power law. The private sector’s size distribution decays faster than the full one, with a distribution that can be better fitted by a truncated power law, although the lognormal cannot be

Table 3 Asymmetric Subbotin estimations for growth-rate distributions, developed from Eq. (6)

Distribution	Estimates (SE)				
	$\hat{b}_l$	$\hat{b}_r$	$\hat{a}_l$	$\hat{a}_r$	$\hat{m}$
Aggregate					
Bin 1: $10^1 \leq S_0 < 10^2$	0.88 (0.01)	0.98 (0.01)	0.05 (0.00)	0.06 (0.00)	0.00 (0.00)
Bin 2: $10^2 \leq S_0 < 10^3$	0.68 (0.01)	0.79 (0.02)	0.05 (0.00)	0.04 (0.00)	0.00 (0.00)
Bin 3: $10^3 \leq S_0 < 10^4$	0.59 (0.03)	0.77 (0.05)	0.03 (0.00)	0.02 (0.00)	0.00 (0.00)
Bin 4: $10^4 \leq S_0 < 10^5$	1.23 (0.47)	1.00 (0.53)	0.05 (0.01)	0.02 (0.01)	0.00 (0.01)
Sector					
Public					
Bin 1: $10^1 \leq S_0 < 10^2$	0.74 (0.03)	0.75 (0.03)	0.04 (0.00)	0.05 (0.00)	0.00 (0.00)
Bin 2: $10^2 \leq S_0 < 10^3$	0.63 (0.03)	0.78 (0.04)	0.04 (0.00)	0.03 (0.00)	0.00 (0.00)
Bin 3: $10^3 \leq S_0 < 10^4$	0.78 (0.07)	1.00 (0.11)	0.02 (0.00)	0.02 (0.00)	0.00 (0.00)
Bin 4: $10^4 \leq S_0 < 10^5$	1.23 (0.47)	1.00 (0.53)	0.05 (0.01)	0.02 (0.01)	0.00 (0.01)
Private					
Bin 1: $10^1 \leq S_0 < 10^2$	0.90 (0.01)	1.00 (0.01)	0.05 (0.00)	0.06 (0.00)	0.00 (0.00)
Bin 2: $10^2 \leq S_0 < 10^3$	0.67 (0.02)	0.87 (0.03)	0.05 (0.00)	0.04 (0.00)	0.00 (0.00)
Bin 3: $10^3 \leq S_0 < 10^4$	0.47 (0.04)	1.20 (0.13)	0.05 (0.00)	0.04 (0.00)	-0.02 (0.00)
Industry (1 bin)					
1—Agriculture, forestry and fishing	0.89 (0.11)	0.93 (0.11)	0.07 (0.01)	0.07 (0.01)	0.00 (0.01)
2—Manufacturing and mining	0.20 (0.00)	0.18 (0.00)	0.02 (0.00)	0.03 (0.00)	0.00 (0.00)
3—Energy and waste	0.57 (0.04)	0.70 (0.05)	0.04 (0.00)	0.04 (0.00)	0.00 (0.00)
4—Construction	0.02 (0.00)	0.02 (0.00)	0.03 (0.00)	0.03 (0.00)	0.00 (0.00)
5—Commerce and communication	0.64 (0.01)	0.70 (0.01)	0.05 (0.00)	0.05 (0.00)	0.00 (0.00)
6—Financial and company services	0.86 (0.01)	1.01 (0.02)	0.06 (0.00)	0.07 (0.00)	0.00 (0.00)
7—Education and research	0.63 (0.02)	0.80 (0.03)	0.04 (0.00)	0.05 (0.00)	0.00 (0.00)
8—Human health and social work	0.71 (0.03)	0.90 (0.03)	0.04 (0.00)	0.05 (0.00)	0.00 (0.00)
9—Cultural services	0.88 (0.02)	0.93 (0.02)	0.05 (0.00)	0.05 (0.00)	0.00 (0.00)
10—Public administration and others	0.50 (0.02)	0.52 (0.02)	0.03 (0.00)	0.04 (0.00)	0.00 (0.00)

Whole time period. Aggregate, by sector and by industry

ruled out. Consequently, the power-law distribution we observe in the aggregated size distribution results from the superposition of the power laws of the sector distributions at low sizes, while the range of large organizations is dominated by the functional form of the public sector distribution.

The corresponding growth-rate plots are shown in Fig. 4. The tails in the private sector are similar to the full distribution, with less variation across the initial-size curves. The parameter estimations are quantitatively similar as well (see Table 3), for initial-size bins 1 and 2. The body of the public sector growth-rate distribution looks qualitatively like the aggregated one and the variance decay is similar in quantitative terms. The tails are generally more symmetric in the case of the public sector.

The aggregate size distribution and the body of the growth-rate distribution are dominated by the form of the public sector distributions, while the functional form of the tails in the growth-rate distribution is dominated by the private sector. We can relate this to the time evolution of sizes by sector (cfr. Fig. 1) in which we saw that the size of the public

Fig. 4 Conditional growth-rate probability density by public (left) and private sector (right). Laplace MLE fit for reference. Whole time period

sector fluctuates on average less than the private. Had we looked only at the aggregated time evolution, we could not have seen that it results from a superposition of similarly-size growth and declines, as indicated by the symmetry of the tails.

4.2.3 Distributions by industry

Another dimension is industry, or more precisely, a categorization of workplace activities in terms of its gross industrial sector. The analysis could be performed with higher level of activity disaggregation, but over such a long period of time, categories are created and eliminated, and the bureau of statistics changes its classification standards, so in practice the gross industrial sector is the only viable alternative for longitudinal analysis. Only divisions larger or equal than 10 employees are considered.⁴

We plot the size distributions by industry in Fig. 5. We note that the exponent estimates are above 1 for the most part, and close to the public sector estimate for ‘7—Education and research’ and ‘8—Human health and social work’. The public sector is dominated by these two activities. Public administration is also a dominant activity; in this case though the estimate for the minimum size is one order of magnitude higher than the rest, so the fit does not hold to the same extent in this sector.

Regarding alternative functional forms for the CCDF, the lognormal offers on a 5% level a better fit for ‘1—Agriculture, forestry and fishing’, ‘6—Financial and company services’, ‘7—Education and research’ and ‘10—Public administration and others’. This illustrates how sometimes it is not necessary to complicate the estimation beyond the simple lognormal, a point made by Alstott et al. (2014). The truncated power law is in general a statistically significant better fit, except for ‘10—Public administration and others’ where the power law is preferred, and ‘5—Commerce and communication’ where the test

⁴ This cut does not affect the size estimations, but does influence the tails of the growth-rate distributions by adding large fluctuations and giving “fatter” tails i.e. lower b estimations. Estimations for the other parameters are not substantially affected.

Fig. 5 Size distributions by industry. Whole time period. Divisions with 10 or more employees. Power-law MLE fit for reference

log-likelihood difference is not significant on a 5% level. This suggests that, broken down by industry, most of the distributions have an upper exponential cut-off in size, although at different size ranges indicating that different industrial activities have different characteristic sizes.

Fig. 6 Conditional growth-rate probability density by industry. Whole time period. Divisions with 10 or more employees. Laplace MLE fit for reference

The growth-rate PDFs by industry are plotted in Fig. 6. Since some categories contain too few data points to allow for more than one initial-size bin, all data in each industry are pooled together in one single initial-size bin. Here the variation across industries becomes more apparent. All probability densities are centered around zero growth. Regarding the

Fig. 7 Summary plot of estimations for sector and industry. Estimation for power-law exponent in size distribution ($\hat{\alpha}$) in horizontal axis, and estimation for shape parameters of asymmetric Subbotin growth-rate distribution ($\hat{b}_l$ hollow squares, $\hat{b}_r$ solid circles) in vertical axis. Point estimates and error bars are shown for each case. Sector point estimates obtained as averages over initial-size bin estimations. Dotted lines for reference: typical exponent for power-law size distribution ($\alpha = 1$) and Laplace growth-rate distribution ($b = 1$)

shape parameter estimates, we observe that none is larger than 1 (including error bars). So the distributions are either close to the Laplace form—like for ‘1—Agriculture, forestry and fishing’, ‘6—Financial and company services’ and ‘9—Cultural services’—or of fatter tails, ‘4—Construction’ and ‘2—Manufacturing and mining’ among the lowest estimates. The normal distribution ($b = 2$) is definitely out of range for all of these distributions, so the extreme events observed in reality decay more slowly than a normal distribution, and are consequently more likely to take place. The scale estimates $\hat{\alpha}$ also vary, from some narrow densities as in the case of ‘4—Construction’ and ‘10—Public administration and others’, to high values as in ‘1—Agriculture, forestry and fishing’ and ‘6—Financial and company services’. The degree of symmetry is also interesting to observe. Most industries exhibit quite symmetrical growth-rate PDFs, although in general the left tail tends to be fatter than the right one, meaning an overall excess in decline events. Some industries like ‘2—Manufacturing and mining’ lean towards the positive side.

4.2.4 Synthesizing sector and industry statistics

Size and growth are two intertwined aspects of growth dynamics. Therefore, it is natural to look for patterns in the parameters of the respective distributions. We plot the two main estimations, namely the size exponent $\hat{\alpha}$ and the Subbotin shape parameter $\hat{b}$ in Fig. 7.

Comparing industries, we can see that ‘1—Agriculture, forestry and fishing’ falls close to ‘6—Financial and company services’, despite the fact that they are two very distinct activities of different importance in terms of employed population in the Stockholm Region. Also, there is a cluster formed by ‘3—Energy and waste’, ‘5—Commerce and communication’ and ‘10—Public administration and others’, again three different activities. Industries ‘2—Manufacturing and mining’, ‘4—Construction’ and ‘9—Cultural services’ stand isolated.

Looking at the decomposition of sectors by industry, we observe that the point estimation for the public sector falls close to industries ‘7—Education and research’ and ‘8—Human health and social work’, two large public-sector activities. Interestingly, ‘10—Public administration and others’ does not fall close to public sector, so the overall dynamics is not as influenced by this industry as it is by the previous two. The private sector estimate falls close to the first cluster of ‘1—Agriculture, forestry and fishing’ falls close to ‘6—Financial and company services’, plus ‘5—Commerce and communication’. The industry ‘2—Manufacturing and mining’, despite employing more than 10% of the aggregated employee population, is far from the two sectors. It is as well interesting that the tails appear less fat in the aggregated picture, when they are rather a superposition of tails of different magnitudes when broken down by industrial activity.

4.3 Movement distributions

Our data provide a unique opportunity to analyze the dynamics of organizational size statistics in terms of inter-organizational movements. This is more informative than just looking at the time evolution of aggregate quantities and complements our previous observations. It also sheds light into one important, yet not often discussed limitation in the analysis of growth-rate statistics: by binning the data by initial size, all data points are pooled into one distribution. That is to say, these are usually aggregate plots that look at the whole time period for a given organizational population, but do not inform on the time point the rates come from. Neither does it inform on the rate composition. If we denote the absolute employee movements during a given year by $M_{1,\text{in/out}}$, the size that given year is by definition of movements equal to:

$$S_1 = S_0 + M_{1,\text{in}} - M_{1,\text{out}}. \quad (7)$$

So looking at the growth rate alone gives information on the net movement, but not on how this net movement came about. The reason for these analyzes not being abundant is probably the lack of data that allows for tracking of people flows in time. We plot in Fig. 8 the PDF of absolute incoming/outgoing employee movements, using the same plotting logic as in the growth-rate plots. We observe that the densities shift to higher movement ranges as the initial-size bin increases; this is true for incoming as well as for outgoing. Small organizations are qualitatively different from the rest. Movements decay faster in the outgoing plot; this fact most likely reflects our cut-off at sizes smaller than 10 employees. The other distributions follow a lognormal pattern, that is to say the logarithm of the movements follows a normal distribution, as shown by the MLE fit in the plot. The few data points in the uppermost initial-size bin do not follow a clear pattern.

We also note that only small and medium organizations have zero incoming movements and only the small ones has zero outgoing movements. Furthermore, the fraction of zero movements is very low in all case, as shown in the annotation in the figure: 7% and less than 1% in small respectively medium for incoming and almost 7% in small for outgoing.

Fig. 8 Conditional movement probability densities. Incoming (left) and outgoing (right). Whole time period. The plots follow the same logic as the growth-rate distribution plots. Zero movements shifted to 0.1 so they show up in the plot. Lognormal MLE fit for comparison (considering only non-zero movements). Fraction of zero movements by initial-size bin annotated in lower-right corner

This means that the large amount of zero growth rates (the center of the Laplace distribution) do not result from the organization not having any movements, but come instead from a cancellation of similar employees incoming and outgoing movements.

5 Discussion

In this study, we have used unique Swedish register data to inform organizational growth models by studying the interplay between growth statistics by sector and industry, and inter-organizational movements in the Stockholm Region for a period of 14 years. Heavy-tailed patterns exist for both size and growth-rate distributions, at the aggregate level and at the level of sector/industrial activity. This is an indication of a system where organizations interact and their sizes depend on each other's size.

The implications of this fact are quite relevant. From the point of view of size, a heavy-tailed organizational world differs from a “Gaussian” world in the fact that medium and large sizes occur more often, hence the “fatter” tail. Moreover, the heavy tail holds for several orders of magnitude, so that very large, but also middle-size organizations are allowed. In a normally distributed world only a handful of large companies would coexist with an overwhelming majority of average-size ones. The real world thus allows for much higher resource concentration in fewer organizations. From the growth rate point of view, non-Gaussian rate distributions allow for growth of large magnitude to take place relatively more frequently. This suggests that the power law is a better fit for the first three orders of magnitude. The exponential cut-off at higher orders might reflect global constraints on the possibility to have very large organizations due to overall population constraints of the Stockholm Region or of characteristic sizes in each industrial activity.

Aggregate statistical properties can be decomposed by sector. The public sector dominates the size distribution for large sizes and the growth statistics at small growth rates, while the private sector contributes to small-size statistics and dominates in the tails of the growth-rate distribution. A deeper discussion on the difference growth statistics for public

and private sectors is important to complement the quite private-centered view of empirical reality that currently exists in the literature. For instance, it is well established nowadays that the power law distribution is a good approximation to the size distribution; see e.g. Delli Gatti et al. (2008). Here, we see that a classical lognormal distribution might offer an alternative fit in the cases of sector distributions, ‘1—Agriculture, forestry and fishing’, ‘6—Financial and company services’, ‘7—Education and research’ and (marginally) ‘10—Public administration and others’.

Having observed these differences in the growth statistical patterns of the public and private sectors, one may speculate on the reason for their existence. One possibility is that the mechanisms at play are different for the public and private sectors. The public sector statistics in Stockholm, as seen in the results, are dominated by the health and education industries, which are composed of organizations and maintained to provide welfare and education services, not to seek profit in the market as the private sector. This could explain why organizations in the public sector are larger and have less extreme events than the private ones. It would be interesting to explore this in other datasets. In this discussion about generative mechanisms, it is important not to fall into the reductionism of trying to impute the observed functional form of the distributions to one single stochastic mechanism.

We further broke down the distributions by industry, and show how the estimation for different industries cluster in the parameter space of the estimations, and how the sector estimations are closely related to some of them, while other differ considerably. This clustering does not follow directly from the industry’s importance in terms of size. One thing we may conclude from size statistics is that different categories (sectors, industries) show statistics with varying functional forms. It is important to note how this analysis differs from traditional regressions to study growth processes. In that context, one would control for sector/industry and in practice that would imply adding them as a covariate. From the functional point of view, this would mean shifting the same regression function by an amount determined by the respective estimated coefficient. This logic differs fundamentally from what we do here. By separating each group we are able to capture differences in the distributions.

The total aggregated statistical pattern is better thought as a superposition of statistical distributions both across sector/industry. This kind of superposition of different distributions pointing to differentials in underlying dynamics can be related to recent findings in the econophysics literature on income distribution. Yakovenko and Barkley Rosser (2009) showed that income-based economic transactions lead to an exponential distribution in wealth, while investment-based transaction leads to a power-law distribution. The author theorizes that small incomes are dominated by an additive process leading to an exponential distribution, while large income is dominated by a multiplicative one leading to a power-law distribution.

A limitation of our research approach is that it focuses on growth rather than on the start and end points of the process, namely, the birth and death of organizations. The growth rate in Eq. (4) does not capture cases where the size is zero, thus excluding the mentioned start and end points. Our data does furthermore not contain information on the causes of organizational death, e.g. bankruptcy, fusion, etc. If such information was available, many interesting questions could be studied. For instance, one could consider fused and divided organizations as one overarching entity, like the concept of a “superfirm” in Bottazzi et al. (2001, 2002); Bottazzi and Secchi (2006b) and see which impact this has on growth statistics. Classical economic studies of firms look at such events, e.g. (Mansfield 1962). Firm birth and death are also an essential component within the organizational ecological approach (Hannan and Freeman 1989) and more recently (Hannan et al. 2007).

Another limitation is that our study looks at employment relations, and is therefore restricted to the labor market sphere of the organizational world. That is not to say that, e.g. voluntary associations are totally excluded, but rather that only the people that are formally employed by them appear in the data. This should be taken into account when interpreting the results, because interactions with the rest of the organizational population might be relevant for particular dynamics. Additionally, the movement statistics of Stockholm Region reflect at least partially the dynamics of the local economy. Thus, the changes we observe at the region level do not necessarily imply a change in sector/industry for the whole country.

Finally, another way to look at movements in this problem is from a network perspective. It is possible to construct a network where the nodes are the organizations, and a directed link exists at a given year if one or more employees have moved from one organization to another one that year. This is a very promising approach, which can give new insights on organizational dynamics in terms of the structure created inter-organizational movements. We will develop this approach in another article.

Acknowledgements The author would like to thank Peter Hedström and Yvonne Åberg for assistance with the database, and Fredrik Liljeros and Sara Roman for useful comments on earlier versions of this work.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Alstott, J., Bullmore, E., Plenz, D.: Powerlaw: a python package for analysis of heavy-tailed distributions. *PLoS ONE* **9**(1), e85777 (2014)
- Amaral, L.A.N., Buldyrev, S.V., Havlin, S., Leschhorn, H., Maass, P., et al.: Scaling behavior in economics I: empirical results for company growth. *J. Phys. I (France)* **7**, 621–633 (1997a)
- Amaral, L.A.N., Buldyrev, S.V., Havlin, S., Maass, P., Salinger, M.A., et al.: Scaling behaviour in economics: the problem of quantifying company growth. *Phys. A* **244**(1–4), 1–24 (1997b)
- Amaral, L.A.N., Buldyrev, S.V., Havlin, S., Salinger, M.A., Stanley, H.E.: Power law scaling for a system of interacting units with complex internal structure. *Phys. Rev. Lett.* **80**(7), 1385–1388 (1998)
- Axtell, R.L.: Zipf distribution of U.S. firm sizes. *Sci. New Ser.* **293**(5536), 1818–1820 (2001)
- Bottazzi, G.: Subbotools user's manual. In: Working Paper 2004/14. Laboratory of Economics and Management, Pisa (2004)
- Bottazzi, G., Secchi, A.: Common properties and sectoral specificities in the dynamics of U.S. manufacturing companies. *Rev. Ind. Organ.* **23**, 217–232 (2003a)
- Bottazzi, G., Secchi, A.: A stochastic model of firm growth. *Phys. A* **324**(1–2), 213–219 (2003b)
- Bottazzi, G., Secchi, A.: Explaining the distribution of firm growth rates. *RAND J. Econ.* **37**(2), 235–256 (2006a)
- Bottazzi, G., Secchi, A.: Gibrat's law and diversification. *Ind. Corp. Change* **15**(5), 847–875 (2006b)
- Bottazzi, G., Secchi, A.: A new class of asymmetric exponential power densities with applications to economics and finance. *Ind. Corp. Change* **20**(4), 991–1030 (2011)
- Bottazzi, G., Dosi, G., Lippi, M., Pammolli, F., Riccaboni, M.: Innovation and corporate growth in the evolution of the drug industry. *Int. J. Ind. Organ.* **19**(7), 1161–1187 (2001)
- Bottazzi, G., Cefis, E., Dosi, G.: Corporate growth and industrial structures: some evidence from the Italian manufacturing industry. *Ind. Corp. Change* **11**(4), 705–723 (2002)
- Castaldi, C., Dosi, G.: The patterns of output growth of firms and countries: scale invariances and scale specificities. *Empir. Econ.* **37**, 475–495 (2009)
- Chakrabarti, B.K., Chakraborti, A., Chatterjee, A. (eds.): *Econophysics and Sociophysics. Trends and Perspectives*. Wiley, Berlin (2006)

- Clauset, A., Shalizi, C.R., Newman, M.E.J.: Power-law distributions in empirical data. *SIAM Rev.* **51**(4), 661–703 (2009)
- De Fabritiis, G., Pammolli, F., Riccaboni, M.: On size and growth of business firms. *Phys. A* **324**(1–2), 38–44 (2003)
- Dosi, G.: Statistical regularities in the evolution of industries: a guide through some evidence and challenges for the theory. In: Malerba, F., Brusoni, S. (eds.) *Perspectives on Innovation*. Cambridge Scholars, Cambridge (2007)
- Fagiolo, G., Napoletano, M., Roventini, A.: How do output growth-rate distributions look like? Some cross-country, time-series evidence. *Eur. Phys. J. B* **57**(2), 205–211 (2007)
- Fu, D., Pammolli, F., Buldyrev, S.V., Riccaboni, M., Matia, K., et al.: The growth of business firms: theoretical framework and empirical evidence. *Proc. Natl. Acad. Sci. USA* **102**(52), 18,801–18,806 (2005)
- Gallegati, M., Keen, S., Lux, T., Ormerod, P.: Worrying trends in econophysics. *Phys. A* **370**(1), 1–6 (2006)
- Gatti, D.D., Di Guilmi, C., Gallegati, M., Gaffeo, E., Giulioni, G., Palestrini, A.: Scaling laws in the macro-economy. *Adv. Complex Syst.* **11**(1), 131–138 (2008)
- Gibrat, R.: *Les Inégalités Économiques*. Sirey, Paris (1931)
- Hannan, M.T., Freeman, J.H.: *Organizational Ecology*. Harvard University Press, Cambridge (1989)
- Hannan, M.T., Pólos, L., Carroll, G.R.: *Logics of Organization Theory: Audiences, Codes, and Ecologies*. Princeton University Press, New York (2007)
- Ishikawa, A.: Derivation of the distribution from extended Gibrat's law. *Phys. A* **367**, 425–434 (2006a)
- Ishikawa, A.: Pareto index induced from the scale of companies. *Phys. A* **363**(2), 367–376 (2006b)
- Keitt, T.H., Amaral, L.A.N., Buldyrev, S.V., Stanley, H.E.: Scaling in the growth of geographically subdivided populations: invariant patterns from a continent-wide biological survey. *Philos. Trans. R. Soc. B* **357**, 627–633 (2002)
- Lotti, F., Santarelli, E., Vivarelli, M.: Does Gibrat's law hold among young, small firms? *J. Evol. Econ.* **13**(3), 213–235 (2003)
- Luttmer, E.G.J.: Models of growth and firm heterogeneity. *Annu. Rev. Econ.* **2**, 547–576 (2010)
- Mandelbrot, B.B.: The Pareto–Lévy law and the distribution of income. *Int. Econ. Rev.* **1**(2), 79–106 (1960)
- Mandelbrot, B.B.: *The Fractal Geometry of Nature*. W. H. Freeman and Company, New York (1983)
- Mansfield, E.: Entry, Gibrat's law, innovation, and the growth of firms. *Am. Econ. Rev.* **52**(5), 1023–1051 (1962)
- Mantegna, R.N., Stanley, H.E.: *An Introduction to Econophysics: Correlations and Complexity in Finance*. Cambridge University Press, Cambridge (2000)
- Matia, K., Fu, D., Buldyrev, S.V., Pammolli, F., Riccaboni, M., Stanley, H.E.: Statistical properties of business firms structure and growth. *Europhys. Lett.* **67**(3), 498–503 (2004)
- Mondani, H., Liljeros, F., Holme, P.: Fat-tailed fluctuations in the size of organizations: the role of social influence. *PLoS ONE* **9**(7), e100527 (2014)
- Newman, M.E.J.: The structure and function of complex networks. *SIAM Rev.* **45**(2), 167–256 (2003)
- Plerou, V., Gopikrishnan, P., Amaral, L.A.N., Meyer, M., Stanley, H.E.: Scaling of the distribution of price fluctuations of individual companies. *Phys. Rev. E* **60**(6), 6519–6529 (1999)
- Podobnik, B., Horvatic, D., Pammolli, F., Wang, F., Stanley, H.E., Grosse, I.: Size-dependent standard deviation for growth rates: empirical results and theoretical modeling. *Phys. Rev. E* **77**(5), 056,102 (2008)
- Santarelli, E. (ed.): *Entrepreneurship, Growth, and Innovation. The Dynamics of Firms and Industries*. Springer, New York (2006)
- Simon, H.A.: On a class of skew distribution functions. *Biometrika* **42**(3–4), 425–440 (1955)
- West, M.: On scale mixtures of normal distributions. *Biometrika* **74**(3), 646–648 (1987)
- Yakovenko, V.M., Barkley Rosser, J.: Colloquium: statistical mechanics of money, wealth, and income. *Rev. Mod. Phys.* **81**(4), 1703–1725 (2009)