



Munich Personal RePEc Archive

Measuring productivity dispersion: a parametric approach using the Lévy alpha-stable distribution

Yang, Jangho and Heinrich, Torsten and Winkler, Julian and Lafond, François and Koutroumpis, Pantelis and Farmer, J. Doyne

Oxford Martin Programme on Technological and Economic Change, University of Oxford, Institute of New Economic Thinking at the Oxford Martin School, University of Oxford, Department for Business Studies and Economics, University of Bremen, Mathematical Institute, University of Oxford, Smith School for Enterprise and the Environment, University of Oxford, Santa Fe Institute

30 September 2019

Online at <https://mpra.ub.uni-muenchen.de/96474/>

MPRA Paper No. 96474, posted 18 Oct 2019 07:19 UTC

Measuring productivity dispersion: a parametric approach using the Lévy alpha-stable distribution

Jangho Yang^{1,2}, Torsten Heinrich^{2,3,4}, Julian Winkler^{1,2}, François Lafond^{1,2,5}, Pantelis Koutroumpis^{1,2}, and J. Doyne Farmer^{1,2,4,6}

¹*Oxford Martin Programme on Technological and Economic Change, University of Oxford*

²*Institute of New Economic Thinking at the Oxford Martin School, University of Oxford*

³*Department for Business Studies and Economics, University of Bremen*

⁴*Mathematical Institute, University of Oxford*

⁵*Smith School for Enterprise and the Environment, University of Oxford*

⁶*Santa Fe Institute*

September 30, 2019

Abstract

Productivity levels and growth are extremely heterogeneous among firms. A vast literature has developed to explain the origins of productivity shocks, their dispersion, evolution and their relationship to the business cycle. We examine in detail the distribution of labor productivity levels and growth, and observe that they exhibit heavy tails. We propose to model these distributions using the four parameter Lévy stable distribution, a natural candidate deriving from the generalised Central Limit Theorem. We show that it is a better fit than several standard alternatives, and is remarkably consistent over time, countries and sectors. In all samples considered, the tail parameter is such that the theoretical variance of the distribution is infinite, so that the sample standard deviation increases with sample size. We find a consistent positive skewness, a markedly different behaviour between the left and right tails, and a positive relationship between productivity and size. The distributional approach allows us to test different measures of dispersion and find that productivity dispersion has slightly decreased over the past decade.

Keywords: productivity, dispersion, distribution, heavy-tail, Lévy stable distribution

JEL codes: D2, O3, J24, R12

1 Introduction

Productivity is highly heterogeneous across countries and firms. At the national and regional level it motivates the work of growth and development economics representing a natural predictor of income differences. As such it has been linked to physical and human capital (Mankiw et al. 1992), technology (Hall & Jones 1999) and institutions (Acemoglu & Dell 2010). Within markets and firms, the literature has attributed productivity differences to competition, innovation and the quality of regulatory oversight (Aghion et al. 2005). More recently, productivity measurements have been at the forefront of economic and policy attention due to the increasing role of service industries, intangible capital and the limitations of national accounting (Goldin et al. 2018, Haskel & Westlake 2018).

But how do we *measure* productivity and its dispersion? Our aim in this paper is to find a good parametric model for the distribution of productivity levels and growth. In theory, firm labor productivity is the ratio of value added divided by the labour inputs used. As such, even at this granular level, productivity represents an aggregation of several micro-level processes such as tasks, sales, employees, and contracts. The Central Limit Theorem (CLT) states that, when independent random variables with *finite variance* are aggregated, the distribution of their normalized sum is the normal distribution. The *Generalized* CLT relaxes the assumption of finite variance, and in this case the only possible non-trivial limit of the normalized sum is the Lévy alpha-stable distribution (Nolan 1998). As for the CLT, the assumption that the micro variables are i.i.d can be partly relaxed, and while strong correlations imply a breakdown of the theorem, it remains widely applicable to real-world situations. Therefore, the GCLT makes the Lévy alpha-stable distribution a strong candidate for the resulting distribution.

To test this hypothesis, we employ a comprehensive dataset of around 9 million European firms for the period 2006-2015 using the Lévy alpha-stable distribution. The Lévy alpha-stable distribution has four parameters (location, scale, asymmetry, and tail), so it is flexible enough to capture the complex pattern of the productivity distributions, yet summarizes the properties of the data using four parameters only. In addition, it can be estimated easily using only five empirical quantiles. We compare the fit of the Lévy alpha-stable distribution to major contenders, and in particular the 4-parameter Subbotin distribution (also called asymmetric exponential power), a frequently used model for firm growth rates (Bottazzi et al. 2007, Bottazzi & Secchi 2011). Using a number of different fitting criteria, we find that the Lévy alpha-stable performs markedly better.

Across every dimension in our sample (country, industry, size, year) we find clear evidence that productivity distributions are heavy-tailed and skewed. This has important practical implications as it challenges the key measures of dispersion used in the literature so far. Essentially, heavy-tails imply that the theoretical variance is infinite, and thus the sample standard deviation increases with sample size - turning it to a poor and often misleading measure of dispersion. Figure 1 demonstrates this effect in the empirical data using 4 million observations on labor productivity levels in France, Italy, Germany, and Spain. Smaller sub-samples have a smaller standard deviation, on average, than larger sub-samples. The larger the sample, the more likely it is that the sample includes an extreme value, driving up the standard deviation. Figure 1 shows the mean sample standard deviation, but also the 5% and 95% quantiles (sampling is repeated 20,000 times for each subsample). The 5% and 95% quantiles are about an order of magnitude apart, which emphasizes that the standard deviation estimates are also very volatile between subsamples of the same size. A similar behavior applies also for the standard deviation of the productivity changes and of the productivity growth rates, providing further support for these findings.

If the standard deviation is a poor measure of dispersion, what should we use? There is no perfect answer to this, but the four-parameters of the Lévy alpha-stable provides a useful distinction between the




Figure 1: **Measured standard deviation in sub-samples of firm labour productivity.** We construct a distribution by pooling together the productivity levels of firms in France, Italy, Germany, and Spain (which have the same currency), for all years. Then, for each subsample size N , we compute the standard deviation of 20,000 subsamples, and report the average (black dots). In Appendix B, we explain that on i.i.d. data the theoretical scaling would be $N^{\frac{1}{\alpha}-\frac{1}{2}}$. The blue line shows that this scaling holds empirically here (we use our estimate of the tail parameter $\hat{\alpha} = 1.3$, and the intercept calculated by simulating i.i.d. random samples of size 20,000, based on all estimated parameters).

tail parameter, which measures the frequency of extreme observations, and the scale parameter, which measures the spread of the distribution. The asymmetry (i.e. skewness) of the distribution is measured by a separate parameter. To emphasize the relevance of this distinction, we provide simulations showing that the inter-quantile range (IQR), perhaps the most widely used dispersion statistic together with the standard deviation, mixes up skewness and dispersion: it can change when skewness changes, even if the scale and tail parameter do not change.

We analyze mostly productivity levels and productivity change, rather than growth. The reason for this is that, perhaps surprisingly, we find that the share of negative productivity levels is non negligible (5%). Keeping all values on productivity levels makes growth rates hard to interpret when they involve negative starting or ending values. Thus, in the paper we work mostly with productivity *changes*, i.e. the change in (say) euros of value-added per worker from one year to the next. We check that our results hold with growth rates when excluding negative productivity levels.

In an effort to alleviate the issues related to negative and extreme values, productivity levels are often log-transformed (Bartelsman & Wolf 2018). A first consequence of this is that the statistician is forced to exclude important information about firms that survive despite a negative value-added (which is not a result of measurement errors). The same applies for the log growth rates. A second, less obvious effect is that logarithms tend to highlight the differences between low productivity firms (left tail) and attenuate those among high productivity ones (right tail). Especially in the case of dispersion measurements the log transformation needs to be studied with caution. To highlight this, we compare labor productivity and log-labor productivity IQR for the five largest European countries and find contradicting evidence from

the two distributions. In fact the latter provides support for an increase in productivity dispersion over the past decade while the former shows the opposite. This result for the actual productivity distribution contradicts a long strand of the literature supporting that productivity dispersion has increased (Andrews et al. 2016, Haldane 2017, Cetto et al. 2018).

Further, we compare productivity levels and growth between countries, sectors, and firm sizes, highlighting the heterogeneity across these dimensions. Sectors with higher capital intensity, such as utilities, mining, and finance and real estate tend to be more dispersed and have a higher central value in both productivity levels and change. Different size classes exhibit noticeably different productivity distributions. Larger firms tend to have a more dispersed pattern, with a heavier tail and larger scale, and a higher central value in both productivity distributions.

Our results relate to four strands of literature. The first is the large literature on the distribution of firm sizes (Ijiri & Simon 1977, Axtell 2001) and their growth rates (Bottazzi et al. 2007, Bottazzi & Secchi 2011, Schwarzkopf et al. 2010, Holly et al. 2013), where the types of relevant distributions have been studied in detail, together with relationships between certain moments (typically, it is found that the volatility of growth rates scales with firm size (Calvino et al. 2018)). Most related to our work is the work on estimating productivity distributions such as (Aoyama et al. 2010), and most closely Gaffeo (2008) who estimates a Lévy alpha-stable distribution for the growth rates of sector-level total factor productivity in the US, and, like us, finds a tail exponent lower than two, suggesting an infinite second moment. This result persists when repeating the exercise on a larger sample of countries, with mixed results about asymmetry (Gaffeo 2011). In the present paper, the breadth of the Orbis Europe data improves the statistical power of these estimates, while shedding light on dispersion within industries.

Second, our results on productivity levels are relevant to the literature that links the misallocation of factors of production with productivity dispersion (Hsieh & Klenow 2009, Bartelsman et al. 2013, Foster et al. 2018), for several reasons. In this paper, we critically discuss measures of dispersion in detail, showing how some statistics may be misleading and what parts of the distribution can affect dispersion. Additionally, our work contributes to the current debate in this literature about the actual sources of dispersion, beyond misallocation, suggesting that models are necessary to interpret the data (Hsieh & Klenow 2017). We strongly suggest that theoretical models should aim to derive a Lévy alpha-stable distribution, where ideally the parameters would be interpreted in terms of misallocation or other sources¹. Finally, the misallocation literature is affected by serious data cleaning and measurement issues (Nishida et al. 2016). By suggesting a parametric framework for the underlying distribution, we also open the possibilities of detecting outliers more reliably (and in particular, avoiding the removal of extreme value when they are indeed expected).

Third, our results on productivity growth are important for the literature on the micro-economic origins of aggregate fluctuations. The seminal study by Gabaix (2011) suggested that shocks to specific entities do not “average out” in the aggregate. The mechanism in Gabaix’s paper is that because some firms are very large, the aggregate shock may eventually be dominated by shocks to the largest firms. Here, even though we do not focus on growth rates for the reasons mentioned above, we find heavy tails in the cross sectional distributions of productivity growth rates, suggesting that aggregate fluctuations can arise from large shocks to firms, rather than shocks to large firms. (Schwarzkopf et al. 2010, Gaffeo 2011).

Fourth, our results relate to the work of statistical agencies around the world, who are developing

¹There is an important theoretical literature on productivity dispersion, which recognizes that the right tail of the distribution of productivity levels is a power law, and explains this mostly as the result of innovation and imitation processes, see e.g. Ghigliano (2012), Lucas & Moll (2014), Perla & Tonetti (2014) and König et al. (2016)

procedures to make publicly available summary statistics from micro data, and in particular industry-level productivity dispersion (Cunningham et al. 2018, Berlingieri, Blanchenay, Calligaris & Criscuolo 2017). A key contribution in this paper is to make a suggestion regarding what statistics should be released, namely 5 quantiles necessary to perform Mc Culloch’s (1986) quantile estimation of the four Lévy alpha-stable parameters.

The paper is organized as follows: We describe the data in Section 2 and discuss the GCLT and our methods in Section 3. Section 3 also describes the Lévy alpha-stable distribution and compares it with the Subbotin distribution. Sections 4-5 look into subsamples of firm productivity across firm sizes, countries, years and sectors. Section 6 concludes.

2 Data and descriptive statistics

2.1 Data sources

In this section, we present the dataset and some basic patterns of the empirical data. We use data from the Orbis Europe database, compiled by Bureau van Dijk, which includes the balance sheet and the profit-loss statement of approximately 9 million firms across Europe.

Unlike other widely used firm-level data such as Compustat and Worldscope, Orbis Europe records a large number of small and medium sized firms (SME) that are often not publicly traded. As a consequence, the Orbis Europe database is regarded as more representative of the national economies and as able to capture more realistic firm-level heterogeneity. See Table 1 for the number of observations per country-year.²

Table 1: Number of observations per country-year (after cleaning).

	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015
Bulgaria	-	-	30,747	31,839	35,445	82,474	88,826	93,748	98,571	101,871
Czech Republic	13,454	40,732	36,103	53,061	57,578	63,127	65,641	71,680	72,478	-
Estonia	-	-	11,490	10,661	13,155	14,184	15,728	17,089	18,597	-
Finland	-	23,027	27,838	30,715	30,305	39,020	41,496	43,543	47,180	45,473
France	63,844	115,821	125,991	132,116	146,452	143,029	120,376	141,885	181,717	135,966
Germany	-	26,178	28,150	29,481	31,207	35,491	54,606	64,886	38,390	33,954
Hungary	-	30,844	-	68,925	46,624	48,887	98,182	96,156	102,308	103,976
Italy	22,756	160,815	245,208	216,291	177,640	383,149	431,245	462,610	514,626	-
Portugal	30,465	156,264	164,832	171,027	175,664	179,703	186,245	195,041	205,753	-
Romania	12,163	89,185	99,519	99,167	102,780	114,581	121,781	111,547	123,479	29,468
Slovakia	-	14,333	21,464	40,546	41,700	43,280	42,790	48,072	56,472	63,803
Slovenia	-	-	-	-	15,323	15,964	15,835	17,206	18,671	19,398
Spain	72,916	197,141	227,617	242,518	244,035	252,091	258,405	268,558	283,765	281,029
Sweden	-	64,480	117,356	123,207	126,798	131,091	135,186	138,180	142,357	-
United Kingdom	-	27,450	50,756	52,676	55,811	57,442	60,733	64,151	67,937	70,446

Notes: The table records the number of observations per country-year for all 15 countries. Note that the sample size varies by country. Among 5 largest European countries (France, Germany, Italy, Spain, and United Kingdom), Germany and United Kingdom have relatively a small sample size compared to France, Italy, and Spain. For example, in 2014, Germany has only 38,390 firms, while Italy has 514.625 firms.

In processing and cleaning the data, we only use unconsolidated and consolidated data without com-

²For a more detailed discussions on the advantages and drawbacks of the Orbis Europe database, see Kalemli-Ozcan et al. (2015).

panion statements³ in order to avoid double counting, after which we eliminate duplicated firm-year pairs.

We use various deflators provided by the EUKLEMS project (Jäger & The Conference Board 2018) to deflate sales (GO_P , the EU KLEMS output deflator), to deflate all capital assets (Ip_GFCF , the EU KLEMS capital deflator), and to deflate all income and value-added related variables (VA_P , the EU KLEMS value-added deflator). Finally, we regard negative values for total assets, fixed assets, sales, wage, and employment as missing values. For a detailed discussion on data cleaning, see Appendix A.

Finally, the top and bottom 0.25% of observations are cut to mitigate the influence of outliers in visualizing the results. The estimation result is not substantially sensitive to extreme tails since our estimation method is based on quantile statistics.⁴

2.2 Construction of variables

For each firm i and year t , we compute value-added (Y_{it}) as the sum of deflated wages and deflated Earnings Before Interest and Taxes (EBIT). Denoting employment by (L_{it}), the two measures of firm-level productivity used in this paper are

$$LP_{it} = \frac{Y_{it}}{L_{it}}, \quad (1)$$

$$\Delta LP_{it} = LP_{it} - LP_{it-1}, \quad (2)$$

The first measure of productivity in Eq. (1) is labour productivity (LP_{it}), simply the ratio of value-added to the number of employees. The second measure in Eq. (2) is the change in labour productivity (ΔLP_{it}). Both are expressed in units of currency of the country where the firm is located.

Researchers in productivity analysis often proceed to a data transformation, and analyze the natural logarithm of labor productivity. However, firm-level labor productivity can be negative both theoretically and empirically when output is measured as value-added. From the income perspective of the definition of value-added, the sum of cashflow and wages can be negative since the firm’s loss (negative profits) can be greater than wages. Even from the output perspective, the firm’s intermediate cost can be sometimes greater than revenues in a given fiscal year. Empirically, a sizeable share of firms has negative labor productivity. Table 2 records the proportion (%) of negative observations per country-year in our data. Overall, 5% of the firms in our sample have negative productivity. As a more extreme case, nearly 20% of UK firms during the financial crisis experience negative productivity. For this reason, we use the log-labor productivity and log growth rate only as a point of comparison, to relate our results to the literature and widespread empirical practice.⁵ The full details of the log variables can be found in Appendix G separately, and in the main text our variables are LP and ΔLP .

³For some corporate groups that consist of multiple firms, both consolidated and unconsolidated data is provided, i.e. data is listed for the entire group and again for each of the firms.

⁴The relative difference in the estimated parameters in different sub-samples stay almost the same whether we cut the tails or not.

⁵It is worthwhile to note that, for the log-transformed data, the Levy alpha-stable distribution is outperformed by our reference model Asymmetric Exponential Distribution (AEP) as will be explained later in detail.

Table 2: Proportion (%) of negative observations per country-year (after cleaning).

	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015
Bulgaria	-	-	3.32	4.83	5.01	7.48	7.25	7.48	7.27	6.68
Czech Republic	4.39	3.50	3.86	4.92	4.85	5.15	5.38	5.72	5.47	-
Estonia	-	-	6.34	9.25	7.75	5.50	5.60	5.73	5.78	-
Finland	-	1.92	1.96	2.31	2.16	2.55	2.48	2.62	2.92	3.02
France	1.37	1.45	1.57	1.73	1.60	1.59	1.84	2.05	2.10	2.07
Germany	-	1.19	1.41	1.58	1.36	1.44	1.52	1.66	1.84	1.73
Hungary	-	4.32	-	7.89	6.78	7.63	9.67	7.98	7.61	7.34
Italy	2.55	2.08	2.43	2.98	2.65	3.24	4.44	4.91	5.76	-
Portugal	7.15	5.25	5.67	5.71	5.23	6.27	7.77	7.95	8.54	-
Romania	14.71	9.16	9.72	12.34	14.03	13.08	14.41	14.42	13.33	11.65
Slovakia	-	3.96	4.33	7.84	7.04	7.45	7.45	7.84	7.64	7.06
Slovenia	-	-	-	-	2.87	2.88	2.89	3.09	2.94	2.88
Spain	5.01	3.65	3.89	4.35	4.24	4.78	5.48	5.62	5.16	4.46
Sweden	-	4.25	4.20	4.29	4.12	4.03	4.01	4.06	4.02	-
United Kingdom	-	15.38	22.88	21.68	5.76	3.95	3.91	3.82	3.80	3.85

2.3 Characteristic patterns of the distributions

To motivate Lévy alpha-stable distributions as a model for LP (labor productivity) and ΔLP (labor productivity change), we first show qualitatively that the distributions are heavy-tailed and asymmetric.




Figure 2: **Labour productivity in France: 2006-2015.** The graph on the left displays the histogram of labour productivity levels, and the one on the right displays labour productivity changes. Solid lines indicate the fitted Lévy alpha-stable distributions.

Figure 2 shows the distributions for France for the years covered by the data on a semi-log scale with Lévy alpha-stable fit. When plotted in this way, a normal distribution would appear as an inverse U, and a Laplace distribution as tent-shaped. A simple visual inspection rules out these two standard models. We observe the following five general characteristics of the empirical distribution of the labor productivity and its change, which we will discuss in detail in Sections 4-5.

Unimodality. All distributions appear unimodal.

Large support. The support of the distributions is very large, which means that some firms are vastly more productive than others, and firms experience vastly different productivity changes. Firms may have negative productivity if the sum of profits and wages is negative, since the value-added is imputed from these measures.

Right-skewness of levels. The levels of labor productivity are highly asymmetric, with a pronounced right skewness, which indicates that many firms are much more productive than the ‘typical’ (modal, average or median) firm. Productivity change is more symmetric. We will find that in general, productivity change is slightly left skewed.

Heavy tails. To examine tail behaviour in detail, Figures 16 and 17 in Appendix F.1 show the tails on a log-log plot, i.e. the most and least productive 5% of firms respectively. The right tails appear remarkably linear in log-log scale, suggesting a power law scaling on a relatively long range. This is less clear for the left tails.

Persistence. Comparing the results for different years shows that while there is some movement from year to year, overall the shape of the distribution is very persistent, as one would expect (We come back to this in Section 4, Fig. 5).

2.4 Testing for Infinite Variance

The heavy-tailed nature of the data means that moments higher than the tail exponent do not exist. Trapani (2016) devised a formal test for the existence of any finite moment p . The testing procedure exploits the divergent nature of non-finite sample moments. A full overview of the testing procedure and the results are available in Appendix C. Table 3 presents the test statistics under the null hypothesis that the second moment of firm-level labour productivity is infinite for the 2015 year sample. (see Appendix C for the results in other years, which are similar). Values of the test statistic close to zero indicate that the second moment is likely to be infinite, as indicated by a high p.value.

Variable	Moment	France	Germany	Italy	Spain	United Kingdom
LP	First	598.26 (0.00)	61.10 (0.00)	1210.42 (0.00)	616.46 (0.00)	185.10 (0.00)
LP Change	First	317.65 (0.00)	26.41 (0.00)	677.92 (0.00)	122.73 (0.00)	22.09 (0.00)
LP	Second	0.00 (0.98)**	0.00 (1.00)**	0.10 (0.75)**	0.00 (0.99)**	0.00 (0.98)**
LP Change	Second	0.00 (0.98)**	0.00 (0.97)**	0.00 (0.99)**	0.00 (1.00)**	0.00 (0.97)**

Table 3: Test statistics ϑ_{nr} for Trapani’s test for infinite moments of order 1 and 2 in the levels and changes of LP. P.values underneath in parentheses. Test is run on Firm-level data from ORBIS Europe for 2014.

First of all, the test clearly rejects that the first moment of LP and $LPchange$ are infinite. Second, the test fails to reject the null hypothesis that the second moment is infinite for any reasonable level of confidence.

These results demonstrate that labor productivity at the firm level will be poorly modeled by any density function that relies on the second moment as a parameter. The next section presents one candidate model to explain both the body and the tails of the distributions.

3 Models and estimation methods

To account for the characteristics of the distribution described in the previous section (unimodality, skewness, heavy tails), we consider the Lévy alpha-stable distribution (Nolan 1998) as our main model for the firm-level productivity. As we will explain, the Lévy alpha-stable distribution is based on the generalized central limit theorem and is a natural candidate for our productivity data since the firm’s productivity can be understood as the sum of all productivities of the sub-units of the firm. Furthermore, the tails of the Lévy alpha-stable distribution asymptotically approach a power law and thus could explain the scaling of the standard deviation in the empirical data as shown Figure 1. To compare the model performance of the Lévy alpha-stable distribution, we use another highly flexible four-parameter model, the skew generalized normal family often called the *asymmetric Subottin* or *asymmetric exponential power (AEP)* distribution (Bottazzi & Secchi 2011), as a reference model. We will explain the AEP in more detail in Section 3.4.1. See Bottazzi & Secchi (2011) for a thorough discussion on the AEP.

3.1 The Lévy alpha-stable distribution

Mitchell (1915), who studied commodity and security prices, is perhaps the earliest reference on departure from Gaussian distributions in economic data. As more granular data were collected by empirical economists this disconnection became more obvious, “simply because the empirical distributions of price changes are usually too “peaked” to be relative to samples from Gaussian population” (Mandelbrot 1963). Mandelbrot also pioneered the use of a Lévy stable distribution for incomes and city sizes (Mandelbrot 1960, 1963), while others applications followed soon afterwards (Fama 1965, Samuelson 1967, Embrechts et al. 1997).

The Generalized Central Limit Theorem (GCLT henceforth) states that the only possible non-trivial limit of normalized sums of independent and identically distributed terms is the Lévy alpha-stable distribution. As many physical and economic phenomena, including firm productivity, can be modelled as linear combinations of variables, this property makes Lévy alpha-stable distributions a strong candidate in many cases. We are interested in aggregations of distributions of the form

$$X = \frac{\sum_{n=1}^N X_n - a_n}{b_n}, \tag{3}$$

where X_n are N i.i.d. random variables with the same distribution as X and a_n, b_n are normalization constants. In the case of firm productivity, X_n can be interpreted as the productivity of sub-units in the firm, e.g. the productivity of each division, task or contract. The derivation of the Lévy alpha-stable distribution as the limiting distribution of X is longwinded. For details, see Gnedenko & Kolmogorov (1954), Bouchaud & Potters (2003), and Uchaikin & Zolotarev (2011).

The Lévy alpha-stable distribution is a four-parameter distribution with parameters $\alpha, \beta, \gamma,$ and δ in one of Nolan’s (Nolan 1998) parameterizations (often called \mathbf{S}_0 parameterization) being interpreted as tail exponent, skew, scale, and shift parameters. The density function exists in closed form only exists in a few special cases such as $\alpha = 2$ (Gaussian), $\alpha = 1$ (Cauchy), and $\alpha = 0.5$ (standard Lévy). In the

general case, we can express the Lévy alpha-stable distribution using its characteristic function (Nolan 1998)

$$\varphi(t) = \mathbb{E}[\exp(itX)] = \begin{cases} \exp(-\gamma^\alpha |t|^\alpha [1 + i\beta \tan(\frac{\pi\alpha}{2}) \operatorname{sgn}(t) ((\gamma|t|)^{1-\alpha} - 1)]) + i\delta t & \alpha \neq 1 \\ \exp(-\gamma|t|[1 + i\beta \frac{2}{\pi} \operatorname{sgn}(t) \log(\gamma|t|)] + i\delta t) & \alpha = 1 \end{cases} \quad (4)$$

where i is the imaginary unit, $t \in \mathbb{R}$ is the argument of the characteristic function, and $\operatorname{sgn}(t)$ is a sign function. The parameter $\alpha \in (0, 2]$ (tail exponent or stability index) determines the tail behavior. The lower α the thicker the tail. $\beta \in [-1, 1]$ is a skewness parameter; the higher β the more right-skewed the distribution. $\beta = 0$ gives a symmetric density. $\gamma \in [0, +\infty]$ (scale) determines the width of the density, while the parameter $\delta \in (-\infty, +\infty)$ is the central location of the density. This particular parameterization is useful for numerical work and statistical inference since the characteristic function is continuous in all four parameters.

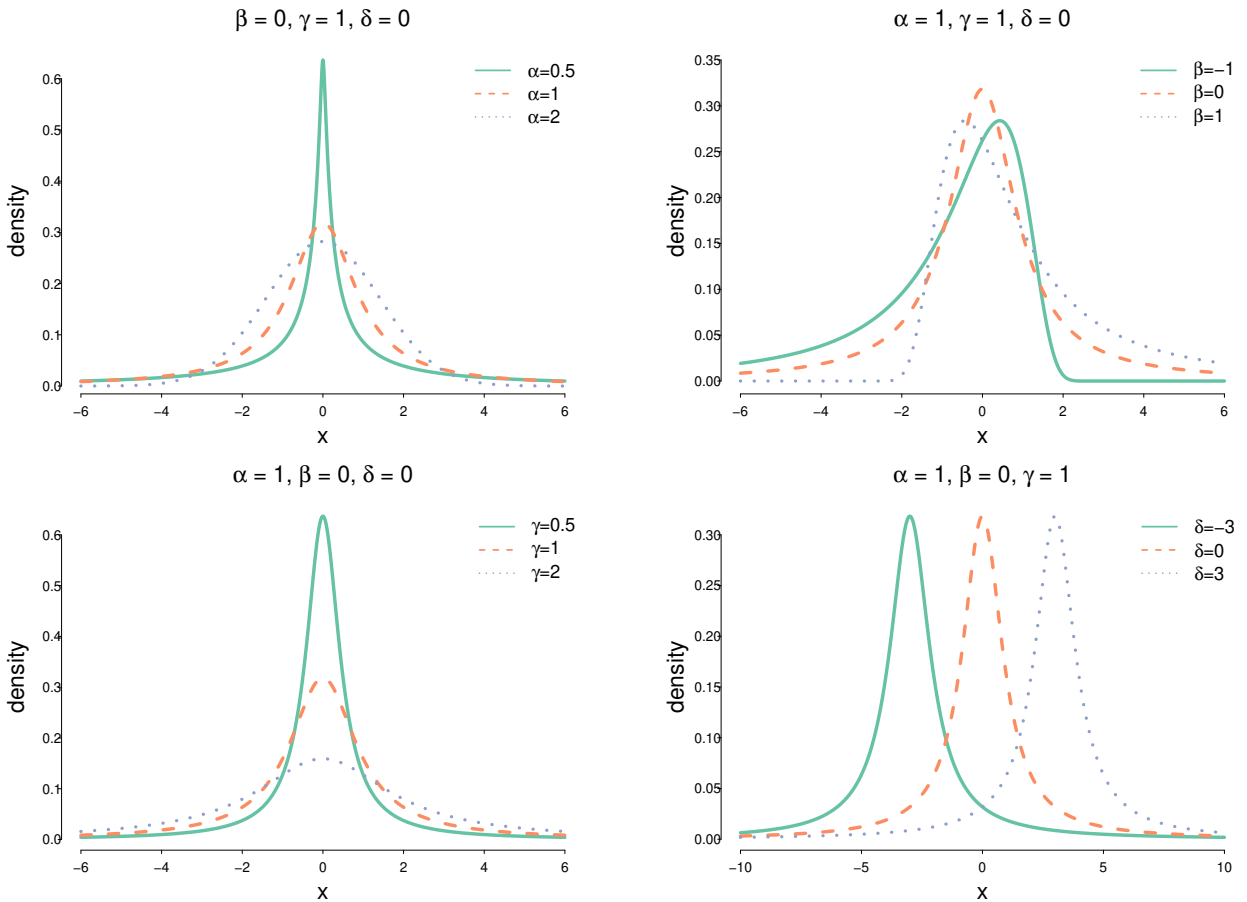


Figure 3: **Probability densities of Lévy alpha-stable distributions by parameter values.** The four plots demonstrate the shape of the Lévy alpha-stable distribution for given parameter values. The top left plot shows varies the α parameter (tail exponent), which determines how heavy the tail is. The top right plot varies the β parameter (skew parameter). The bottom left plot varies the γ parameter, which determines the ‘scale’, or ‘width’, of the distribution. The last plot varies the δ parameter, which shifts the location of the modal value of the distribution.

Figure 3 shows the probability density function of the Lévy alpha-stable distribution with three different values for the four parameters. For example, the top left panel shows the distribution for three different values of α , and fixing $\beta = 0, \gamma = 1, \delta = 0$. A higher α corresponds to a heavier tail. Note that for $\alpha = 2$ the distribution becomes a Gaussian distribution with variance $2\gamma^2$. For $\alpha = 1, \beta = 0$ the distribution becomes a Cauchy distribution with scale parameter γ and location parameter δ . For $\beta = 0$ the distribution is symmetric, for $\beta > 0$ and $\beta < 0$ it is left- and right-skewed respectively (except when $\alpha = 2$, as the skew parameter β vanishes in that case).⁶

3.2 Dispersion parameters in the Lévy alpha-stable distribution

One of the key characteristics of the Lévy alpha-stable distribution is that its typical parametrizations separate the tail behavior from the overall scale of the distribution. While the scale parameter, γ , indicates the width in the body of the distribution, the tail parameter, α , captures the relative prevalence of extreme values. In our present context this distinction becomes particularly important when measuring the “dispersion” of the productivity distribution as often discussed in the recent literature. Many contributions appear to find a rising productivity dispersion (Andrews et al. 2016, Berlingieri, Blanchenay & Criscuolo 2017a, Haldane 2017, Cette et al. 2018). In the unimodal heavy-tailed distribution, the dispersion can be driven either by the scale or by the tail. A higher scale parameter γ (keeping the tail parameter constant) indicates a more equal density and thus a wider body of the distribution. From this perspective, the closer the distribution is to the uniform distribution ($\gamma \rightarrow \infty$), the higher the degree of dispersion. The bottom left panel in Figure 3 graphically shows the effect of the scale parameter γ .

In contrast, the tail parameter α (keeping the scale parameter constant) indicates the relative prevalence of the extreme values in the distribution. For the unimodal distribution to have a heavy-tail and thus allow for more realization of extreme values, it needs to have a relatively narrow body. Put differently, the higher the tail parameter, the more peaked the distribution becomes so that the density centers around a narrower body. The tail is correspondingly wider, decaying as a power law with exponent α .⁷ As a result, extreme values are very common and can dominate any moment of the distribution. The top left panel in Figure 3 shows that the lower α , the more peaked the distribution becomes given the same scale parameter. At the same time, the tails become heavier as seen for values $|x| > 4$ in this example.

γ and α indicate qualitatively different aspects of dispersion, providing a richer explanation of the observed patterns. In discussing the estimation results in the following sections, we will compare the estimated α and γ with other conventional dispersion measures such as standard deviation and interquartile ratio (IQR) to demonstrate the strength of the Lévy alpha-stable distribution as a distributional model for understanding the dispersion of heavy-tailed data, in our case specifically the productivity dispersion.

⁶To give a sense of the magnitude of the tail parameter, take the French sample in 2010 whose estimated α parameter is 1.36. The parameter values for β, γ , and δ , are 0.89, 14.9, and 42, respectively. Suppose that the tail becomes heavier and now the α parameter is 1.26 given that other parameters stay the same. In this case, the probability of the firm to reach the productivity of €(1 million)/employee is 1.6 times higher in the latter with $\alpha = 1.26$. The difference increases as we look at more extreme cases. For example, the probability reaching €(10 million) per employee productivity is twice higher in $\alpha = 1.26$.

⁷Here α is the exponent of the complementary cumulative density $1 - CDF(X)$, which indicates what percentile of observations is larger than x . Since we have $P(x > X) = 1 - CDF(x) \approx Cx^{-\alpha}$, scaling the value of the observation with a factor h just scales the tail probability with $h^{-\alpha}$, since $h^{-\alpha}P(x > X) \approx C(hx)^{-\alpha} = h^{-\alpha}Cx^{-\alpha}$. E.g., if $\alpha = 1$, values if x that are twice (i.e., $h = 2$) as large as some reference \tilde{x} will be half as common, no less. Tails of thin tailed distributions decay much quicker.

3.3 Fitting methods

For the estimation of Lévy alpha-stable distributions, we use the quantile-based estimation method by McCulloch (1986) extending an idea of Fama & Roll (1971). This method is computationally faster compared to other standard methods due to the fact that it uses only the sample quantiles for the estimation (McCulloch 1986).⁸ It is reliable and gives an accurate estimator when the sample size is large enough.

The quantile-based method uses five empirical quantiles to construct an unbiased and consistent estimator of the Levy stable parameters (McCulloch 1986). To show this, McCulloch first proved that the scale and location parameters are independent of any interquantile ranges of the stable distribution. Then, he constructed two functions that relate the tail and skewness parameters to specific interquantile ranges:

$$\Phi_1(\alpha, \beta) = \max\left(\frac{x_{0.95} - x_{0.05}}{x_{0.75} - x_{0.25}}, 2.439\right) \quad (5)$$

$$\Phi_2(\alpha, \beta) = \frac{x_{0.95} - x_{0.5}}{x_{0.95} - x_{0.05}} - \frac{x_{0.5} - x_{0.05}}{x_{0.95} - x_{0.05}} \quad (6)$$

Replacing the theoretical quantiles with the sample quantiles, these two functions provide an estimator of α and β based on the tabulated density table of the Lévy alpha-stable distribution (Fama and Roll, 1968; Holt and Crow, 1973). Given the estimated values of α and β , the scale and location parameters are estimated by two auxiliary functions:

$$\Phi_3(\alpha, \beta) = \frac{x_{0.75} - x_{0.25}}{\gamma} \quad (7)$$

$$\Phi_4(\alpha, \beta) = \frac{\delta - x_{0.5}}{\gamma} + \beta \tan\left(\pi \frac{\alpha}{2}\right) \quad (8)$$

Replacing the theoretical quantile with the sample quantiles, γ and δ are fully determined. We use the R-package **StableEstim** (Kharrat & Boshnakov 2016). As McCulloch pointed out, the quantile-based method leads to a consistent estimator. We use sample sizes of at least 10,000 observations for country-year samples, 5,000 for country-firm size samples, and 1,000 for country-sector samples.

To compute standard errors, we use a bootstrapping with 1000 resampling repetitions.⁹

3.4 Model Comparison

3.4.1 Reference model: 4-parameter AEP (Subbotin) distribution

As a reference model to compare the fit of the Levy stable distribution, we will use the AEP distribution, which is also known as the Asymmetric Exponential Power distribution. The AEP distribution generalizes the Gaussian model by augmenting it with additional tail and skewness parameters. The probability

⁸Other alternatives include Maximum Likelihood, which is computationally expensive, and the general method of moments (GMM). Appendix D.2 shows that the quantile method outperforms the GMM method in terms of accuracy and goodness of fit.

⁹The estimated parameters and standard errors are very similar in the quantile-based estimation with the bootstrapping and the Generalized Method of Moments (GMM). However, the GMM tends to be unstable in estimation and performs poorly in terms of accuracy and goodness of fit. Therefore, the paper only reports the results based on the quantile-based estimation with the bootstrapping error.

density function of the distribution is

$$f(x) = \frac{\kappa h}{\sigma(1 + \kappa^2)\Gamma(1/h)} \exp[-(k^{\text{sgn}(x-\xi)}(|x - \xi|/\sigma))^h],$$

where ξ, σ, h and κ are the location, scale, tail and skewness parameters. We use L-moments method to estimate these parameters. For a detailed discussion, see Asquith (2014). For the computation, we use the R-package **lmomco** (Asquith 2018).

3.4.2 Comparison of the fit: SOOFI ID, AIC, and K-fold cross validation

For the goodness of fit and model comparison measure, we will use the K-fold cross-validation (CV) for the out-of-sample comparison (testing the predictive power of the model for new observations), Soofi's (1995) information distinguishability (ID) index and the Akaike (1973) information criterion (AIC) for in-sample comparison (testing the accuracy of model fitting). The higher CV log likelihood, the higher Soofi ID index, and the lower AIC, the better the model performance. A detailed description of each method can be found in Appendix E.

The comparison of the goodness of fit measures for all country-year samples is visualized in Figure 4. It can be seen that the Lévy alpha-stable is - by all three measures - a better fit in almost every one of the country-year samples: the former has a higher predictive power from k-fold CV, higher Soofi-ID, and a lower AIC.




Figure 4: **Goodness-of-fit comparisons of Lévy alpha-stable and AEP distributions in country-year subsamples.** Left panel: Difference of likelihoods of k-fold cross-validation (Lévy alpha-stable minus AEP). Middle panel: Soofi ID scores. Right panel: Difference of AIC values (Lévy alpha-stable minus AEP). Observations in the white areas indicate subsamples for which the Lévy alpha-stable is the better model, those in the shaded area indicate subsamples with better AEP than Lévy alpha-stable fit. According to all three statistics, Lévy alpha-stable performs better than AEP in almost all subsamples.

To understand the overall performance of these two models, Table 8 in Appendix G summarizes the model comparison results by showing the average value of the three measures for the country-level samples. For the labor productivity level, the average difference of Soofi ID index between the Lévy alpha-stable and the AEP distribution is 5.1, implying that the former model extracts 5.1 percentage points more information from the empirical data than the latter model. The relative likelihood per data point both in the AIC and the cross-validation is 0.95. This means that the AEP is 0.95 times as probable as the Lévy

stable distribution for the in-sample and out-of-sample prediction. For the labor productivity change, Lévy stable model extracts 2.2 percentage points more information over average than the asymmetric exponential power distribution. From the average relative log-likelihood, we note that the latter model is 0.98 times as probable as the former model.

Appendix G compares the two models for log-labor productivity and log growth rates. The Lévy alpha-stable is still preferred for log-labor productivity, while the AEP is slightly more preferred for the log growth rate of labor productivity.

4 Estimation result 1: productivity by year

In the following sections, we report parameter fits for LP (labor productivity) and ΔLP (labor productivity change) in subsamples arranged by (1) country and year, (2) country and firm size, (3) country and industry. The Lévy alpha-stable model generally gives an excellent fit. Substantial differences can be found by year, firm size, and industry. We will focus on the five largest EU economy: France, Germany, Italy, Spain, and the United Kingdom.

4.1 Productivity by year

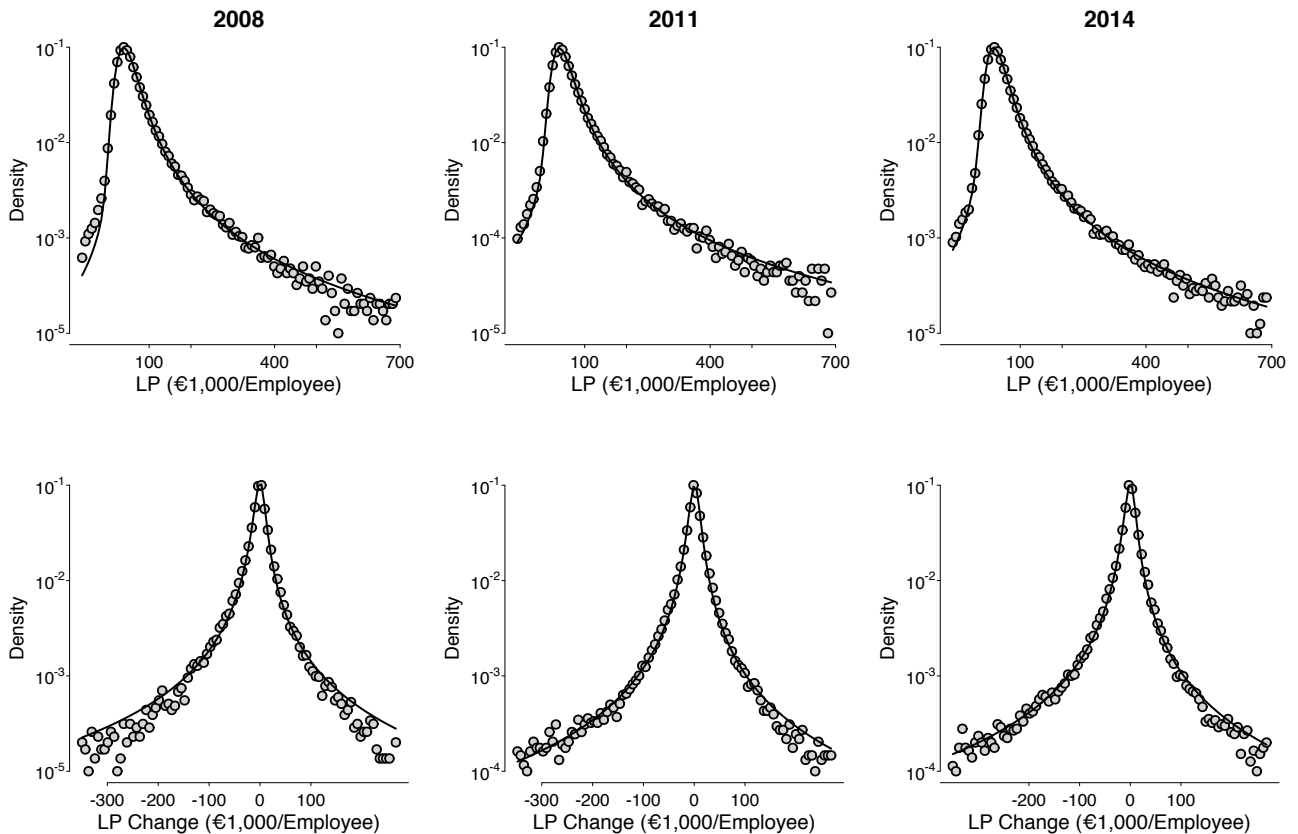


Figure 5: **Distribution of French labor productivity and its change by year with Lévy alpha-stable fits.** Years in 2008, 2011, 2014, which exemplify the other years range from 2007-2015. Solid lines indicate the fitted Lévy alpha-stable distributions, for which the estimated parameters are seen in Figure 6.

Figure 5 shows the original distributions with the fitted lines for three years in the French sample. The Lévy alpha-stable model generally gives a very good fit with the Soofi ID higher than 99% in most cases, while Figure 6 shows the time trend of the estimated four parameters with the bootstrapping standard error.

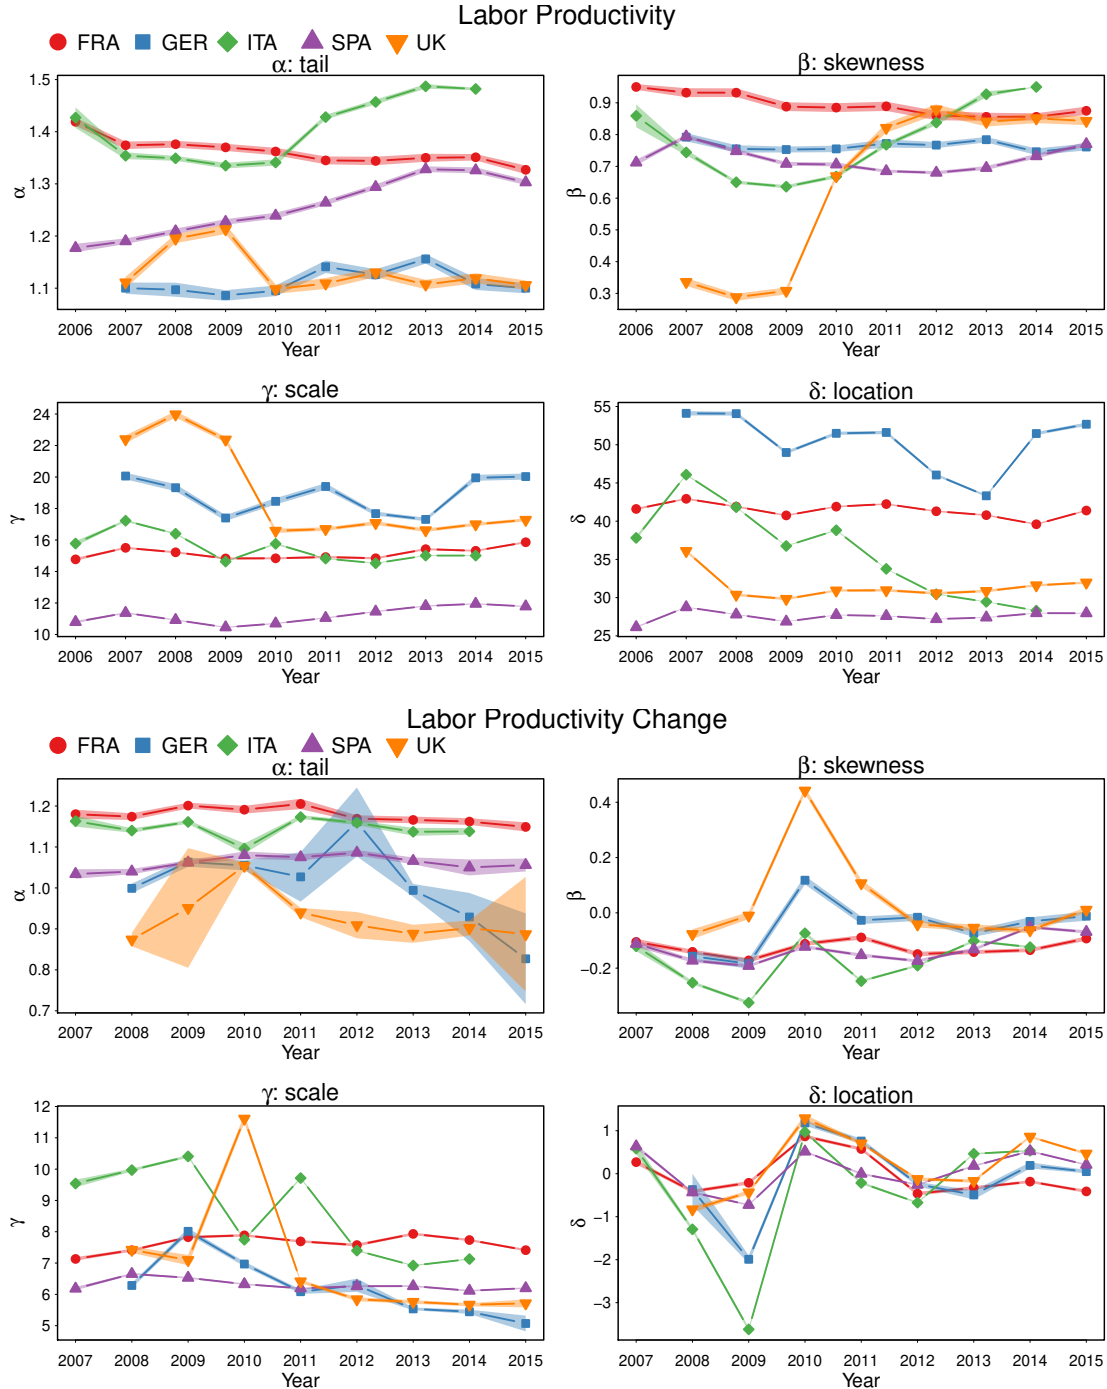


Figure 6: **Estimated parameters for country-year samples.** The four parameters of the fitted Lévy alpha-stable distributions for labour productivity levels (top row) and change (bottom row) are plotted by year (2006-2015 for levels, 2007-2015 for growth). γ and δ are denominated in €1,000/employee, α and β are dimensionless. The shaded area represents the range of ± 1 bootstrapped standard errors.

Productivity level. Before we discuss the yearly variation, it is worth noting that there are considerable differences in the level of estimated parameters among countries. Germany has a relatively low α (heavy tail), high γ (large scale), and high δ (high central value). Considering the positive skewness of the labor productivity distribution, a heavy tail suggests that there is a relatively high presence of extremely successful firms in Germany. A larger-scale parameter with a positive skewness indicates that the body of the productivity distribution of larger firms is spread out across a wider range of support, implying a relatively higher presence of high productivity firms (but not necessarily of extremely successful firms in the right tail) compared to other countries. A higher location parameter also confirms the overall higher labor productivity in Germany. In contrast, Spain has the lowest γ and δ . Especially, the location parameter δ is estimated to be around €30,000 per employee, while Germany has almost twice that.¹⁰ Finally, the skewness parameter in the UK is significantly lower during the financial crisis reflecting the large number of negative productivity firms in the UK as shown in Table 2.

The yearly variation of the estimated parameters for labor productivity has also considerable differences between the countries. The tails of distributions in firm-level productivity have generally thinned in Italy and Spain (higher α), but have not changed noticeably in Germany and the UK, whilst they have grown moderately in France (lower α). Given the positive skewness of the distribution, the result suggests that there has been a relatively smaller presence of extremely successful firms in Spain and Italy. For example, the probability of Italian firms reaching €500,000 per employee is 10 times lower in 2014 than in 2007 (See the supporting material for the yearly distribution of labor productivity for Italy). The change in the skewness parameter β is most dramatic in the UK, jumping from 0.3 to 0.9 after the financial crisis. As discussed before, this result reflects the high presence (around 20%) of negative productivity firms in the UK during the financial crisis, skewing the distribution further to the left. The scale parameter γ of the UK was also significantly higher during the financial crisis due to a large number of negative productivity firms.

Productivity change. There are several interesting patterns in the time trends of productivity change distributions. The most noticeable effect that appears consistently across countries is a slump in the location parameter δ and a decrease in the skewness parameter β during the time of the crisis 2008-2009. The result shows that not only did the central location of the distribution move to the left, but the left tail also became heavier due to a large number of significantly distressed firms during the crisis. The recovery process is rather dramatically reflected in the surge of the skewness parameter and the location parameter in 2010. Especially, Germany and the UK have a relatively high β , indicating a more dramatic positive productivity change in these two countries. It is important to note that the time trend of the location parameter δ captures the aggregate labor productivity growth calculated from the national accounts remarkably well. Fig 22 in the Appendix F.3 compares δ and the labor productivity growth from the Penn World Table database and demonstrates that they have a similar trend in the majority of our samples.

Overall, tails in productivity changes have stayed relatively stable between 2006 and 2015, with the exception of Germany which exhibited a constant decrease in α over the last five years (although not statistically significant). Considering that the skewness parameter β is relatively stable around zero during this period, a heavier tail indicates more extreme outcomes in both left and right tails of the productivity change distribution in Germany.

¹⁰Note that the comparison of the scale and location parameters between the UK and the other countries needs to be done carefully since the UK firms report the financial statements in Pound sterling, not in Euro.

4.2 Pattern of dispersion in recent decades

This section discusses the time trend of our dispersion parameters α and γ of the labor productivity distribution in further detail in the context of the recent debate about a rising productivity dispersion. A recent OECD policy report on productivity dispersion suggested a growing disparity between ‘superstar’ and ‘zombie’ firms (Andrews et al. 2016, Berlingieri, Blanchenay & Criscuolo 2017*b*). The most recent paper (Berlingieri, Blanchenay & Criscuolo 2017*b*) uses variance, percentile ratios (90-10, 90-50, 50-10 ratios), and Inter Quantile Ranges (IQR) of log labor productivity as dispersion measures. Based on the firm-level administrative data collected from multiple European countries, the OECD concludes that there has been widening productivity dispersion among European countries. For example, the 90-10 percentile ratio in log Productivity increased by 12% between 2001 and 2012 according to the report.

To understand the relative change of dispersion in labor productivity over time, we compare our dispersion parameters α and γ with the IQR over our sample period. We chose the IQR and not the percentile ratio as the reference metric of dispersion since the negative values in the left tail of labor productivity distribution make the percentile ratio uninterpretable. The (IQR) is defined as the difference between two given quantiles, and therefore measures how widely the body of the distribution is spread out regardless of the signs of the two quantiles. In comparing these dispersion metrics, we show the results both for labor productivity and log-labor productivity. See Appendix G for the complete estimation result on the log-labor productivity and its change.

Figure 7 shows the mean relative change in the tail and scale parameter of the distribution, and 90-10 interquantile range of labor productivity and log labor productivity for the five countries: France, Germany, Italy, Spain, and the UK. We use the AEP results for the log labor productivity since the AEP outperforms the Lévy alpha-stable distribution for the log variables. The tail and the scale parameters in the AEP are h and σ respectively.¹¹ All the variables are normalized with respect to the initial year in 2007, from which all five countries have 9 consecutive years to 2015. There is a stark contrast between the log and non-log labor productivity. The IQR for log labor productivity shows that the dispersion increased by 9% between 2007 and 2014, followed by a substantial drop in 2015. The scale parameter γ of the Lévy alpha-stable tracks the IQR remarkably well and increased by 12% during this period except for the crisis period, in which the scale parameter did not increase as much as the IQR. From the IQR and the scale parameter, it is reasonable to conclude that the dispersion (of log labor productivity) has indeed increased for the last decade. The tail parameter α has a cyclical pattern, slightly going down during the financial crisis followed by a moderate increase afterward. Since the tail parameter indicates the presence of extreme observations in both tails, the result suggests that the prevalence of extreme values increased during the financial crisis and slowly decreased afterward.

¹¹Note that the estimated parameters of Lévy alpha-stable distribution and the AEP have a very close relationship.

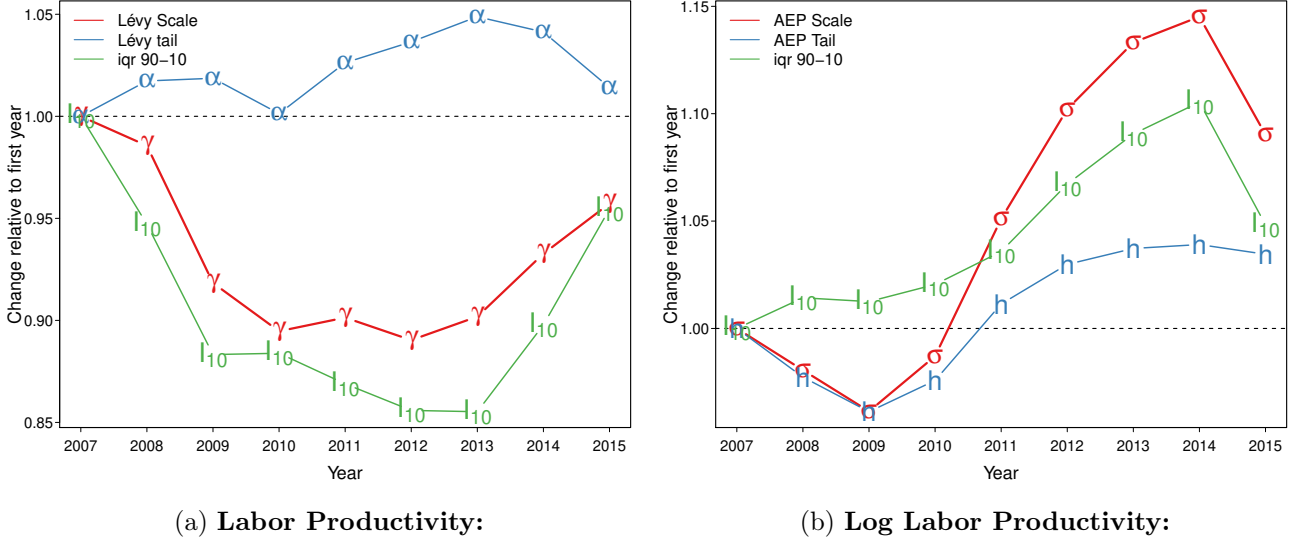


Figure 7: **Relative change of dispersion in labor productivity over time.** The table shows 1) the mean relative change in the tail (α) and the scale (γ) parameters of Lévy alpha-stable distribution, and 90-10 interquantile range of labor productivity and 2) the mean relative change in tail (h) and the scale (σ) of the AEP distribution, and 90-10 interquantile range of log labor productivity for the five countries: France, Germany, Italy, Spain, and the UK.

The result for the non-log labor productivity shows a completely different story. The IQR of labor productivity distribution actually decreased by 5% between 2007 and 2015. The scale parameter γ also dropped significantly from 2007 to 2010, followed by a steep increase afterward. This result suggests that the dispersion of labor productivity actually decreased over the past decade. The tail parameter α also generally increased except for a small drop in 2010, indicating a decrease in dispersion in the tail part.

This stark difference of the dispersion behavior in two different measures of labor productivity does not come from the exclusion of non-positive values from the sample. The time trend of dispersion metrics does not change dramatically whether we exclude non-positive productivity or not (See the supporting material). It is the log transformation itself that changes the time trend of dispersion. Due to the nature of log, the log transformation highlights the difference between firms with very low labor productivity but attenuates the difference between firms with very high labor productivity. Therefore, when the log-transformation is used, a right-skewed distribution with a heavy tail such as the labor productivity distribution is transformed into a more symmetric distribution with moderate tails.

As shown in Figures 6 and 24, the skewness and tail index for France, for example, changes from 0.9 and 1.4 to 0.1 and 1.65 when taking logs, meaning that the right-skewness almost disappeared and the tail gets much thinner. Therefore, from the perspective of the distributional characteristics, the labor productivity and log-labor productivity have two completely different distributions and therefore we cannot expect similar results for the pattern of dispersion in the two distributions.

More specifically, the log transformation puts the variable on a relative scale. While $x_{10\%} - x_{90\%}$ is the difference, $\log(x_{90\%}) - \log(x_{10\%})$ is the log of the ratio: $e^{\log(x_{90\%}) - \log(x_{10\%})} = \frac{e^{\log(x_{90\%})}}{e^{\log(x_{10\%})}} = \frac{x_{90\%}}{x_{10\%}}$. For example, the 90-10 IQRs of labor productivity and log labor productivity for Germany in 2007 are around 123,000 €/employee and 1.54, respectively. This means that the top 10 percent firm produces €123,000 more per employee than the bottom 10 percent firm does in Germany. On a relative scale (i.e. using

log labor productivity), the same top 10 percent firm is 4.7 times more productive than the bottom 10 percent firm since $\exp(1.54) = 4.7$. The relative change of the dispersion in 2014 with respect to the year 2007 is 0.99 for labor productivity (€123,000 \rightarrow €122,000) and 1.04 for log-labor productivity (1.54 \rightarrow 1.60). This means that the absolute difference in labor productivity between the top and the bottom firms decreased by 1% while its relative difference (in a log scale) increased by 4%. This discrepancy arises when both 10% and 90% quantile values decrease but the latter decreases more than the former. Since the log transformation accentuates the left tail, this decrease in the value of the 10% quantile of the labor productivity can result in an increase in its distance from the 90% quantile value in the log scale even when the absolute distance between the 90-10% quantiles is smaller.¹²

This result indicates that the dispersion measure based on the log-transformed variable needs to be analyzed with caution. Appendix I provides more details and examples.

5 Estimation result 2: productivity by size and industry

5.1 Productivity by size

Figure 9 shows the estimated four parameters with the bootstrapping standard error. The original distributions with the fitted lines are visible in Figure 8a. Both figures show that the size-wise variation of the estimated parameter is pronounced.

Productivity level. The parameter α is smaller in the samples of the large and very large firms compared to the small and medium sized firms, suggesting heavier tails for those firms.¹³ This shows that there is a relatively high presence of extremely successful and extremely unsuccessful firms in the large and very large firm groups.¹⁴ Considering the positive skewness of the labor productivity distribution, the result implies that large firms have a higher chance of being substantially more productive than other firms. The scale parameter γ has a higher value as the size of the firm increases, meaning that the body of the productivity distribution of larger firms is spread out across a wider range of support. Considering the positive skewness of the labor productivity distribution, the result implies that larger firms have a relatively higher presence of high productivity firms than the smaller firms.

The skewness parameter β tends to be higher for the large and very large samples but the pattern is not as clear as for the other parameters. Note that the distribution of the small and medium-sized firms in the UK is less right-skewed, reflecting the fact that there are more negative productivity firms in the UK compared to other countries as shown in Table 2. Finally, the location parameter δ has a higher value as the size of the firms increases, meaning that the productivity of larger firms tends to have a higher central value.

Figure 25a in Appendix G shows the log of labor productivity after removing negative values. The overall pattern is similar except for the scale parameter in the distribution of the small firms. The scale

¹²For example, the 10% and 90% quantiles of labor productivity (excluding the non-positive values) in the full sample across all countries in our data set are 32,663 €/employee and 155,958 €/employee in 2007 and 29,265 €/employee and 152,226 €/employee in 2014, showing that both of 10% and 90% quantile values decreased from 2007 to 2014 but the 90% quantile value decreased more than the 10% value. The log-transformed values at 10% and 90% changed from 10.42 and 11.96 in 2007 to 10.33 and 11.93 in 2014. The change in the left tail is more highlighted in the log transformation (10.42 \rightarrow 10.33) than the right tail (11.96 \rightarrow 11.93) even though the absolute change is higher in the 90% quantile values (155,958 - 152,226 = 3732) than the 10% values (32,663 - 29,265 = 3398).

¹³Note that there is a slight overestimation of the tail of the small firms. The Soofi ID is still very high due to the fact that the KL divergence puts relatively less weight on the small probability area. For example, the observation of 1e-05 log-density is 10,000 times less likely than the observation of 1e-01 log-density.

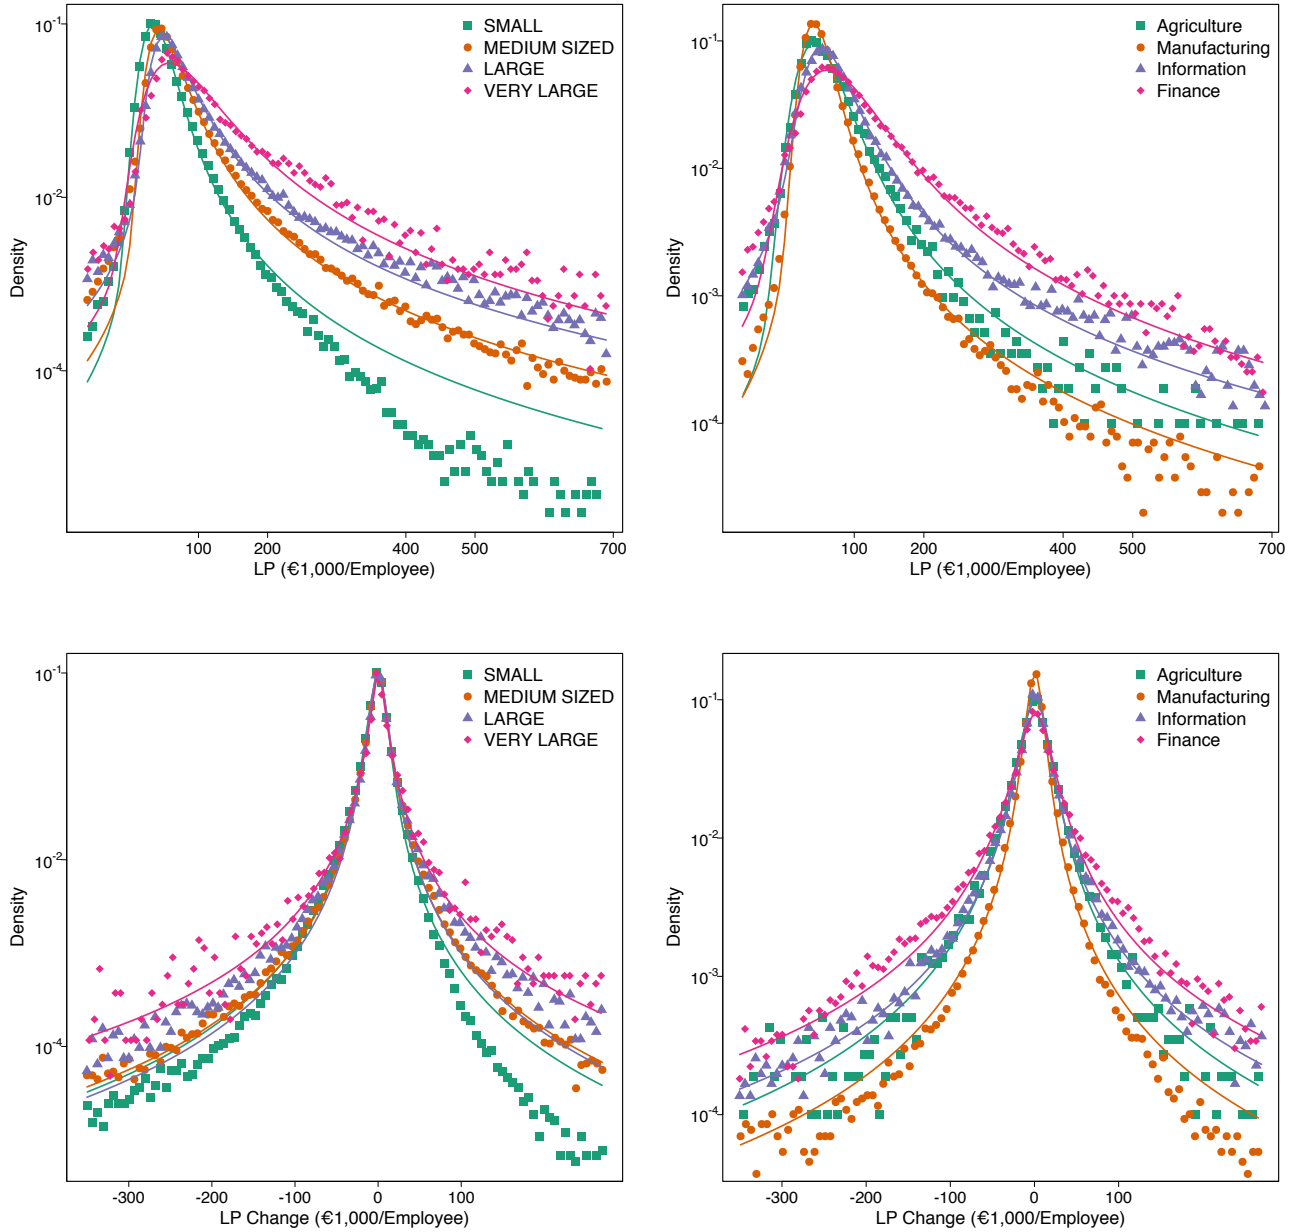
¹⁴Note that the variation in the α by size becomes ambiguous when we remove the small size firms.

parameter is not the smallest anymore and is sometimes greater than that of large firms. This is because the log transformation highlights the firms with very low labor productivity that are prevalent in the small firm group. This is also confirmed by the fact that the log labor productivity distribution has a negative skewness for small size firms for all countries.

Productivity change. The pattern for the parameter α of productivity change is less obvious than in the case of productivity levels. But generally, the large and very large firms tend to have a heavier tail for most countries. This shows that larger firms tend to have extremely high and extremely low productivity changes. As for the productivity level, the scale parameter γ tends to be higher for the large and very large firms, meaning that the body of the distribution of productivity change for larger firms is spread out across a wider range of the support. The fact that the variation in productivity change is greater for the larger sized firms may appear unintuitive and goes against the established literature on the relationship between the firm size and the growth volatility where researchers often find higher volatility for small-sized firms (Stanley et al. 1996, Schwarzkopf et al. 2010, Calvino et al. 2018). This diverging finding may come from the fact that we are studying labor productivity, not firm size, but it may also come from the fact that we are using the first difference (ΔLP) as our measure for productivity change, while the growth volatility literature mostly uses the growth rates (usually, measured as changes in the log). As shown in Figure 25b in Appendix G, the growth rate (log changes) of labor productivity does have a smaller scale for large and very large firms in our data as well. Therefore, we can conclude that large firms tend to have a higher variation in the absolute change in productivity but have a smaller variation in its relative change.

The skewness parameter β and the location parameter δ tend to be higher for the large and very large firms for most countries, indicating a higher presence of firms with positive productivity change and a higher overall productivity change for these firms. Figure 26 in Appendix G shows that the same pattern persists for the log growth rate of productivity.

Summary. Overall, the size variation of productivity distribution is noticeable. Larger firms tend to have a more dispersed pattern (heavier tail and larger scale) with a higher central value in their productivity distribution, both in levels and changes. Also, larger firms tend to have a more positively skewed pattern in productivity change.



(a) **Size:** The top panel shows the distribution of the level of labour productivity in France and firm size. The second row computes the changes in labour productivity. Solid lines indicate the fitted Lévy alpha-stable distributions, for which the estimated parameters are available in Figure 9. Firm size categorisations can be found in Appendix A.5. All sub-groups here meet the threshold observation count of 5,000.

(b) **Industry:** The first graph shows the distributions of the level of labour productivity in France by industry. The second computes the changes in labour productivity. Solid lines indicate the fitted Lévy alpha-stable distributions, for which the exact parameterisations are available in Figures 10a and 10b. Firms are categorised by their provided Nace Rev. 2 classification code. An overview of the industry headers and corresponding descriptions is available in Appendix A.5. All groups here meet the threshold observation count of 1,000.

Figure 8: Distributions of labor productivity in France with Lévy alpha-stable fits.

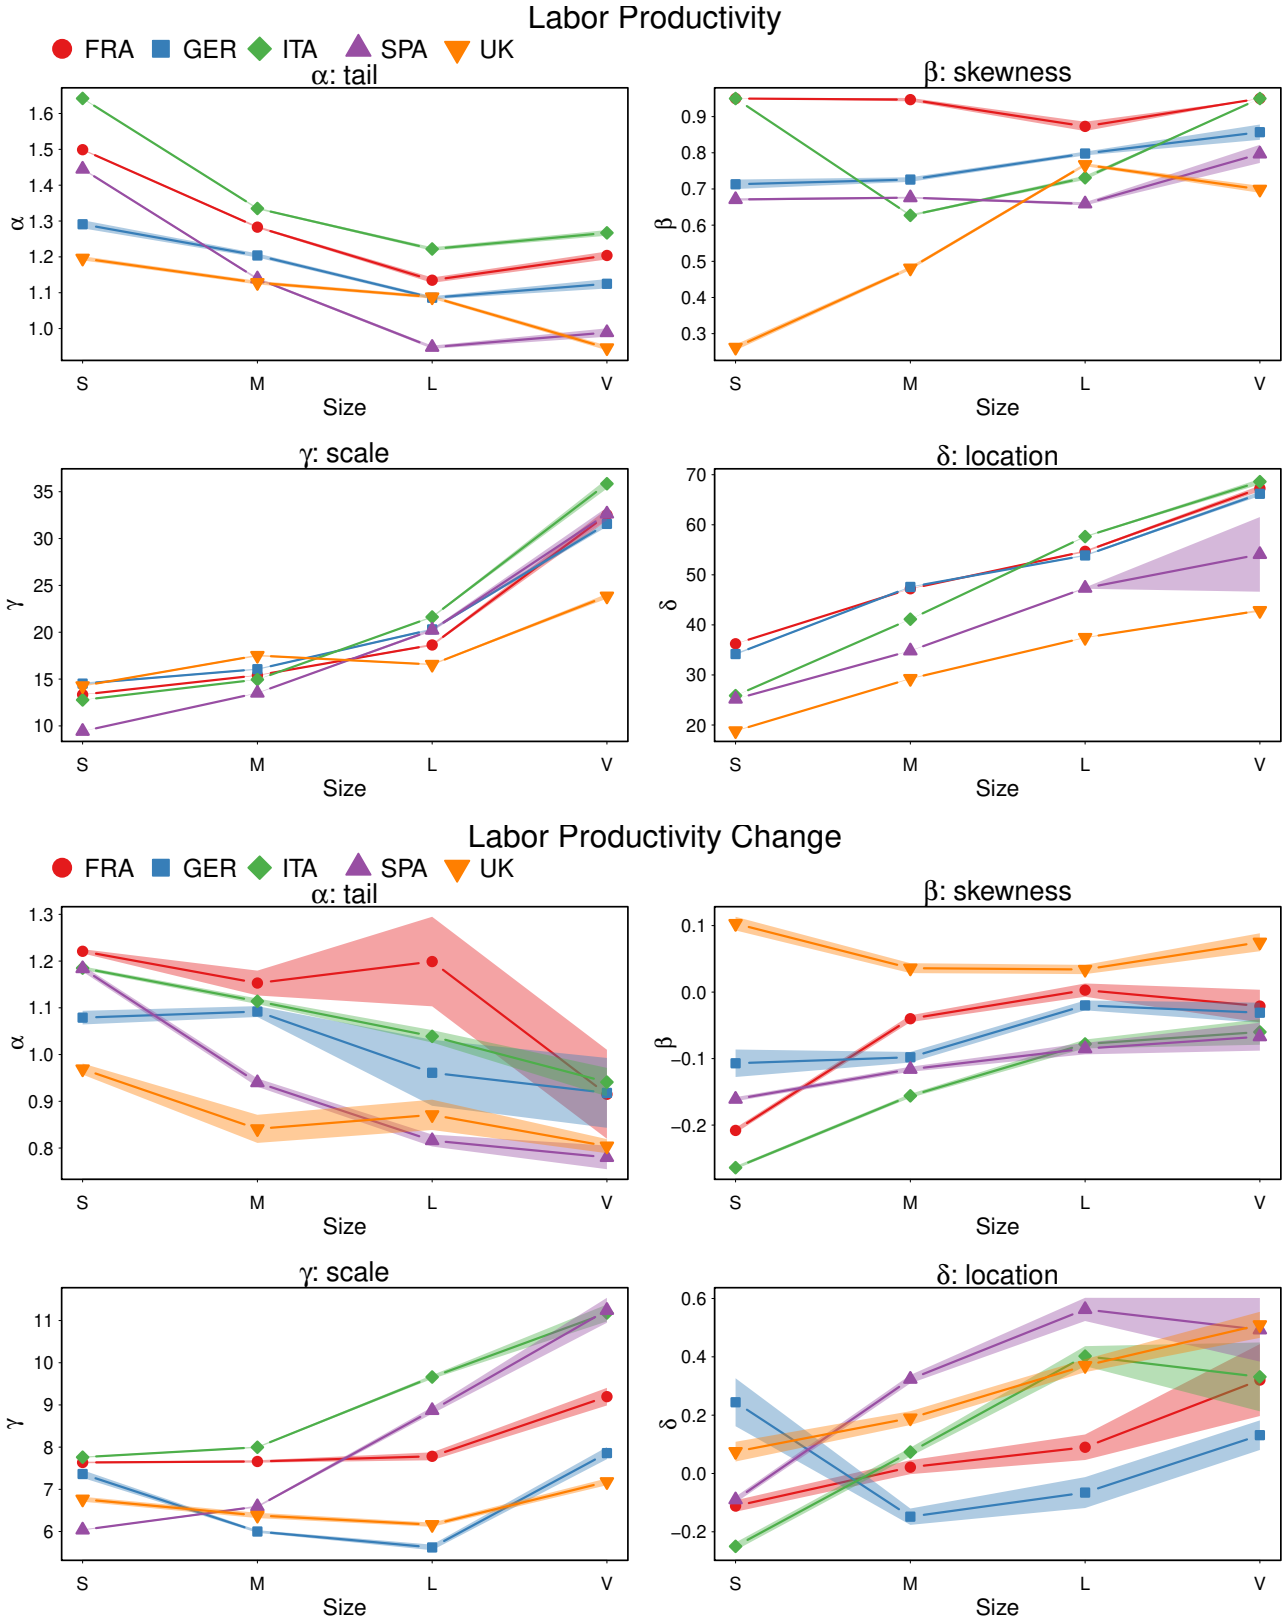


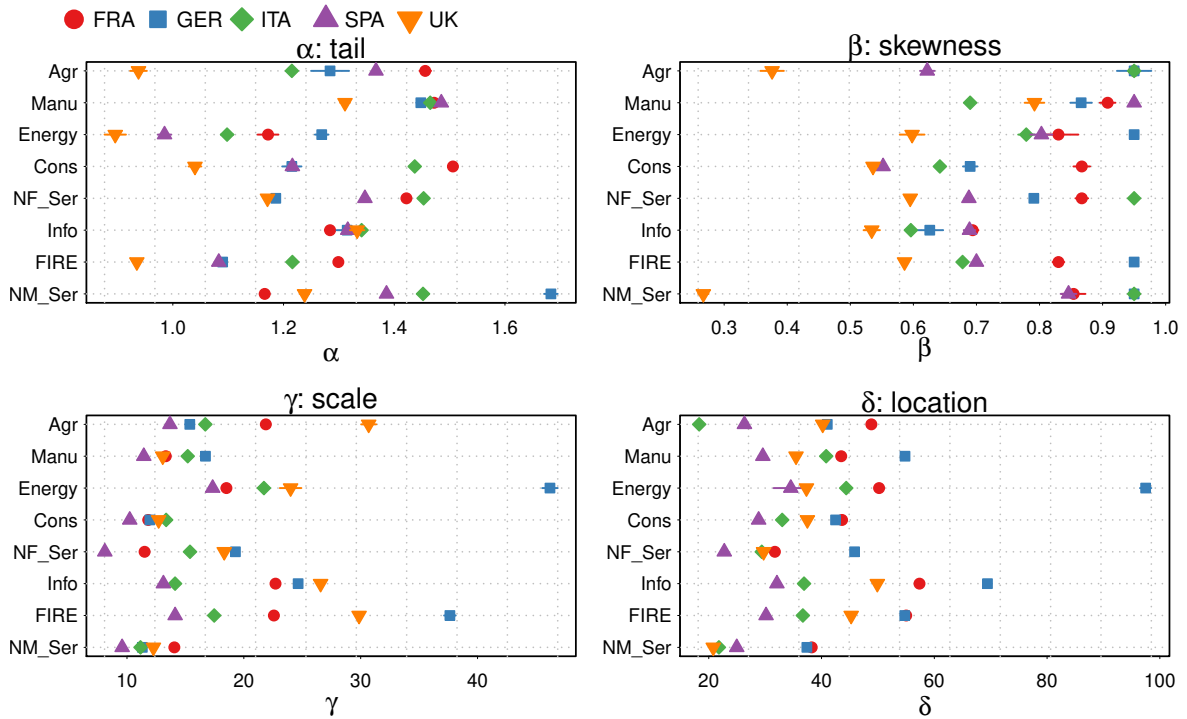
Figure 9: **Estimated parameters country-size samples.** The four parameters of the fitted Lévy alpha-stable distributions for labor productivity levels (top row) and growth (bottom row) are plotted by firm size (small, medium, large, very large). γ and δ are denominated in €1,000/employee, α and β are dimensionless. The shaded area represents the ± 1 bootstrapped standard errors.

5.2 Productivity by industry

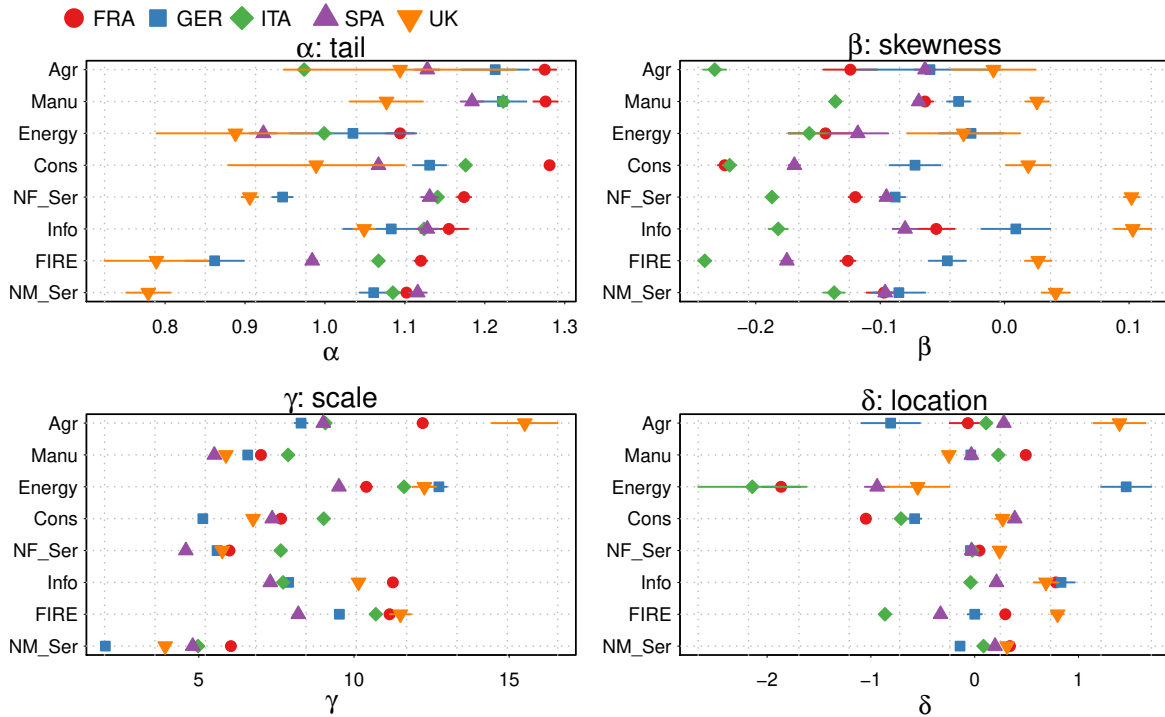
Figures 10a and 10b show the estimated mean value of the four parameters and the bootstrapping standard error of all 5 countries, while Figure 8b shows the original distribution for 4 representative industries with the fitted lines. The fit is generally good as the year and the size samples. We repeat the same exercise for a more disaggregated industry classification in Appendix H, which has similar results.

Productivity levels. The industry-wise variation of the estimated parameters for labor productivity distribution is generally noisy. For the tail behavior, the *FIRE* (Finance, Insurance, Real Estate, and Professional Service) and *Energy* sectors tend to have heavier tails than others. No clear pattern exists for the skewness parameter except that all sectoral productivity distributions are right-skewed. *FIRE* and *Info* (Information and Technology) tend to have a high scale parameter. *Info* has a relatively higher location parameter as well. Overall, the *FIRE* sector has a relatively high degree of dispersion driven by both in the tail and the body as indicated by the low tail parameter α and the high scale parameter γ .

Productivity changes. Results for ΔLP are similar to the result for the levels, *FIRE* firms tends to have a heavier tail and a larger scale γ , suggesting a high degree of dispersion both in the tail and the body part of the distribution. Interestingly, ΔLP are generally left-skewed, with exception for industries in the UK. The location parameters are relatively well centered around zero, although large cross-country differences emerge for *Agr* (Agriculture), *Energy*, and *Info*. The *Info* sector tends to have a higher location parameter.



(a) Labour Productivity Levels



(b) Labor Productivity Changes

Figure 10: **Estimated parameters for country-industry sub-samples.** γ and δ are denominated in €1,000/employee, α and β are dimensionless. The horizontal bar represents the ± 1 bootstrapped standard errors.

By observing the dispersion pattern at the industry level, we hypothesized that dispersion is higher in capital intensive industries. Figure 11 shows the scatterplot of the estimated α and γ against the aggregate capital-labor ratio (Fixed capital/Number of Employees) at the industry level for 5 countries. We use the original non-aggregated NACE classification to obtain more observations of the country-industry pairs. The same result can be found for the aggregate classification.

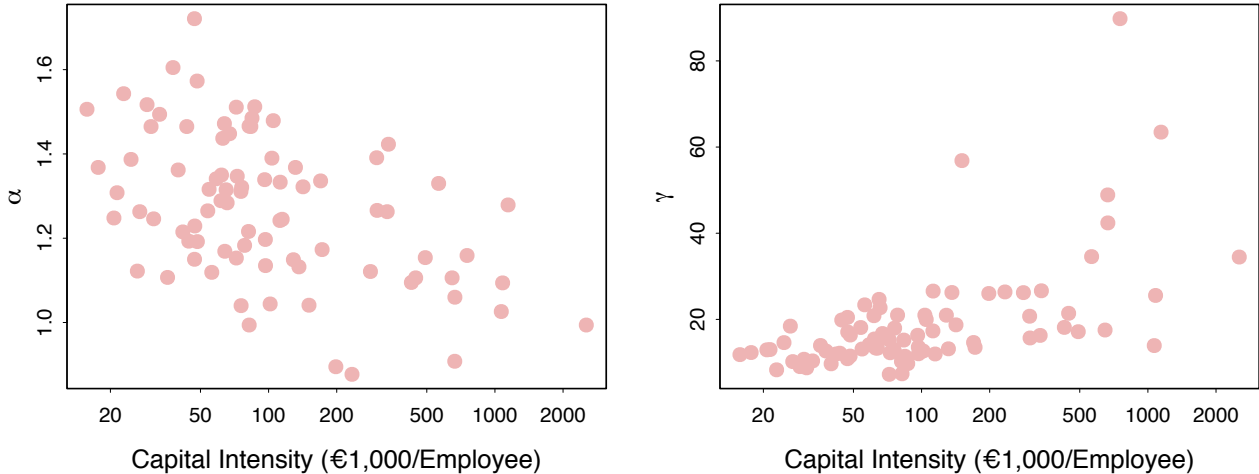


Figure 11: **Industry-level productivity dispersion and capital intensity**

It is clear from Figure 11 that the higher the industry-level capital intensity is, the larger the dispersion in the firm-level labor productivity (the lower α and the higher γ). An interesting research agenda would be to explore how other determinants of dispersion (competition, misallocation, institutions, etc.) affect γ and α separately.

6 Conclusion

The distributions of firm-level productivity levels and growth follow persistent patterns. These distributions are unimodal, have a large support, are asymmetric and heavy-tailed. The most important consequence of this is that measuring dispersion is a delicate question: standard deviations are poor metrics because second moments do not exist, and interquantile ranges are hard to interpret due to the asymmetry.

We propose the Lévy alpha-stable distribution as a sensible distributional model for labor productivity, motivated by the generalized Central Limit Theorem on the one hand, and by the empirical measurement of an infinite standard deviation in our productivity data, on the other hand.

Good distributional models make it possible to offer a richer picture of dispersion. While the scale parameter captures the overall width of the distribution, the tail parameter captures the occurrence of the extreme events. These are qualitatively distinct aspects of dispersion.

In almost all sub-samples of Orbis Europe firm-level data for the EU with sufficiently many observations, excellent fits of the distribution can be obtained with the Lévy alpha-stable model. This holds for sub-samples by country and year, by country and firm size group, and by country and industry for an observation period spanning ten years, the 2008 financial crisis, and the subsequent Euro crisis.

Parameter values obtained in the fits are relatively concentrated and suggest mostly heavy tails distributions with a tail exponent typically slightly above $\alpha = 1$, but always $\alpha < 2$, indicating that the LP and ΔLP distributions do not have a finite variance.

Strong patterns can be observed in the sub-samples by year and firm size group: Larger size groups have more heavy tailed labour productivity distributions with higher scale and location (central value) parameters. They have generally lower skew and higher location in their productivity growth distributions. Noticeable differences can be identified as well between industries. Sectors with higher capital intensity, such as utilities, mining, and finance and real estate tend to be more dispersed and have a higher central value in both productivity levels and growth.

Our finding that Lévy alpha-stable distributions are a good distributional model for LP and ΔLP has implications for a variety of ongoing debates ranging from the productivity slowdown, to misallocation and the origins of aggregate fluctuations. It provides a starting point for more realistic models of firm dynamics, which should attempt to match the properties of these distributions.

Acknowledgement

We would like to acknowledge funding by Baillie Gifford, the Institute for New Economic Thinking and the Rebuilding Macroeconomics project, which is funded by the ESRC.

References

- Acemoglu, D. & Dell, M. (2010), ‘Productivity differences between and within countries’, *American Economic Journal: Macroeconomics* **2**(1), 169–88.
- Aghion, P., Bloom, N., Blundell, R., Griffith, R. & Howitt, P. (2005), ‘Competition and innovation: An inverted-u relationship’, *The Quarterly Journal of Economics* **120**(2), 701–728.
- Akaike, H. (1973), ‘Maximum likelihood identification of Gaussian autoregressive moving average models’, *Biometrika* **60**(2), 255–265.
- Andrews, D., Criscuolo, C. & Gal, P. N. (2016), ‘The best versus the rest’, *OECD Productivity Working Papers* .
- Aoyama, H., Yoshikawa, H., Iyetomi, H. & Fujiwara, Y. (2010), ‘Productivity dispersion: facts, theory, and implications’, *Journal of economic interaction and coordination* **5**(1), 27–54.
- Asquith, W. (2018), *lmomco - L-moments, censored L-moments, trimmed L-moments, L-comoments, and many distributions*. R package version 2.3.1.
- Asquith, W. H. (2014), ‘Parameter estimation for the 4-parameter Asymmetric Exponential Power distribution by the method of L-moments using R’, *Computational Statistics & Data Analysis* **71**, 955 – 970.
- Axtell, R. L. (2001), ‘Zipf distribution of U.S. firm sizes’, *Science* **293**(5536), 1818–1820.
- Bartelsman, E., Haltiwanger, J. & Scarpetta, S. (2013), ‘Cross-country differences in productivity: The role of allocation and selection’, *American Economic Review* **103**(1), 305–34.

- Bartelsman, E. J. & Wolf, Z. (2018), Measuring productivity dispersion, *in* E. Grifell-Tatjé, C. K. Lovell & R. C. Sickles, eds, ‘Oxford Handbook of Productivity Analysis’, Oxford University Press, United Kingdom.
- Berlingieri, G., Blanchenay, P., Calligaris, S. & Criscuolo, C. (2017), ‘The multiprod project: A comprehensive overview’, *OECD Science, Technology and Industry Working Papers* **2017**(4), 1.
- Berlingieri, G., Blanchenay, P. & Criscuolo, C. (2017a), ‘The great divergence(s)’, *OECD Productivity Working Papers* .
- Berlingieri, G., Blanchenay, P. & Criscuolo, C. (2017b), ‘The great divergence(s)’, *OECD Science, Technology and Industry Policy Papers* **39**.
- Bottazzi, G., Cefis, E., Dosi, G. & Secchi, A. (2007), ‘Invariances and diversities in the patterns of industrial evolution: Some evidence from italian manufacturing industries’, *Small Business Economics* **29**(1-2), 137–159.
- Bottazzi, G. & Secchi, A. (2011), ‘A new class of asymmetric exponential power densities with applications to economics and finance’, *Industrial and Corporate Change* **20**(4), 991–1030.
- Bouchaud, J.-P. & Potters, M. (2003), *Theory of Financial Risk and Derivative Pricing: From Statistical Physics to Risk Management*, 2 edn, Cambridge University Press.
- Burnham, K. P. & Anderson, D. R. (2004), *Model Selection and Multimodel Inference*, Springer, New York.
URL: <https://doi.org/10.1007/b97636>
- Calvino, F., Criscuolo, C., Menon, C. & Secchi, A. (2018), ‘Growth volatility and size: A firm-level study’, *Journal of Economic Dynamics and Control* **90**, 390–407.
- Carrasco, M., Florens, J.-P. & Renault, E. (2007), ‘Linear inverse problems in structural econometrics estimation based on spectral decomposition and regularization’, *Handbook of Econometrics* **6**, 5633–5751.
- Cette, G., Corde, S. & Lecat, R. (2018), ‘Firm-level productivity dispersion and convergence’, *Economics Letters* **166**, 76–78.
- Cunningham, C., Foster, L., Grim, C., Haltiwanger, J., Pabilonia, S. W., Stewart, J. & Wolf, Z. (2018), Dispersion in dispersion: Measuring establishment-level differences in productivity, Working Papers 18-25, Center for Economic Studies, U.S. Census Bureau.
- Embrechts, P., Kluppelberg, C. & Mikosch, T. (1997), Time series analysis for heavy-tailed processes, *in* ‘Modelling Extremal Events,’ Springer, Berlin, Heidelberg, pp. 371–412.
- Fama, E. F. (1965), ‘Portfolio analysis in a stable Paretian market’, *Management Science* **11**(3), 404–419.
- Fama, E. F. & Roll, R. (1971), ‘Parameter estimates for symmetric stable distributions’, *Journal of the American Statistical Association* **66**(334), 331–338.
- Foster, L., Grim, C., Haltiwanger, J. C. & Wolf, Z. (2018), Innovation, productivity dispersion, and productivity growth, Working Paper 24420, National Bureau of Economic Research.

- Gabaix, X. (2011), ‘The granular origins of aggregate fluctuations’, *Econometrica* **79**(3), 733–772.
- Gaffeo, E. (2008), ‘Lévy-stable productivity shocks’, *Macroeconomic Dynamics* **12**(3), 425–443.
- Gaffeo, E. (2011), ‘The distribution of sectoral TFP growth rates: International evidence’, *Economics Letters* **113**(3), 252–255.
- Ghiglino, C. (2012), ‘Random walk to innovation: Why productivity follows a power law’, *Journal of Economic Theory* **147**(2), 713–737.
- Gnedenko, B. & Kolmogorov, A. (1954), *Limit distributions for sums of independent random variables*, Addison-Wesley, Cambridge.
- Goldin, I., Koutroumpis, P., Lafond, F., Rochowicz, N. & Winkler, J. (2018), Why is productivity slowing down?, Technical report, Working Paper, Oxford Martin School, University of Oxford.
- Haldane, A. (2017), ‘Productivity puzzles’, *Speech given by Andy Haldane, Chief Economist, Bank of England at the London School of Economics* .
- Hall, R. E. & Jones, C. I. (1999), ‘Why do some countries produce so much more output per worker than others?’, *The quarterly journal of economics* **114**(1), 83–116.
- Hansen, L. P. (1982), ‘Large sample properties of generalized method of moments estimators’, *Econometrica* **50**(4), 1029–1054.
- Haskel, J. & Westlake, S. (2018), *Capitalism without capital: the rise of the intangible economy*, Princeton University Press.
- Holly, S., Petrella, I. & Santoro, E. (2013), ‘Aggregate fluctuations and the cross-sectional dynamics of firm growth’, *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **176**(2), 459–479.
- Hsieh, C.-T. & Klenow, P. J. (2009), ‘Misallocation and manufacturing TFP in China and India’, *The Quarterly Journal of Economics* **124**(4), 1403–1448.
- Hsieh, C.-T. & Klenow, P. J. (2017), ‘The reallocation myth’, *Fostering a Dynamic Global Economy* .
- Ijiri, Y. & Simon, H. A. (1977), *Skew distributions and the sizes of business firms*, Vol. 24, North Holland, Amsterdam.
- Jäger, K. & The Conference Board (2018), EU KLEMS growth and productivity accounts 2017 release, statistical module 1, Technical report, EU KLEMS Project. Revised 2018, Available online http://www.euklems.net/TCB/2018/Methodology_EUKLEMS_2017_revised.pdf, http://www.euklems.net/TCB/2018/EUKLEMS_2018_revision.pdf.
- Kalemli-Ozcan, S., Sorensen, B., Villegas-Sanchez, C., Volosovych, V. & Yesiltas, S. (2015), How to construct nationally representative firm level data from the orbis global database, Working Paper 21558, National Bureau of Economic Research.
- Kharrat, T. & Boshnakov, G. N. (2016), *StableEstim: Estimate the Four Parameters of Stable Laws using Different Methods*. R package version 2.1.
URL: <https://CRAN.R-project.org/package=StableEstim>

- König, M. D., Lorenz, J. & Zilibotti, F. (2016), ‘Innovation vs. imitation and the evolution of productivity distributions’, *Theoretical Economics* **11**(3), 1053–1102.
- Lucas, R. E. & Moll, B. (2014), ‘Knowledge growth and the allocation of time’, *Journal of Political Economy* **122**(1), 1–51.
- Mandelbrot, B. (1960), ‘The Pareto-Lévy law and the distribution of income’, *International Economic Review* **1**(2), 79–106.
- Mandelbrot, B. (1963), ‘New methods in statistical economics’, *Journal of Political Economy* **71**(5), 421–440.
- Mankiw, N. G., Romer, D. & Weil, D. N. (1992), ‘A contribution to the empirics of economic growth’, *The quarterly journal of economics* **107**(2), 407–437.
- McCulloch, J. H. (1986), ‘Simple consistent estimators of stable distribution parameters’, *Communications in Statistics-Simulation and Computation* **15**(4), 1109–1136.
- Mitchell, W. C. (1915), ‘The making and using of index numbers’, *US BLS* .
- Newman, M. E. (2005), ‘Power laws, Pareto distributions and Zipf’s law’, *Contemporary Physics* **46**(5), 323–351.
- Nishida, M., Petrin, A., Rotemberg, M. & White, T. (2016), ‘Measuring cross-country differences in misallocation’, *US Census Bureau Center for Economic Studies Paper No. CES-WP-16-50* .
- Nolan, J. P. (1998), ‘Parameterizations and modes of stable distributions’, *Statistics & Probability Letters* **38**(2), 187 – 195.
- Nolan, J. P. (2019), ‘Stable distributions - models for heavy tailed data.’, *Unfinished manuscript* .
URL: academic2.american.edu/~jpnolan.
- Perla, J. & Tonetti, C. (2014), ‘Equilibrium imitation and growth’, *Journal of Political Economy* **122**(1), 52–76.
- Royall, R. (2017), *Statistical Evidence*, Routledge.
URL: <https://doi.org/10.1201/9780203738665>
- Samuelson, P. A. (1967), ‘Efficient portfolio selection for Pareto-Lévy investments’, *Journal of Financial and Quantitative Analysis* **2**(2), 107–122.
- Schwarzkopf, Y., Axtell, R. & Farmer, J. D. (2010), ‘An explanation of universality in growth fluctuations’, *An Explanation of Universality in Growth Fluctuations (April 28, 2010)* .
- Soofi, E. S., Ebrahimi, N. & Habibullah, M. (1995), ‘Information distinguishability with application to analysis of failure data’, *Journal of the American Statistical Association* **90**(430), 657–668.
- Sornette, D. (2006), *Critical phenomena in natural sciences: chaos, fractals, selforganization and disorder: concepts and tools*, Springer Science & Business Media.
- Stanley, M., Amaral, L., Buldyrev, S., Havlin, S., Leschhorn, H., Maass, P., Salinger, M. & Stanley, H. (1996), ‘Scaling behaviour in the growth of companies’, *Nature* **379**(6568), 804–806.

- Trapani, L. (2016), 'Testing for (in)finite moments', *Journal of Econometrics* **191**(1), 57 – 68.
- Uchaikin, V. V. & Zolotarev, V. M. (2011), *Chance and stability: stable distributions and their applications*, Walter de Gruyter.
- Yang, J. (2018), 'A quantal response statistical equilibrium model of induced technical change in an interactive factor market: Firm-level evidence in the EU economies', *Entropy* **20**(3).

A Data appendix

This appendix gives a more detailed description of the data.

A.1 Raw data

We use the Orbis Europe firm level database, which is part of the Orbis data provided by Bureau van Dijk. The database encompasses more than 21 million firms of all sizes from 44 different European countries. Around 9 million firms in 15 countries are used in the present analysis; see Table 1.

The variables used are listed in Table 4. The raw data is indexed by the firms' Identification Number (IDNR), the year of reporting (CLOSDATE_year), the size of the firm (COMCAT: small, medium, large, very large), and their four-digit industrial classification under the NACE Rev. 2 classification system (NACE_PRIM_CODE). The key balance sheet items of interest are: employment (EMPL), total assets (TOAS), wages (STAF), depreciation (DEPR), and earnings after depreciation (EBIT).

Table 4: Variables from Orbis/Orbis Europe used for the analysis

Code	Variable	Description
IDNR		Firm's identification number
COMCAT		Firm's size category (small, medium, large, very large)
EBIT	P	Earnings after depreciation (earnings before interest and tax)
STAF	W	Wages (staff costs)
EMPL	L	Employment
DEPR	δ	Depreciation
TOAS	K_A	Capital (Total assets)
NACE_PRIM_CODE		Industrial classification code (NACE Rev. 2)
CLOSDATE_year		Year
	Y	Value-added ($W + P$)
	LP	Labor productivity (Y/L)
	ΔLP	Labor productivity change ($LP_t - LP_{t-1}$)

A.2 Data cleaning

The first step in processing the data is to ensure that no observations have missing values for their reporting IDNR or year. Only unconsolidated and unaccompanied¹⁵ accounts are kept to avoid double counting. Since some accounts provide a default year for various reasons, observations where the reported year is not between 2006 and 2015 are removed. Negative total asset, fixed asset, sales, wages, and employment observations are regarded as missing.

Furthermore, any firms for which employment or value-added are recorded as missing are removed altogether. In addition, country-year samples are required to have at least 10,000 observations. This reduces the list of countries considered from 44 to 23 countries which have at least one year with a sample of at least 10,000 firms. Following this procedure, only countries with over five years of data are used, reducing the sample to 15 countries, see Table 1.¹⁶

¹⁵For some groups that consist of multiple firms, both consolidated and unconsolidated data is provided, i.e. data is listed for the entire group and again for each of the firms.

¹⁶For 29 countries, there was insufficient data: Albania, Austria, Belarus, Belgium, Bosnia and Herzegovina, Croatia, Cyprus, Denmark, Greece, Iceland, Ireland, Kosovo, Latvia, Liechtenstein, Lithuania, Luxembourg, Malta, Monaco, Montenegro, Netherlands, Norway, North Macedonia, Poland, Moldova, Russian Federation, Serbia, Switzerland, Turkey, Ukraine.

A.3 Deflation

The KLEMS database¹⁷ provides the most comprehensive data on deflation for the countries and industries covered by the Orbis sample. In particular, tables `ALL_output_17ii.txt` and `ALL_capital_17i.txt` provide value-added, gross output, and capital deflators for two-digit NACE Rev.2 industries for the countries in the Orbis sample. Depending on its location, industry, and year of reporting, the firms' wages, depreciation, and profits (EBIT) are deflated using the value-added deflator, their total assets are deflated using the capital deflator.

A.4 Main variables of interest

We investigate the distribution of productivity (LP) and of its absolute change over time (ΔLP). See Table 4.

Labor productivity Value-added is defined as the difference between *gross output* and *intermediate inputs*. However, it is impractical to compute it directly from these variables in our case, since intermediate inputs are rarely reported in the Orbis Europe dataset. This is especially true for small firms. Consequently, we impute value-added instead as the sum of profits and wages paid by the firm. Labor productivity is then computed as the ratio of value-added to employment (see Section 2 and Table 4).

Labor productivity change An obvious choice for tracking intertemporal change would be growth, in this case labor productivity growth $(LP_t - LP_{t-1})/LP_{t-1}$. However, this measure has a singularity at $LP_{t-1} = 0$ which makes the scale inconsistent and will, if LP_{t-1} can be negative, result in undefined values and counterintuitive behavior for negative values. Value-added may in general be close to zero and volatile with significant variation between periods which makes labor productivity growth a problematic measure of intertemporal change. In our case the measure is particularly impractical, as it is imputed as the sum of profits and wages that may be negative. Another alternative, log returns, are subject to the same problems. As a consequence, we use labor productivity change $LP_t - LP_{t-1}$ instead.

A.5 Dimensions for subdivision of the data set

To examine the labor productivity distributions for additional structure, we subdivided the observations by country and by year, but also by firm size category, and industry.

Size The challenge in describing the size of firms is the considerable heterogeneity in their use of inputs as well as their outputs. Self-employed individuals may produce higher revenues than local businesses, large stocks of assets may be kept in separate legal entities that do not generate income on paper. In order to provide a more well-rounded definition of firm sizes, Orbis classifies them as follows:

- Very Large (VL): for a firm to be considered as VL, they have to meet **at least one** of the following conditions:
 - Operating Revenue greater than or equal to 100 million EUR,
 - Total assets greater than or equal to 200 million EUR,
 - Employee count greater than or equal to 1000,

¹⁷<http://www.euklems.net>

- Listed as a publicly traded firm.

It should be noted that firms for which the ratio of operating revenue to employee count **or** the ratio of employee count to total assets is less than 100 are excluded from the VL firms.

- Large (L): for a firm to be considered as L, they have to meet **at least one** of the following conditions:
 - Operating revenue greater than or equal to 10 million EUR,
 - Total assets greater than or equal to 20 million EUR,
 - Employee count greater than or equal to 150.

Again, firms for which the ratio of operating revenue to employee count **or** the ratio of employee count to total assets is less than 100 are not counted as L.

- Medium (M): for a firm to be considered as M, they have to meet **at least one** of the following conditions:
 - Operating revenue greater than or equal to 1 million EUR,
 - Total assets greater than or equal to 2 million EUR,
 - Employee count greater than or equal to 15.

Firms for which the ratio of operating revenue to employee count **or** the ratio of employee count to total assets is less than 100 are also excluded.

- Small (S): firms of size S make up the body of firms that do not fit in either of the three categories above.

Industry Besides determining differences in dispersion over time and firm size, another useful dimension to consider is the dispersion of productivity by industry. Each firm reported in Orbis includes its four-digit NACE Rev. 2 industrial classification. We further aggregate the industry classes in two different ways. First, we use the highest aggregation level in the original NACE Rev. 2 industrial classification. This yields around 19 industries. Second, we further aggregate 19 industries into 8 larger industries, which we will use as the main classification in the paper. Table 5 summarizes the industry category and descriptions.

Table 5: Industry Category Codes and Descriptions

Broad Category	Code	Original NACE Rev. 2 Category
Agriculture (Agr)	A	Agriculture, forestry and fishing (Agr)
	B	Mining and quarrying (Mine)
Manufacturing (Manu)	C	Manufacturing (Manu)
Energy (Energy)	D	Electricity, gas, steam and air conditioning supply (Elec)
	E	Water supply; sewerage, waste management and remediation activities (Water)
Construction (Cons)	F	Construction (Cons)
Non-Financial Service (NF-Serv)	G	Wholesale and retail trade; repair of motor vehicles and motorcycles (Whole)
	H	Transportation and storage (Trans)
	I	Accommodation and food service activities (Accom)
	N	Administrative and support service activities (Adm-S)
	R	Arts, entertainment and recreation (Art)
Information (Info)	S	Other service activities (O-Serv)
	J	Information and communication (Info)
	Finance, Insurance, Real Estate, (FIRE)	K
L		Real estate activities (Real)
M		Professional, scientific and technical activities (Prof-S)
Non-Market Service (NM-Serv)	O	Public administration and defence; compulsory social security (Pub-S)
	P	Education (Edu)
	Q	Human health and social work activities (Health)

A.6 Composition of the data sets by country, year, firm size and industry after cleaning.

The composition of the data set by firm sizes (Figure 12) and industry sectors (Figure 13) remains mostly constant for all countries, especially for the five largest countries in the data set (Spain, France, Italy, Germany, and the United Kingdom) for which results are reported in the main text.

In almost all countries, small firms account for the largest share with progressively smaller percentages of medium sized, large and very large firms. Exceptions are Germany and the United Kingdom where only a small number of small firms is represented. The share of self-employed individuals varies a bit more.

Manufacturing, which is of particular economic interest, is in most countries among the most well-represented sectors. The wholesale and retail, construction, and professional services sectors represent other large parts of the data, which reflects the actual business demographics.¹⁸

¹⁸For the business demographics by sector in Europe, see Eurostat (<https://ec.europa.eu/eurostat/data/database>), "Employer business demography by NACE Rev. 2 and NUTS 3 regions" (bd_enace2_r3) database.

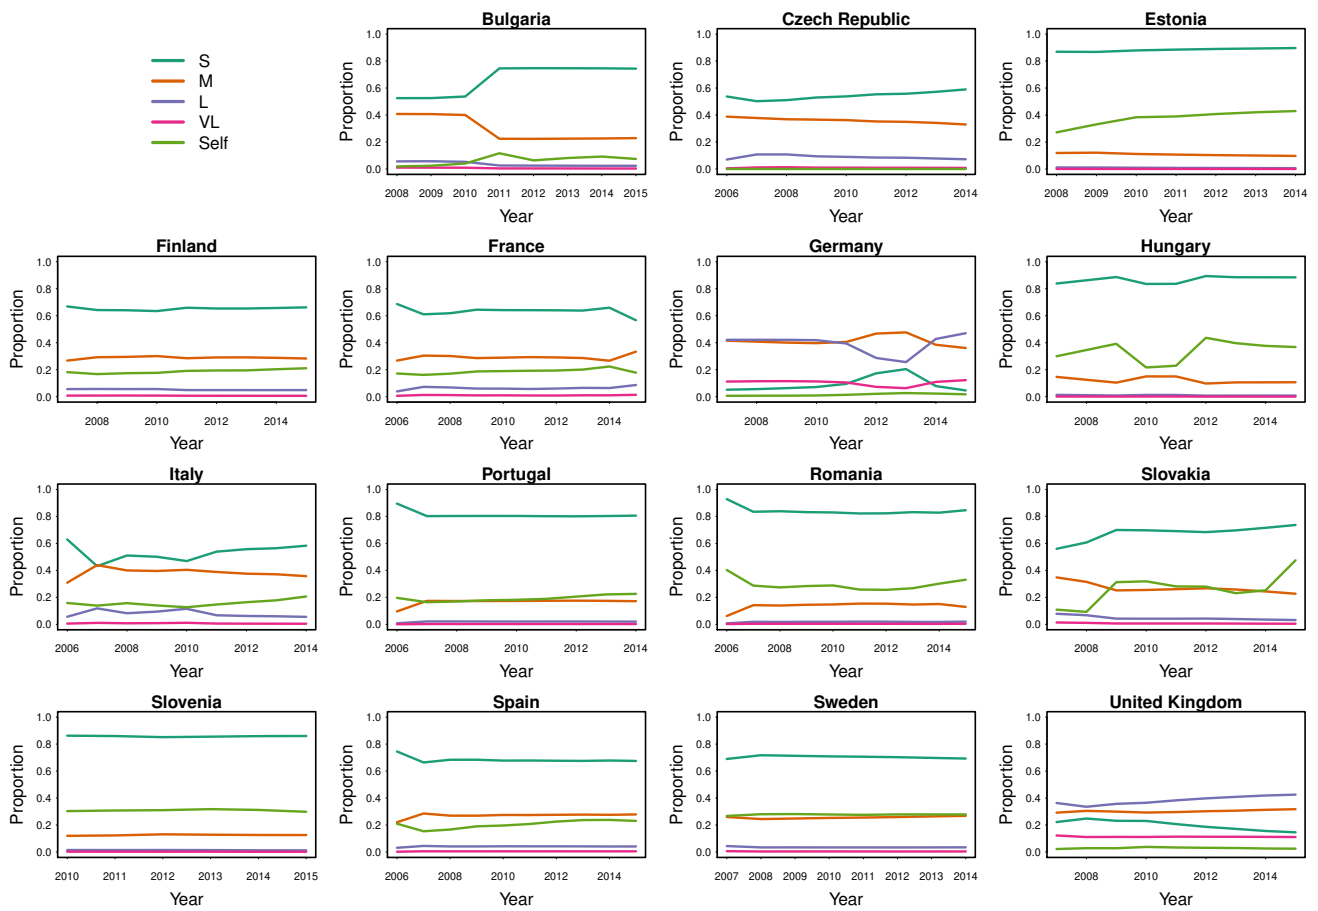


Figure 12: Composition of the data sets by country and firm size after cleaning.

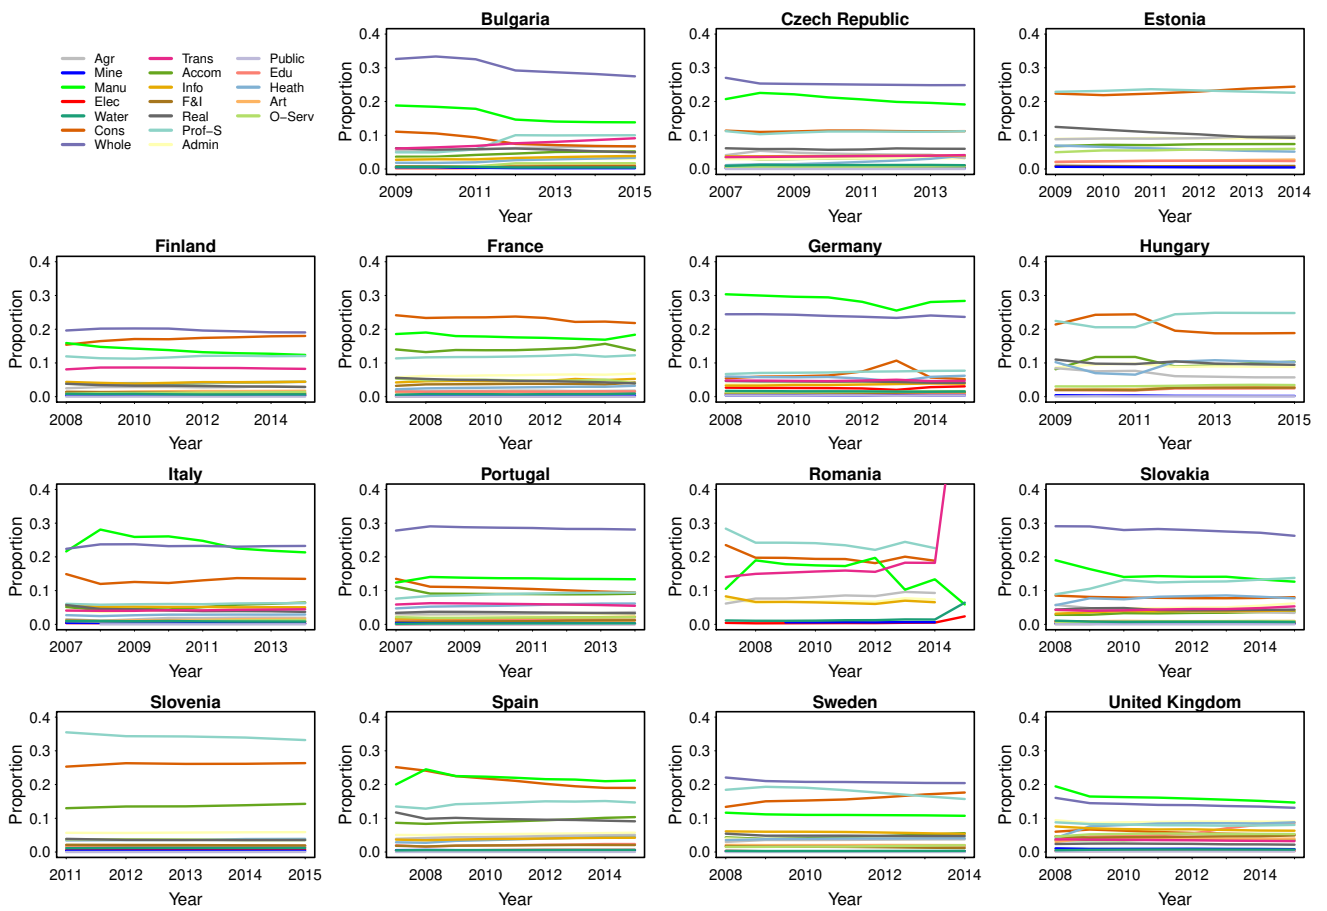


Figure 13: Composition of the data sets by country and firm size

B Scaling of the sample standard deviation with sample size

Here we provide a highly stylized, heuristic derivation for the scaling of the sample standard deviation with sample size in the case of a Lévy alpha-stable distributed random variable. The key to this phenomenon is that because the theoretical moment is infinite, the larger the sample size, the higher is the chance that an extreme event is drawn. These extreme events are so extreme that they dominate the sum of squares from which the variance is computed. Thus, the larger the sample size, the larger the sample variance.

Sornette (2006) and Bouchaud & Potters (2003), for instance, provide more precise statements. Here we expose the argument in the simplest, albeit non rigorous way. We discuss the maximum, but symmetric arguments apply to the minimum.

First, we note that for N large enough, the sample maximum will be dictated by the tail. A key characteristic of the Lévy alpha-stable distribution is that it has power law tails, that is, for large x (Nolan 2019),

$$P(X > x) \sim x^{-\alpha}. \quad (9)$$

Now, in a sample of size N , we would hardly expect to see an extreme value that has chances of occurring less than $1/N$. Thus, we may define the “typical” value of the maximum as the value X_{\max} such that $1/N = P(X > X_{\max})$. Using Eq. 9, we have¹⁹ $1/N \sim X_{\max}^{-\alpha}$, and solving for X_{\max} gives

$$X_{\max} \sim N^{\frac{1}{\alpha}}, \quad (10)$$

For simplicity, let us assume a mean of zero, so that the sample variance is just the sum of squares,

$$\text{Var}(X) \sim \frac{1}{N} \sum_{i=1}^N X_i^2. \quad (11)$$

In a Lévy alpha-stable distribution, the square of the maximum (or minimum, if larger in absolute value) is so large that it dominates the entire sum of squares, such that we may approximate

$$\sum_{i=1}^N X_i^2 \approx X_{\max}^2 \sim N^{\frac{2}{\alpha}}, \quad (12)$$

where the last step uses Eq. 10. Now, inserting Eq. 12 into 11, and taking square root, we find that the standard deviation depends on the sample size as

$$\sqrt{\text{Var}(X)} \sim N^{\frac{1}{\alpha} - \frac{1}{2}}.$$

See the derivations leading to (4.52) in Sornette (2006) for a more rigorous argument.

Note that when $\alpha = 2$, so that the distribution is Gaussian, the sample standard deviation does *not* increase with sample size, as one expects.

¹⁹Sornette (2006) shows that this is the value of the maximum that is not exceeded with probability $1/e \approx 37\%$. Generalizing to the value of the maximum that is not exceeded with probability p implies the condition $\frac{\ln(1/p)}{N} = P(X > X_{\max})$, which does not change the scaling of X_{\max} with N . Newman (2005) shows that the expected value of the maximum also scale as $N^{\frac{1}{\alpha}}$.

C Trapani's (2016) procedure on testing for finite moments

Trapani (2016) suggests a test for the finiteness of arbitrary moments of order p , including fractional (non-integer) moments. As the moment of order p may be infinite, it is unknown whether a limiting distribution of the moments exists. The test can therefore not be applied directly. Following a randomized testing approach, the test therefore adds artificial randomness to manipulate the quantity in question, a moment of order p , to yield a known distribution if the moment is infinite. The resulting test statistic can then be compared to this distribution to obtain a p value for whether or not the null hypothesis, that the moment of order p is infinite, is correct.

The test statistic is derived such that under the null hypothesis H_0 the p th moment of the distribution a sample follows is infinite. For this, the approach starts with the absolute moment with order p over the sample

$$A_p = \frac{1}{n} \sum_{k=1}^n |X_k|^p.$$

To make the resulting test statistic scale-invariant and comparable, the absolute moment must be rescaled,

$$A_p^* = \frac{A_p}{A_\psi^{p/\psi}} \times \frac{(A_\psi^N)^{p/\psi}}{A_p^N},$$

with $\psi \in (0, p)$. A_p^N denotes the p th absolute fractional lower order moment of the standard normal distribution $N(0, 1)$. Next, an artificial random sample ξ of size r^{20} is generated from a standard normal distribution and rescaled,

$$\varphi_r = \sqrt{e^{A_p^*}} \times \xi_r.$$

The intuition here is that φ_r follows a normal distribution with mean zero and a finite variance, as $n \rightarrow \infty$, if A_p^* is finite itself. Thus the problem has been reduced from testing for any moment p , to testing the existence of the variance of the transformed random variable φ_r .

The next step is to generate a sequence ζ_r , given by:

$$\zeta_r(u) \equiv I[\varphi_r \leq u],$$

where $I[\cdot]$ is an indicator function, and $u \neq 0$ is any real number. Under H_0 , $\zeta_r(u)$ will have a Bernoulli distribution with mean $\frac{1}{2}$ and variance $\frac{1}{4}$. This is not the case under the alternative, where $A_p^* < \infty$, as $e^{A_p^*}$ converges to a finite value.

Values for u are picked from some density, but for simplicity it can be taken from a uniform distribution $U(-1, 1)$. As a result, the test statistic of interest is obtained as follows:

²⁰Trapani suggests $r = n^{0.8}$ where n is the number of observations used to compute A_p .

$$\Theta_{nr} \equiv \int_{-1}^1 \frac{1}{2} \vartheta_{nr}^2(u) du,$$

$$\vartheta_{nr}(u) \equiv \frac{2}{\sqrt{r}} \sum_{j=1}^r \left[\zeta_j(u) - \frac{1}{2} \right].$$

Under the null hypothesis that moment p is infinite, $\vartheta_{nr}(u)$ should reduce to zero given that $\zeta_j(u)$ has a mean of two. Θ_{nr} is thus shown to follow a χ^2 distribution with $df = 1$ if moment p is infinite.

The following table shows the test statistics and χ^2 p-values of the infinite second moment test for labor productivity level and change in the five countries in our data set with the largest sample sizes as reported throughout the paper.

The test fails to reject the null hypothesis that the second moment is infinite for any reasonable level of confidence, for almost all country-year sub-samples. The only sub-samples for which the test does not reject the null hypothesis of finite variance are France 2008-10 and Italy 2009-10. Given the excellent other goodness of fit measures presented in Section 3, however, we suspect that this is a data-collection issue.²¹

²¹The test relies on the presence of extreme values to drive the divergence of moments. These may not have been sampled in those years due to various causes or the firms themselves may during the years of the crisis 2008-2010 have employed accounting techniques to avoid extreme values.

C.1 Finding finite moments for firm-level labour productivity measures

Table 6: Testing for Infinite Moments of Firm-level Labour Productivity

Moment	Year	Labour Productivity: LP					Change in Labour Productivity: ΔLP				
		France	Germany	Italy	Spain	UK	France	Germany	Italy	Spain	UK
First: $E(X)$	2006	288.20 (0.00)	- (-)	117.43 (0.00)	251.15 (0.00)	- (-)	- (-)	- (-)	- (-)	- (-)	- (-)
	2007	433.11 (0.00)	80.68 (0.00)	539.81 (0.00)	516.67 (0.00)	63.68 (0.00)	151.54 (0.00)	- (-)	54.06 (0.00)	91.72 (0.00)	- (-)
	2008	470.29 (0.00)	112.15 (0.00)	477.13 (0.00)	561.72 (0.00)	133.62 (0.00)	234.35 (0.00)	0.93 (0.33)	23.84 (0.00)	178.35 (0.00)	20.25 (0.00)
	2009	492.15 (0.00)	127.26 (0.00)	670.87 (0.00)	521.40 (0.00)	118.88 (0.00)	270.48 (0.00)	21.03 (0.00)	62.32 (0.00)	219.09 (0.00)	35.53 (0.00)
	2010	531.25 (0.00)	139.44 (0.00)	572.19 (0.00)	617.44 (0.00)	148.46 (0.00)	285.67 (0.00)	41.93 (0.00)	317.86 (0.00)	266.11 (0.00)	56.73 (0.00)
	2011	520.99 (0.00)	148.63 (0.00)	1022.04 (0.00)	587.63 (0.00)	134.31 (0.00)	279.47 (0.00)	34.25 (0.00)	536.84 (0.00)	149.99 (0.00)	32.65 (0.00)
	2012	453.27 (0.00)	197.96 (0.00)	1085.51 (0.00)	525.25 (0.00)	117.56 (0.00)	251.47 (0.00)	44.62 (0.00)	625.89 (0.00)	181.12 (0.00)	6.95 (0.01)
	2013	505.12 (0.00)	226.12 (0.00)	1136.01 (0.00)	509.91 (0.00)	181.39 (0.00)	270.88 (0.00)	63.10 (0.00)	628.98 (0.00)	119.49 (0.00)	9.83 (0.00)
	2014	598.26 (0.00)	61.10 (0.00)	1210.42 (0.00)	616.46 (0.00)	185.10 (0.00)	317.65 (0.00)	26.41 (0.00)	677.92 (0.00)	122.73 (0.00)	22.09 (0.00)
	2015	478.71 (0.00)	45.99 (0.00)	- (-)	632.09 (0.00)	20.24 (0.00)	253.36 (0.00)	7.43 (0.01)	- (-)	189.83 (0.00)	7.09 (0.01)
Second: $E(X^2)$	2006	0.02 (0.88)	- (-)	0.00 (0.96)	0.01 (0.94)	- (-)	- (-)	- (-)	- (-)	- (-)	- (-)
	2007	0.01 (0.94)	0.00 (0.96)	0.00 (0.97)	0.00 (0.98)	0.00 (0.97)	0.00 (0.97)	- (-)	0.00 (0.96)	0.00 (0.97)	- (-)
	2008	2.62 (0.11)	0.00 (0.97)	0.00 (1.00)	0.00 (0.98)	0.00 (0.97)	0.00 (0.98)	0.00 (1.00)	0.00 (1.00)	0.00 (0.98)	0.00 (0.96)
	2009	14.82 (0.00)	0.00 (0.97)	2.58 (0.11)	0.00 (1.00)	0.00 (0.97)	0.00 (0.98)	0.00 (0.96)	0.00 (1.00)	0.00 (0.98)	0.00 (0.97)
	2010	14.93 (0.00)	0.00 (0.97)	0.29 (0.59)	0.00 (0.98)	0.00 (0.97)	0.00 (0.98)	0.00 (0.97)	0.00 (0.98)	0.00 (0.98)	0.00 (0.97)
	2011	0.44 (0.51)	0.00 (0.97)	0.00 (0.96)	0.00 (0.98)	0.00 (0.97)	0.00 (0.98)	0.00 (0.97)	0.00 (0.99)	0.00 (1.00)	0.00 (0.97)
	2012	0.06 (0.80)	0.00 (0.97)	0.00 (0.99)	0.00 (0.99)	0.00 (0.97)	0.00 (0.98)	0.00 (0.97)	0.00 (0.99)	0.00 (0.98)	0.00 (1.00)
	2013	0.11 (0.74)	0.00 (0.97)	0.00 (0.99)	0.00 (1.00)	0.00 (0.97)	0.00 (0.98)	0.00 (0.97)	0.00 (0.99)	0.00 (1.00)	0.00 (1.00)
	2014	0.00 (0.98)	0.00 (1.00)	0.10 (0.75)	0.00 (0.99)	0.00 (0.98)	0.00 (0.98)	0.00 (0.97)	0.00 (0.99)	0.00 (1.00)	0.00 (0.97)
	2015	0.01 (0.92)	0.00 (1.00)	- (-)	0.00 (0.99)	0.00 (1.00)	0.00 (0.98)	0.00 (0.97)	- (-)	0.00 (0.99)	0.00 (1.00)

Table 7: Testing for Infinite Moments of Firm-level Log Labour Productivity

Moment	Year	Log Labour Productivity: $\log LP$					Change in Log Labour Productivity: $\Delta \log LP$				
		France	Germany	Italy	Spain	UK	France	Germany	Italy	Spain	UK
First: $E(X)$	2006	332.77 (0.00)	- (-)	152.30 (0.00)	0.00 (1.00)	- (-)	- (-)	- (-)	- (-)	- (-)	- (-)
	2007	519.70 (0.00)	171.04 (0.00)	661.24 (0.00)	0.00 (1.00)	0.00 (1.00)	164.08 (0.00)	- (-)	78.56 (0.00)	179.80 (0.00)	- (-)
	2008	552.97 (0.00)	180.25 (0.00)	0.00 (1.00)	0.00 (1.00)	0.00 (1.00)	250.57 (0.00)	86.08 (0.00)	328.57 (0.00)	0.00 (1.00)	89.61 (0.00)
	2009	572.32 (0.00)	186.26 (0.00)	0.00 (1.00)	0.00 (1.00)	0.00 (1.00)	284.56 (0.00)	81.52 (0.00)	350.73 (0.00)	462.42 (0.00)	147.80 (0.00)
	2010	618.84 (0.00)	194.79 (0.00)	709.17 (0.00)	0.00 (1.00)	0.00 (1.00)	332.73 (0.00)	116.72 (0.00)	431.22 (0.00)	524.18 (0.00)	194.44 (0.00)
	2011	607.98 (0.00)	214.34 (0.00)	0.00 (1.00)	0.00 (1.00)	0.00 (1.00)	330.19 (0.00)	119.18 (0.00)	580.91 (0.00)	514.27 (0.00)	183.35 (0.00)
	2012	533.20 (0.00)	295.56 (0.00)	0.00 (1.00)	0.00 (1.00)	0.00 (1.00)	268.65 (0.00)	124.59 (0.00)	705.60 (0.00)	0.00 (1.00)	172.01 (0.00)
	2013	602.20 (0.00)	0.00 (1.00)	0.00 (1.00)	0.00 (1.00)	327.47 (0.00)	301.68 (0.00)	155.13 (0.00)	829.28 (0.00)	0.00 (1.00)	178.24 (0.00)
	2014	724.55 (0.00)	226.50 (0.00)	0.00 (1.00)	0.00 (1.00)	0.00 (1.00)	367.73 (0.00)	128.17 (0.00)	872.81 (0.00)	0.00 (1.00)	206.72 (0.00)
	2015	583.16 (0.00)	206.79 (0.00)	- (-)	0.00 (1.00)	0.00 (1.00)	305.42 (0.00)	120.14 (0.00)	- (-)	0.00 (1.00)	0.00 (1.00)
Second: $E(X^2)$	2006	402.40 (0.00)	- (-)	184.04 (0.00)	0.00 (1.00)	- (-)	- (-)	- (-)	- (-)	- (-)	- (-)
	2007	628.41 (0.00)	206.73 (0.00)	799.39 (0.00)	0.00 (1.00)	0.00 (1.00)	196.81 (0.00)	- (-)	94.22 (0.00)	215.06 (0.00)	- (-)
	2008	668.64 (0.00)	217.87 (0.00)	0.00 (1.00)	0.00 (1.00)	0.00 (1.00)	300.52 (0.00)	102.89 (0.00)	393.94 (0.00)	0.00 (1.00)	106.89 (0.00)
	2009	691.97 (0.00)	225.14 (0.00)	0.00 (1.00)	0.00 (1.00)	0.00 (1.00)	341.35 (0.00)	97.36 (0.00)	420.30 (0.00)	553.11 (0.00)	176.20 (0.00)
	2010	748.27 (0.00)	235.45 (0.00)	856.90 (0.00)	0.00 (1.00)	0.00 (1.00)	399.26 (0.00)	139.78 (0.00)	517.06 (0.00)	627.65 (0.00)	232.69 (0.00)
	2011	735.16 (0.00)	259.07 (0.00)	0.00 (1.00)	0.00 (1.00)	0.00 (1.00)	396.16 (0.00)	142.62 (0.00)	697.16 (0.00)	615.56 (0.00)	219.10 (0.00)
	2012	644.66 (0.00)	357.19 (0.00)	0.00 (1.00)	0.00 (1.00)	0.00 (1.00)	322.12 (0.00)	148.93 (0.00)	845.90 (0.00)	0.00 (1.00)	205.39 (0.00)
	2013	728.00 (0.00)	0.00 (1.00)	0.00 (1.00)	0.00 (1.00)	395.32 (0.00)	361.85 (0.00)	185.38 (0.00)	994.32 (0.00)	0.00 (1.00)	212.73 (0.00)
	2014	875.84 (0.00)	273.64 (0.00)	0.00 (1.00)	0.00 (1.00)	0.00 (1.00)	441.06 (0.00)	153.04 (0.00)	1046.68 (0.00)	0.00 (1.00)	246.92 (0.00)
	2015	704.99 (0.00)	249.86 (0.00)	- (-)	0.00 (1.00)	0.00 (1.00)	366.16 (0.00)	143.30 (0.00)	- (-)	0.00 (1.00)	0.00 (1.00)

D Fitting procedure

Two distributional models are considered in this paper: The Lévy alpha-stable distribution and the 4-parameter asymmetric Subbotin (asymmetric exponential power, AEP) distribution as a reference model. To fit the AEP model, we use the L-moments method Asquith (2014) as implemented in the R-package `lmomco`. A 4-parameter variant was chosen to allow a direct comparison with the Lévy alpha-stable model,

which has four parameters.²² For fitting the Lévy alpha-stable model, we use McCulloch’s quantile method McCulloch (1986) as explained in Section 3.1 in the main text and as implemented in the R package `StableEstim`.

D.1 Comparison of the McCulloch quantile method and the maximum likelihood method

Fitting the parameters of the Lévy alpha-stable distribution using maximum likelihood estimation would be desirable for accuracy, but has prohibitive computation power requirements. Therefore, McCulloch’s computationally much cheaper quantile method was used. To avoid unacceptable deviations from what the accurate parameter fit would be, we only consider subsamples with large numbers of observations (typically $\geq 10,000$; $\geq 1,000$ and $5,000$ respectively for country-industry and country-firm size subsamples).

We conducted a Monte Carlo simulation to determine the expected errors of McCulloch quantile estimation and maximum likelihood estimation fits depending on sample size N . We fitted artificial samples drawn from Lévy alpha-stable distributions with parameters that would be realistic for the distributions considered in the present paper ($\alpha = 1.2$, $\beta = 0.55$, $\gamma = 1.0$, $\delta = 0.0$). The results are shown in Figure 14. The actual parameter value is in red, while the results for MLE and QT are shown in green and blue. The bold lines represent the mean and shaded area represent the error bar from the 95% confidence interval of estimates. It is shown that the mean value and the error bar of the QT estimation get smaller and approaches those of MLE, suggesting that the accuracy of McCulloch quantile estimation and maximum likelihood estimation fits are similar for all parameters for the sample sizes relevant in our case.

D.2 Comparison of the McCulloch quantile method and the generalized method of moments (GMM)

Another computationally less expensive alternative is the generalized method of moments (GMM). To assess the relative performance of the McCulloch quantile method and GMM, we compare the Soofi ID scores of our McCulloch quantile method fits of all subsamples (by country and year, by country and firm size, by country and industry sector) for all 15 countries²³ with those of a GMM fit of the same subsamples using Hansen’s (1982) two-step algorithm with spectral cut-off regularization scheme (Carrasco et al. 2007). If a method produces high quality fits, the Soofi ID scores would be concentrated at or very close to score 100. So plotting the density of Soofi ID scores of the fits obtained with both fitting methods illustrates the quality of the methods. The densities, shown in Figure 15, demonstrate that McCulloch quantile method consistently performs better across all types of subsamples (by country and industry, by country and year, and by country and firm size).

E Comparison of the fit: SOOFI ID, AIC, and K-fold cross validation

For the goodness of fit and model comparison measure, we will mainly use Soofi’s (1995) information distinguishability (ID), along with the Akaike (1973) information criterion (AIC) and K-fold cross-validation.

²²Another possibility would be to use more powerful variant of the AEP (the 5-parameter one) and to penalize the model with greater number of parameters.

²³These are the fits given in the main text for five countries in Figures 6, 9, 28, 29

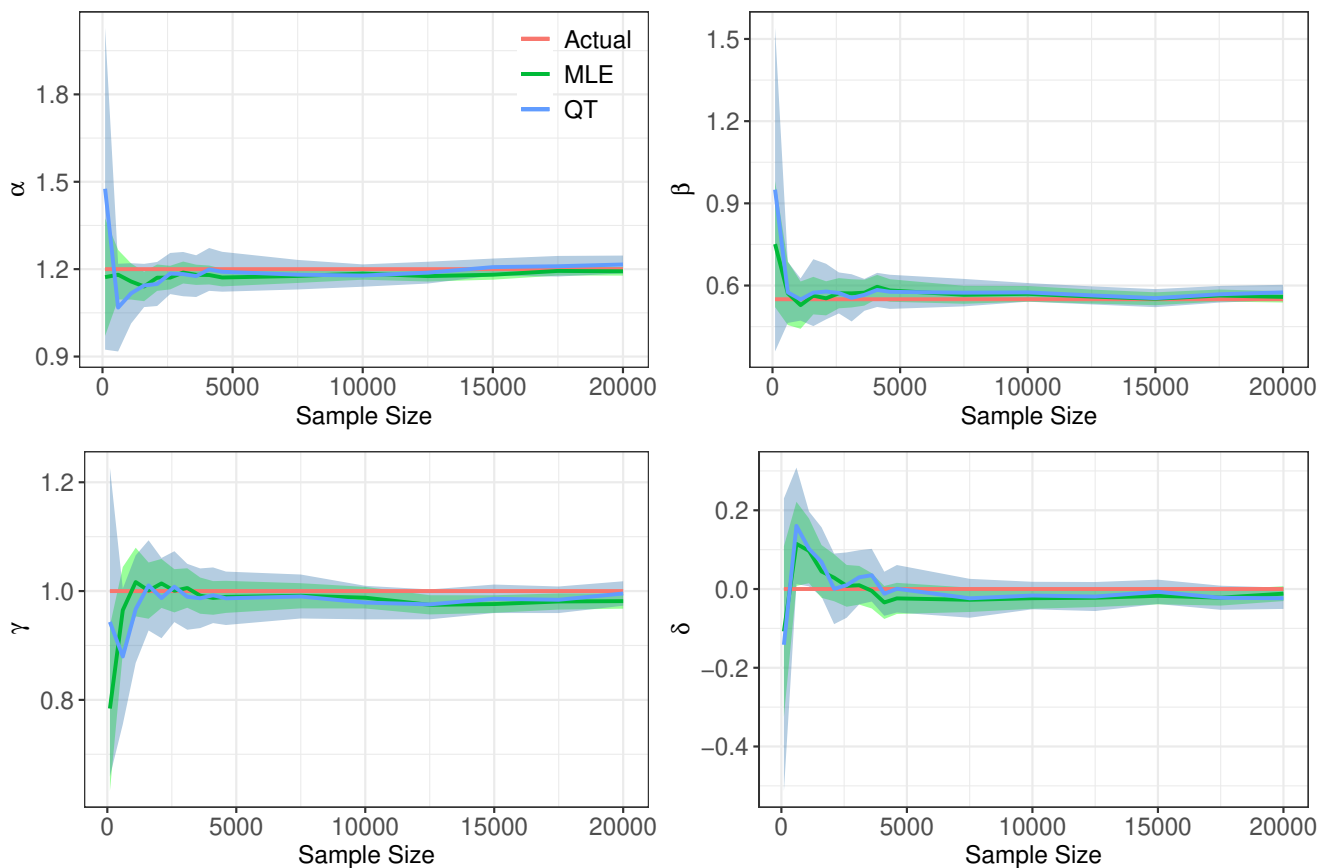


Figure 14: Accuracy of Lévy alpha-stable parameter estimations using McCulloch's quantile method (QT) and maximum likelihood estimation (MLE) in comparison.




Figure 15: Soofi ID scores of our McCulloch quantile estimation fits given in figure 6 with those of a GMM fit of the same subsamples (by country and industry, by country and year, and by country and firm size).

Soofi ID. The Soofi information distinguishability (ID) index (Soofi et al. 1995) compares the actual empirical distribution $p(x)$ against the predicted density of x from a candidate model $q(x)$. Formally, the

Soofi ID over a set of observations x with fit $p(x)$ and corresponding probability model with parameters θ , $q(\theta, x)$, is defined as:

$$ID(p||q|\theta) = \exp[-D_{\text{KL}}(p||q|\theta)],$$

where $||$ is the divergence operator²⁴ and $D_{\text{KL}}(p||q)$ is the Kullback-Leibler (KL) divergence. The KL divergence is, in turn, is defined as follows:

$$D_{\text{KL}}(p||q) = \sum_i p(x_i) \log \frac{p(x_i)}{q(x_i)},$$

where the information content of p is compared to that of q . The Soofi ID was first suggested as a measure of distinguishability of multiple distributions p with the same parameters θ . Since it employs the distance of these from the corresponding maximum entropy distribution q , it can, however, be employed as goodness of fit measure, as suggested in, among others, Yang (2018). The Soofi ID index is a normalized measure, so $0 \leq ID(p||q|\theta) \leq 1$. We use a rescaled version multiplied by 100. We can interpret the value of the rescaled Soofi ID index as the percentage of the empirical distribution that is recovered by the model. If the Soofi ID index is 98, for example, this means that the model recovers 98% of the information in the data.²⁵ When there is no deviation between the empirical density, $p(x)$, and the predicted density from the model, $q(x)$, the Soofi ID score becomes 100, meaning that the fitted model explains 100% information from the data. The lower the Soofi ID, the less predictive power the model has.

AIC. Akaike's information criterion (AIC) for model selection is defined as (Akaike 1973)

$$\text{AIC} = 2k - 2 \log[\bar{L}], \quad (13)$$

where k is the number of parameters to estimate and $\log[\bar{L}]$ is the log-likelihood given the estimated parameters of the model. Note that the AIC can be derived from the KL divergence (Akaike 1973). The AIC provides the relative superiority of a certain model among the candidates, not the absolute quality of the model. The lowest value in a comparison of fits indicates the best model in the selection. Since Lévy alpha-stable distribution and the AEP distribution have the same number of parameters, the AIC difference will be equivalent to a log-likelihood ratio. In comparing two models, we will show the relative likelihood per data point defined as $\exp((\text{AIC}_{\text{Levy}} - \text{AIC}_{\text{AEP}})/2N)$ where AIC_{Levy} and AIC_{AEP} are the AIC of the Lévy alpha-stable model and the AEP model, respectively, and N is the number of observations. When the Lévy alpha-stable model performs better and has a lower AIC, the relative likelihood is always smaller than 1 and tells us how much the AEP model is as probable as the Lévy alpha-stable model for an in-sample prediction. For example, if the average relative likelihood is 0.95, this means that the inferior model is 0.95 times probable relative to the superior model in explaining the data. See Royall (2017) and Burnham & Anderson (2004) for a detailed discussion on the interpretation of the relative likelihood.

K-fold cross-validation While Soofi ID and AIC are used for an in-sample prediction (testing the accuracy of model fitting), K-fold cross-validation performs an out-of-sample prediction (testing the predictive power of the model for new observations). In a simple K-fold cross-validation, the sample is partitioned into K equally sized subsamples. The estimation of the given model is then repeated K times

²⁴I.e., $p||q|\theta$ is to be read as divergence of p and q given θ .

²⁵Note that the original Soofi ID is defined as $1 - \exp[-D_{\text{KL}}(p||q|\theta)]$ and therefore the interpretation is reversed. In our measure, a higher value implies that the actual distribution is similar to the maximum entropy distribution.

holding the K -th subsample as a test set and using the other subsamples as a training set. The average log-likelihood of all K test sets given different models is recorded for M iterations. In comparing the Lévy alpha-stable model and AEP model distribution, we choose $K = 10$ and $M = 10$. That is, we randomly partition the country-year sample into 10 equally sized subsamples and carry out 10 estimations sequentially using each of 10 subsamples as a holdout. We repeat this 10 times and compare the average log-likelihood of each model. In comparing the cross-validation results, we will also show the relative likelihood per data point defined as $\exp((CV_{AEP} - CV_{Levy})/N)$ where CV_{Levy} and CV_{AEP} are the k -fold CV results of the Lévy alpha-stable and AEP model, respectively. The interpretation of the relative likelihood is the same as before. When the Lévy alpha-stable model performs better and has a higher likelihood, the relative likelihood is always smaller than 1 and tells us how much the AEP model is as probable as the Lévy alpha-stable model for an out-of-sample prediction.

F Additional results

In addition to the results presented in the main part, we present more detailed statistics on some aspects in the present section. We discuss the tail behavior of LP and ΔLP in Section F.1 and the comparison of macro-level records of labor productivity growth and the location parameter δ in the Lévy alpha-stable fit of ΔLP in Section F.3.

F.1 Tail behaviour

In this appendix, we take a closer look at the tail behavior.

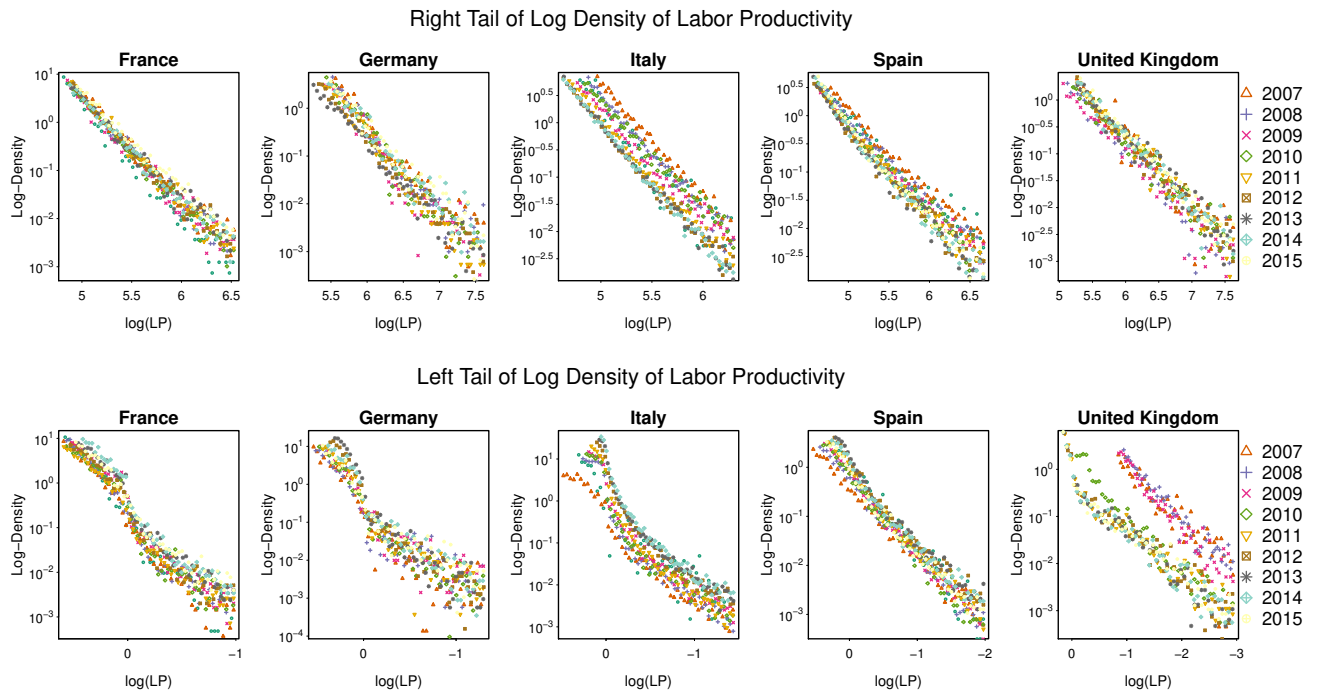


Figure 16: **Left and right tails of labour productivity levels distributions.**

Figure 16 plots the tails of the labour productivity distribution, i.e. the most and least productive 5% of firms respectively, for selected countries and by year; Figure 17 repeats this exercise for labour

productivity change. All observations are plotted on a log-log scale. Tails are calculated as the distance from the mode of the distribution. This enables us to take the log of the left tail. The right tails appear remarkably linear in log-log scale. For the left tails, this is less homogeneous. The form is still linear for the some countries (e.g., for the United Kingdom and for Spain in the case of the labour productivity) while exhibiting more variation around an approximately linear decay for others. The strongly linear form of the right tails with more variation in the left persists when arranging the sub-samples by industry or firm size, as seen in Figures 16-21 (all series with observation count exceeding a minimum of 10,000 are included).




Figure 17: Left and right tails of labour productivity change distributions.

F.2 Tail plots by size and sector




Figure 18: Left and right tails of labour productivity levels distributions, by firm size. Firm sizes denoted as S (Small), M (Medium), L (Large), V (Very large).




Figure 19: Left and right tails of labour productivity change distributions, by firm size. Firm sizes denoted as S (Small), M (Medium), L (Large), V (Very large).

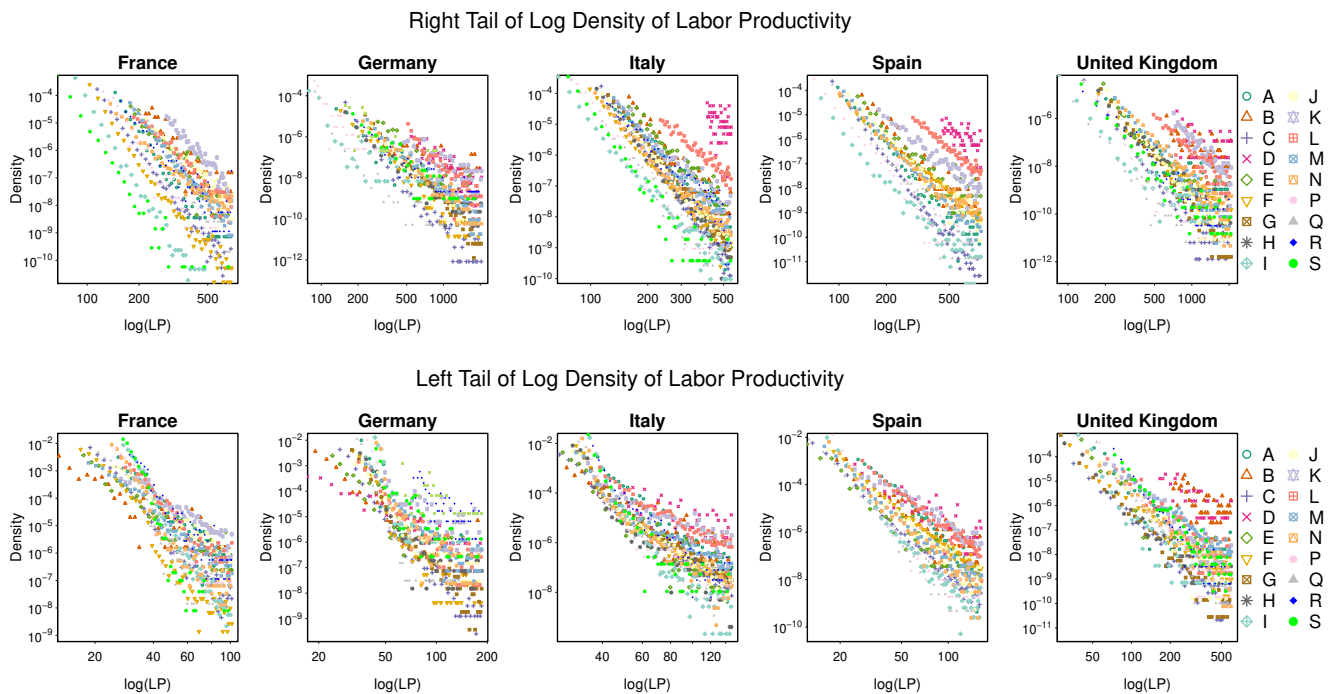


Figure 20: Left and right tails of labour productivity levels distributions, by industry. Industry codes as listed in table 5.




Figure 21: **Left and right tails of labour productivity change distributions, by industry.** Industry codes as listed in table 5.

F.3 The location parameter δ and the aggregate labor productivity

A comparison of our yearly estimates for the central moment of labor productivity change ΔLP , the location parameter δ in our Lévy alpha-stable fits with aggregate labor productivity change data from the *Penn World Tables* shows good agreement of both time series for almost all countries (see Figure 22). This indicates (1) that the firm-level productivity change data used in this paper is sufficiently representative and (2) that the parameters obtained by the Lévy alpha-stable fits are meaningful and predictive.

The figure only shows data for ten countries, as data for the other five countries considered in the present analysis (Bulgaria, Estonia, Hungary, Slovenia, Slovakia) have too short time span for a meaningful comparison.

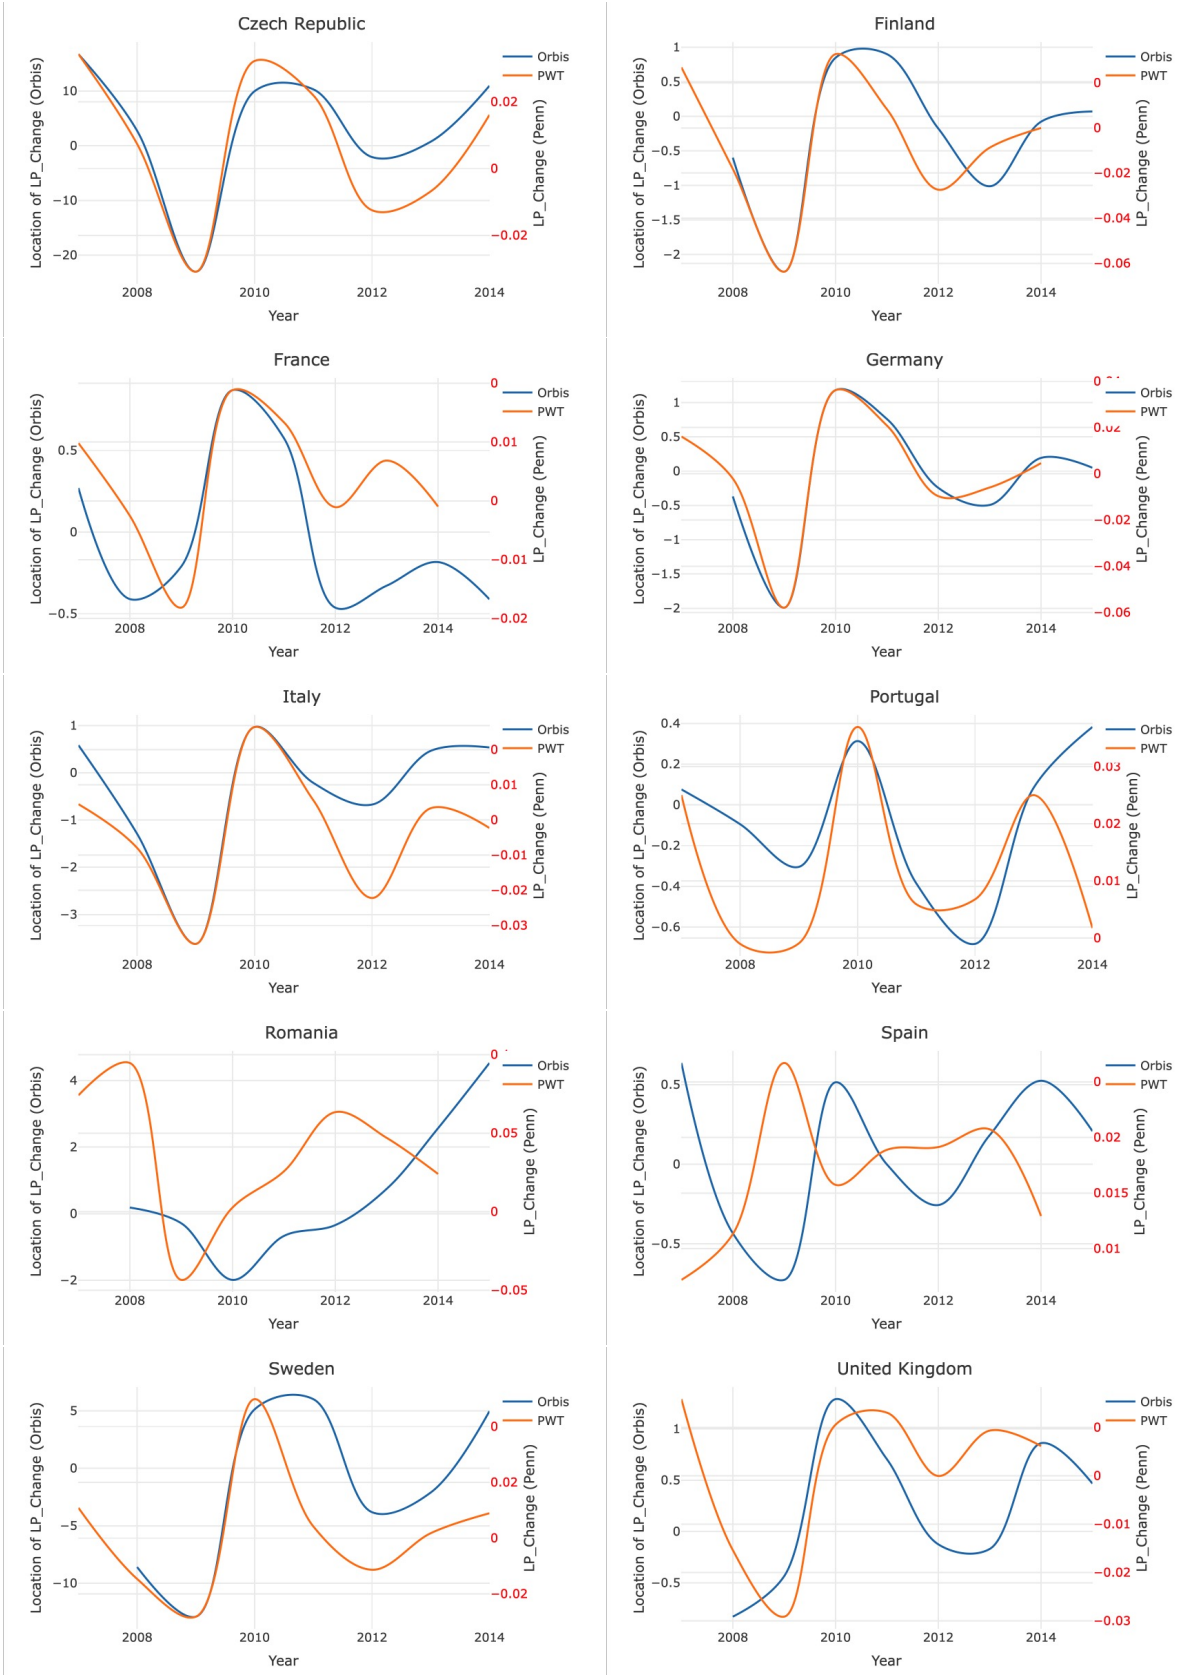


Figure 22: Location parameter δ of yearly Lévy alpha-stable fits of firm-level labor productivity change in Obiris Europe data compared to aggregate labor productivity change data from *Penn World Tables* for the ten countries with sufficient data quality both in Obiris Europe and in the *Penn World Tables*.

G Estimation result for the log variables

This section reports on the estimation result for the log variables. The two measures of log firm-level productivity used in this section are

$$\text{Log}(LP_{it}) = \text{Log}\left(\frac{Y_{it}}{L_{it}}\right), \quad (14)$$

$$LP_{git} = \text{Log}(LP_{it}) - \text{Log}(LP_{it-1}), \quad (15)$$

The first measure of productivity in Eq. (14) is the log labour productivity (LP_{it}), simply the log ratio of value-added to the number of employees. The second measure in Eq. (15) is the log growth rate as the log ratio of LP_{it} to LP_{it-1} . Both measures are calculated after removing non-positive values from the data.

Trapani's test for the first and second moment of these variables is presented in Table 9. The results are ambiguous, suggesting that both moments are infinite for Spain and the UK, and finite for France and Germany. Looking at Table 2, we may conjecture that the surprising results for Spain and the UK come from data truncation, because there is a higher share of negative observations in Spain and the UK than Germany and France (roughly, 5% rather than 2%). This intuition appears to be confirmed by looking at the full results in Table 7. We see that for Italy, which has a number of negative observations rising from about 2 to 5%, the test for the first moment for log LP and log $LPchange$ rejects infinite moment in early years (when there are few negative observations), and does not reject in later years (when there are many negative observations).

Variable	Moment	France	Germany	Italy	Spain	U.K.
Log LP	First	724.55	226.50	0.00	0.00	0.00
		(0.00)	(0.00)	(1.00)**	(1.00)**	(1.00)**
Log LP Change	First	367.73	128.17	872.81	0.00	206.72
		(0.00)	(0.00)	(0.00)	(1.00)**	(0.00)
Log LP	Second	875.84	273.64	0.00	0.00	0.00
		(0.00)	(0.00)	(1.00)**	(1.00)**	(1.00)**
Log LP Change	Second	441.06	153.04	1046.68	0.00	246.92
		(0.00)	(0.00)	(0.00)	(1.00)**	(0.00)

Table 9: Test statistics ϑ_{nr} for Trapani's test for infinite moments of order 1 and 2 in the levels and changes of LP and log LP. P.values underneath in parentheses. Firm-level data collected from ORBIS Europe for 2014.

Figures 23, 25a, and 25b show the original distributions with the fitted lines for year, size, industry samples for France as an example. The Lévy alpha-stable model generally gives a very good fit with the Soofi ID higher than 99% in most cases. Figures 24, 26, 27a, and 27b shows the the estimated four parameters with the bootstrapping standard error. Detailed comparisons with the non-log variables can be found in Sections 4 and 5.




Figure 23: **Distribution of French labor productivity and its change by year with Lévy alpha-stable fits.** Years in 2008, 2011, 2014, which exemplify the other years range from 2007-2015. Solid lines indicate the fitted Lévy alpha-stable distributions, for which the estimated parameters are available in Figure 24.

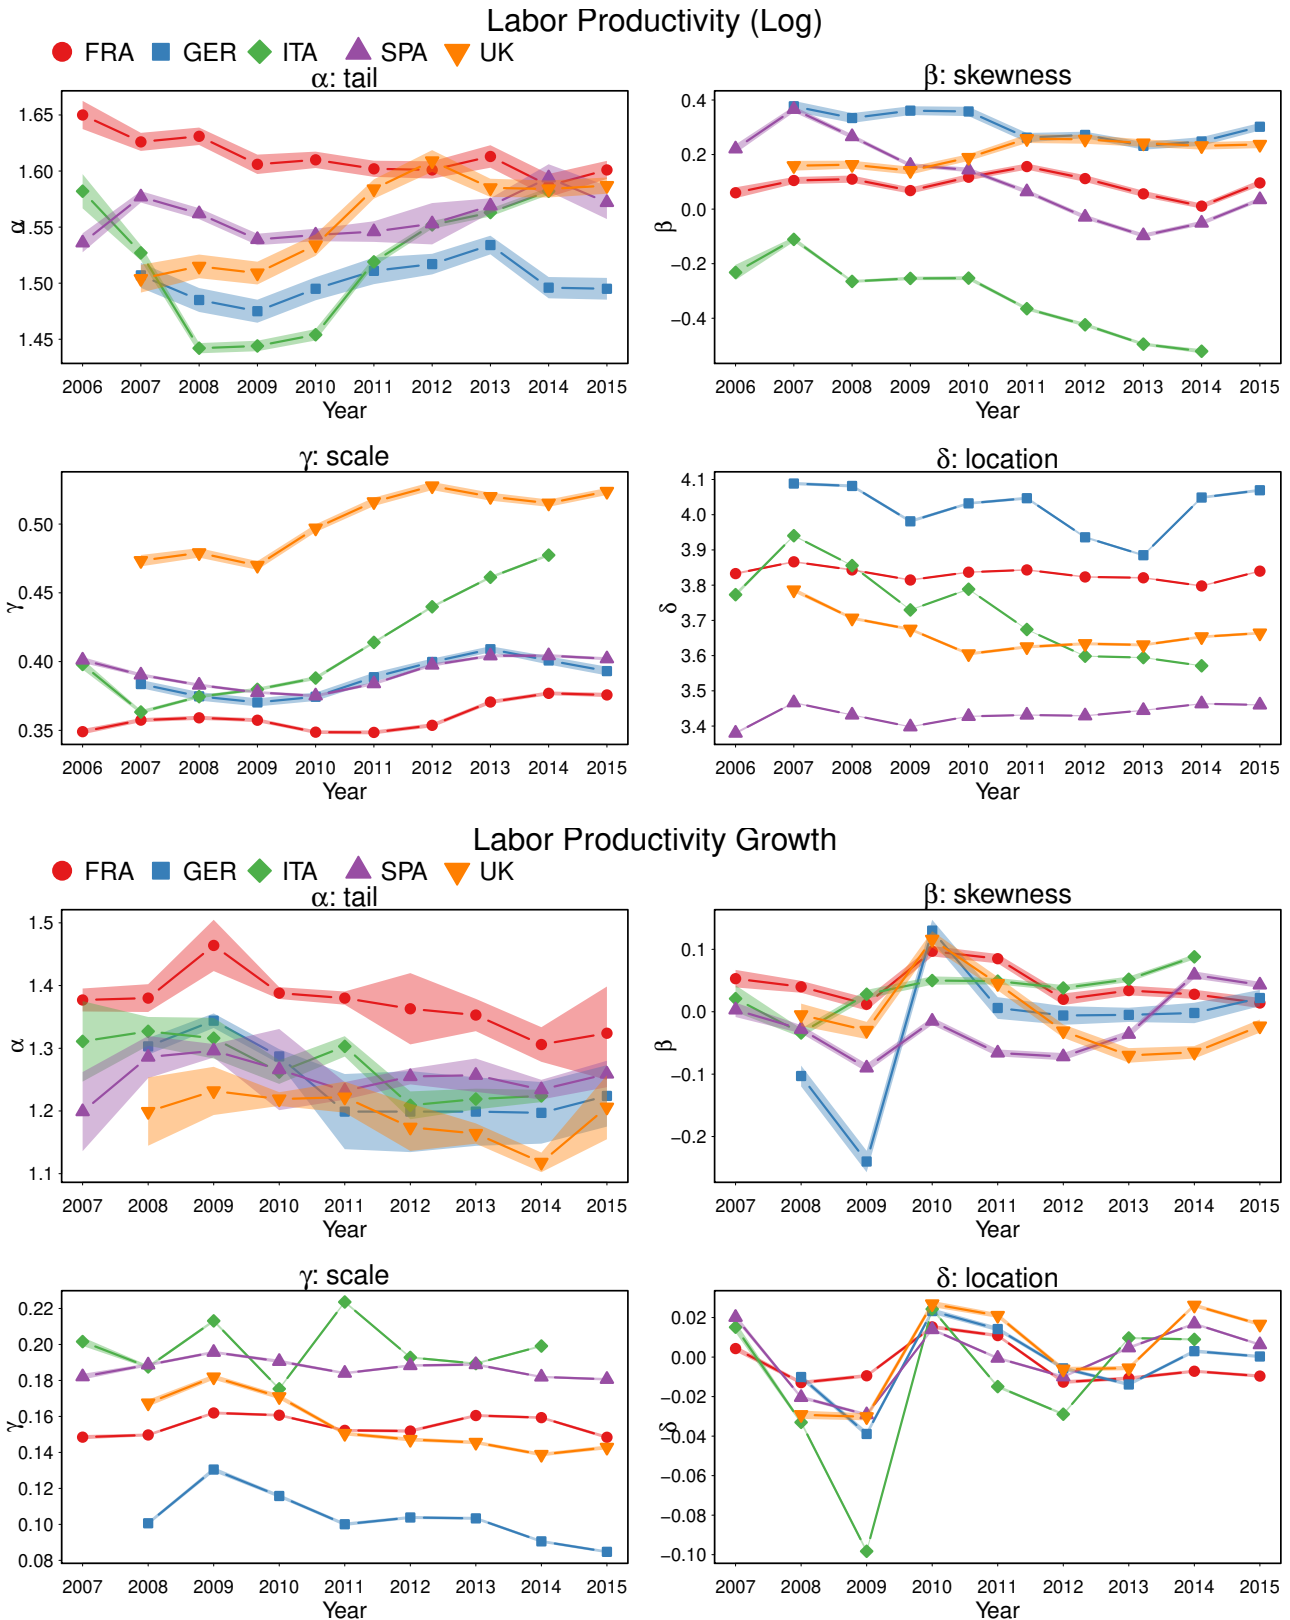
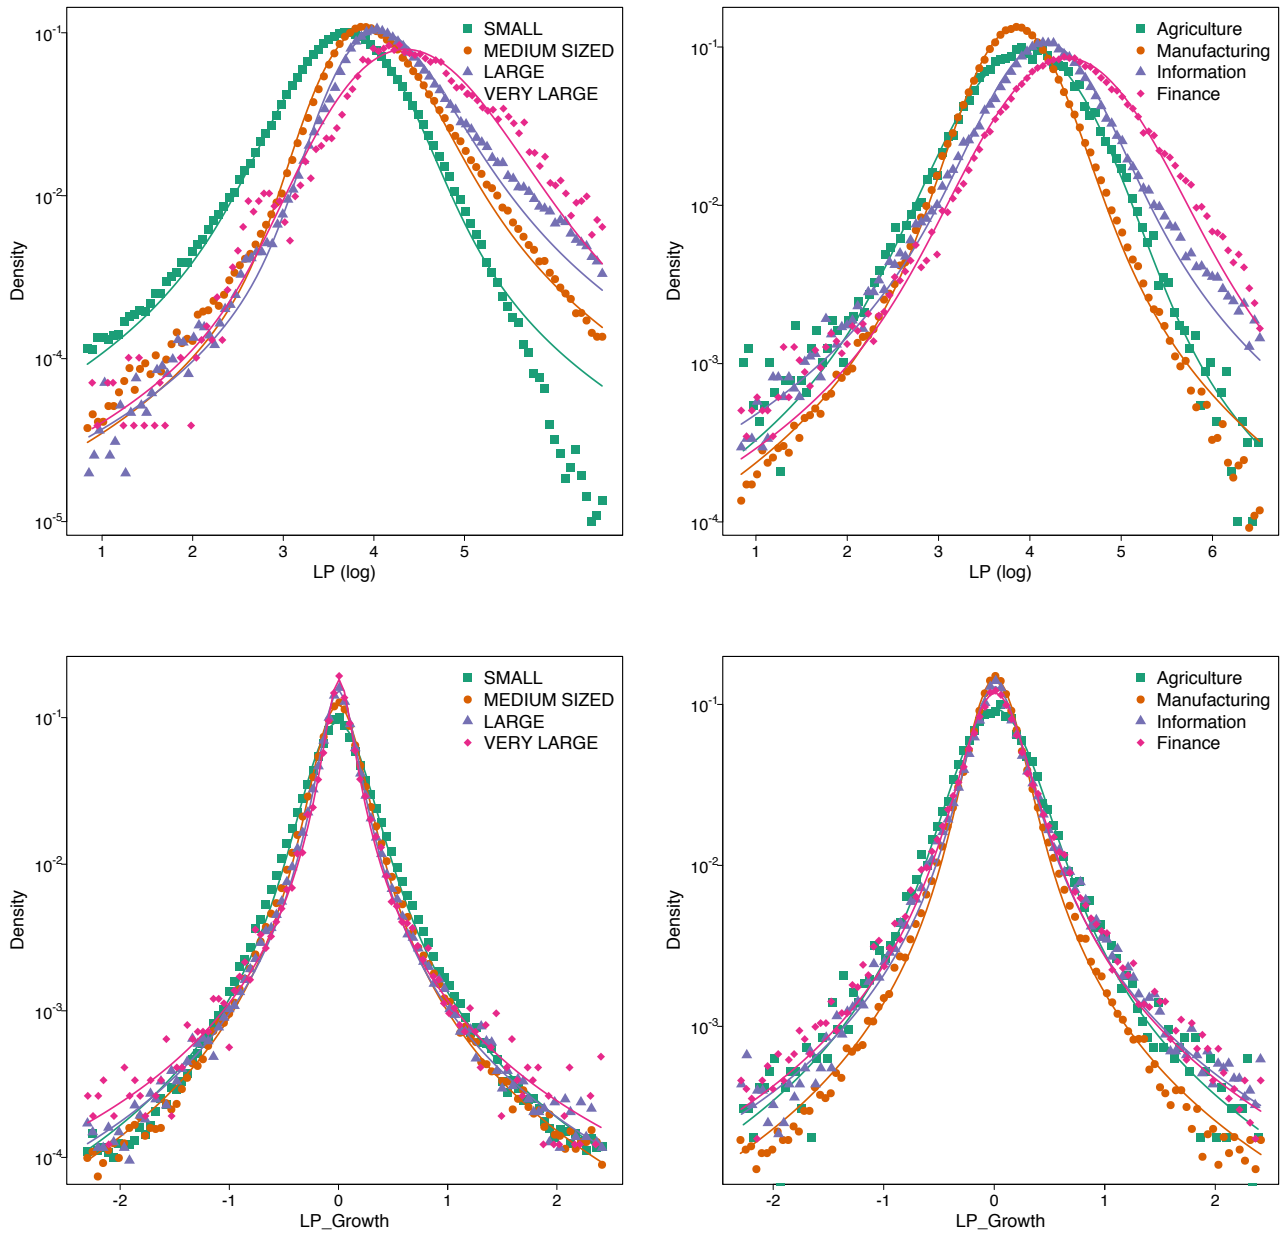


Figure 24: **Estimated parameters for country-year samples.** The four parameters of the fitted Lévy alpha-stable distributions for labour productivity levels (top row) and growth (bottom row) are plotted by year (2006-2015 for levels, 2007-2015 for growth). γ and δ are denominated in €1,000/employee, α and β are dimensionless. The shaded area represents the ± 1 bootstrapped standard errors.



(a) **Size:** The top panel shows the distribution of the level of labour productivity in France and firm size. The second row computes the changes in labour productivity. Solid lines indicate the fitted Lévy alpha-stable distributions, for which the estimated parameters are available in Figure 26. Firm size categorisations can be found in Appendix A.5. All sub-groups here meet the threshold observation count of 5,000.

(b) **Industry:** The first graph shows the distributions of the level of labour productivity in France by industry. The second computes the changes in labour productivity. Solid lines indicate the fitted Lévy alpha-stable distributions, for which the exact parameterisations are available in Figures 27a and 27b. Firms are categorised by their provided Nace Rev. 2 classification code. An overview of the industry headers and corresponding descriptions is available in Appendix A.5. All groups here meet the threshold observation count of 1,000.

Figure 25: Distributions of labor productivity in France with Lévy alpha-stable fits.

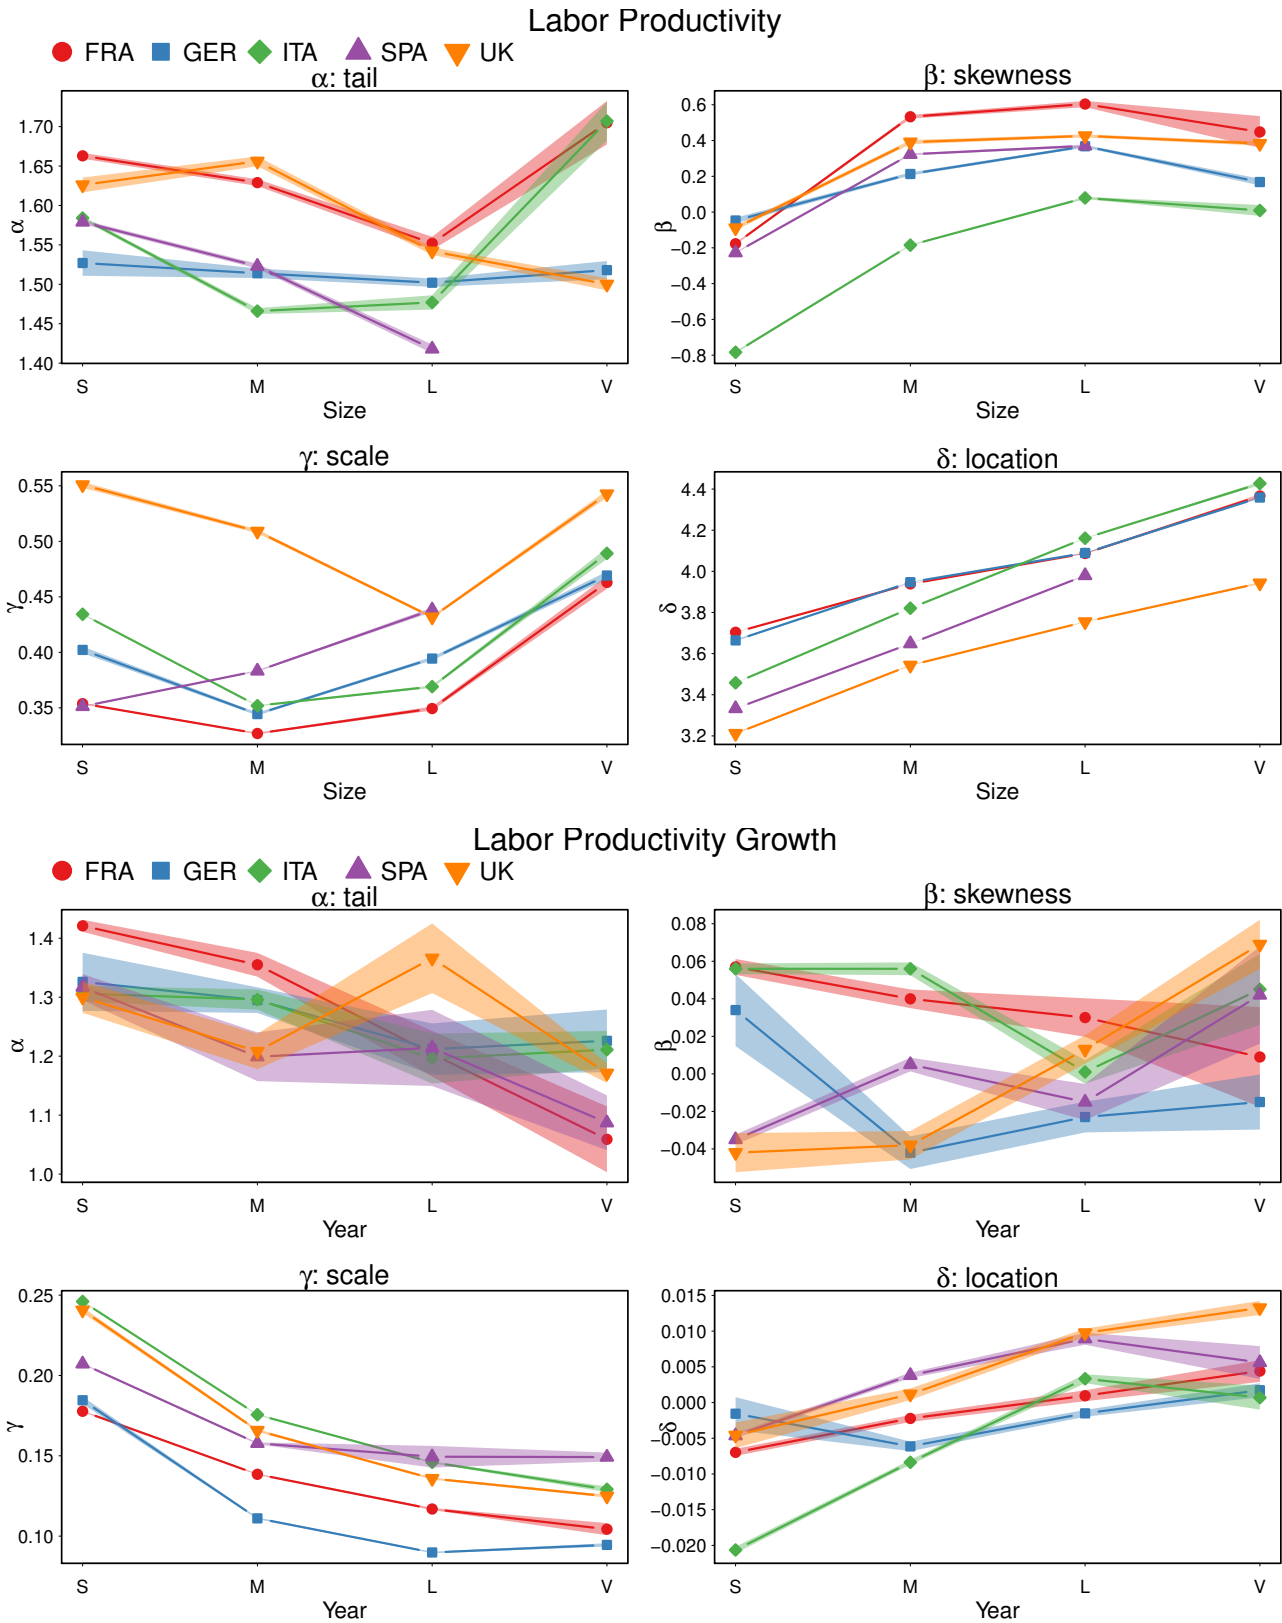
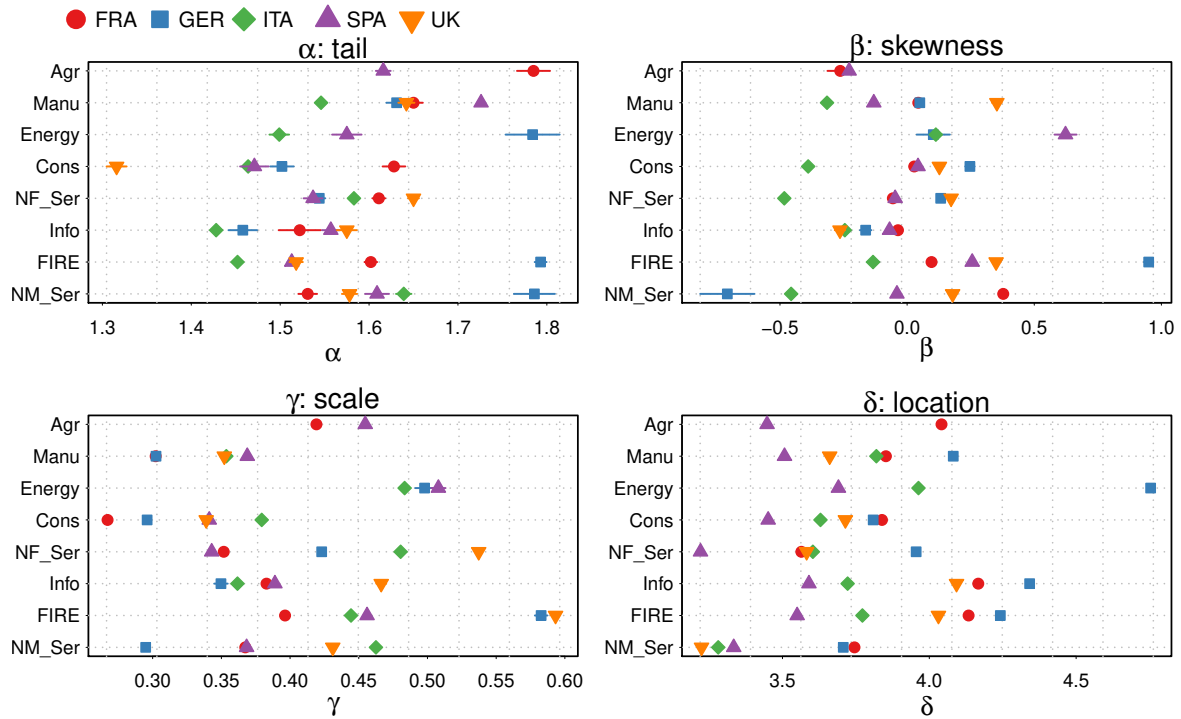
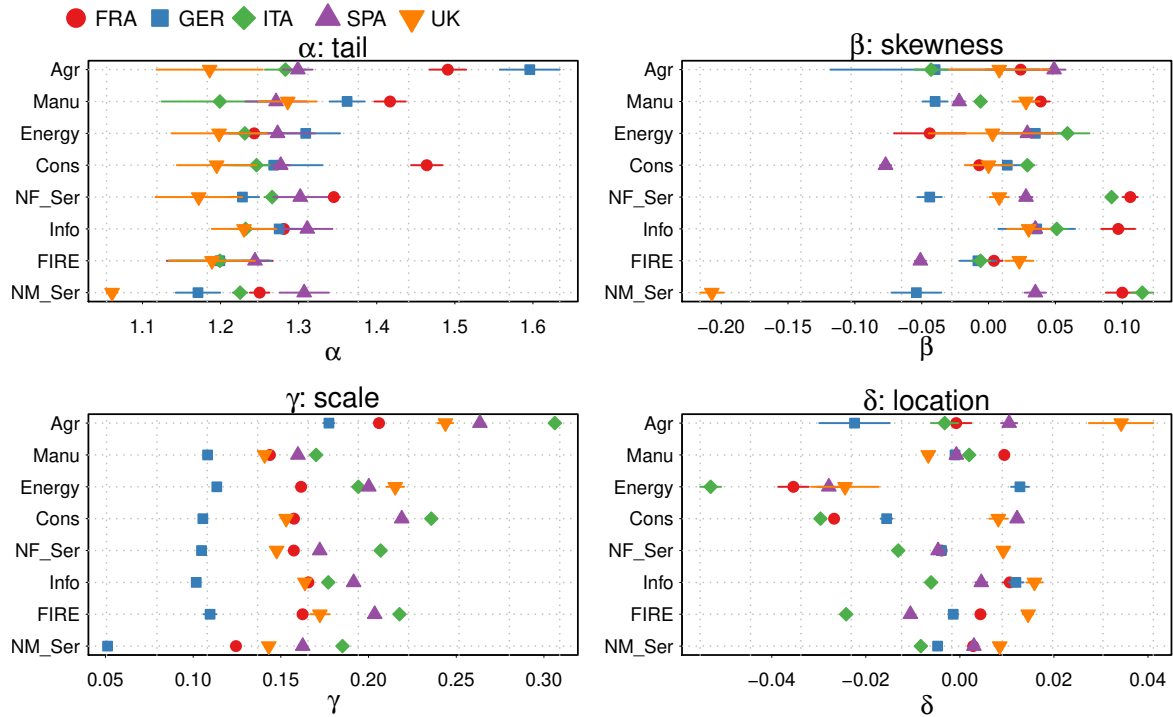


Figure 26: **Estimated parameters country-size samples.** The four parameters of the fitted Lévy alpha-stable distributions for labor productivity levels (top row) and growth (bottom row) are plotted by firm size (small, medium, large, very large). γ and δ are denominated in €1,000/employee, α and β are dimensionless. The shaded area represents the ± 1 bootstrapped standard errors.



(a) Log Labour Productivity Levels



(b) Log Labor Productivity Growth

Figure 27: **Estimated parameters of log variables for country-industry sub-samples.** The four parameters of the fitted Lévy alpha-stable distributions for labour productivity growth are plotted by industry. The data shown here only uses samples for France, Germany, Italy, Spain, and the United Kingdom. Parameter estimates for certain sub-samples are not reported if sample size does not meet the required threshold of 1000 observations. γ and δ are denominated in $\text{€}1,000/\text{employee}$, α and β are dimensionless. The horizontal bar represents the ± 1 bootstrapped standard errors.

Table 10 summarizes the model comparison results by showing the average value of the three measures for the country-level samples. For the log labor productivity level, the average difference of Soofi ID index between the Lévy alpha-stable and the AEP distribution is -0.5% , implying that the latter model extracts 0.5 percentage points more information from the empirical data than the latter model. The relative likelihood per data point both in the AIC and the cross-validation is 1.01. This means that the AEP is 1.01 times as probable as the Lévy stable distribution for the in-sample and out-of-sample prediction. For the log growth rate of labor productivity, the AEP model extracts 0.2 percentage points more information on average than the Lévy stable distribution. From the average relative log-likelihood, we note that the former model is 1.01 times as probable as the former model. Overall, the AEP is generally preferred over Lévy alpha-stable distribution for both of the variables. However, the difference of the model performance is very small compared to the result for non-log variables.

H Industry results for the original NACE classification

The following figures show the estimation results for the country-industry sample using the original NACE ver.2 industry classification. See Table 5 for the full detail of the classification.

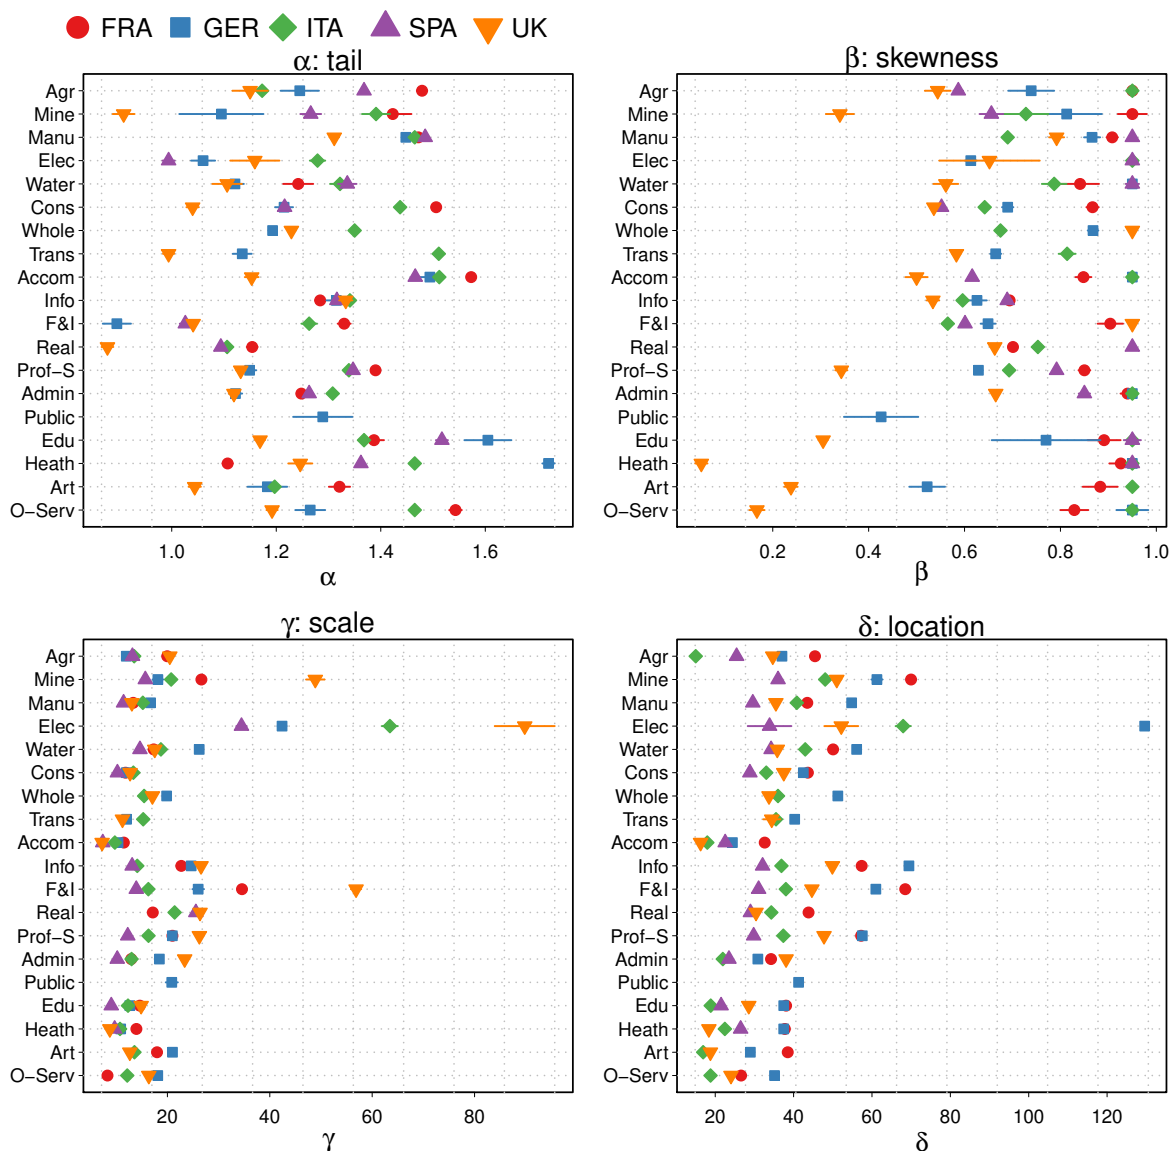


Figure 28: **Estimated parameters for country-industry sub-samples, Labour Productivity Levels.** The four parameters of the fitted Lévy alpha-stable distributions for labour productivity growth are plotted by industry. The data shown here only uses samples for France, Germany, Italy, Spain, and the United Kingdom. Parameter estimates for certain sub-samples are not reported if sample size does not meet the required threshold of 1000 observations. γ and δ are denominated in €1,000/employee, α and β are dimensionless. The horizontal line represents ± 1 bootstrapping standard deviation.




Figure 29: **Estimated Parameters with Bootstrapping Errors for Country-Industry Sub-Samples, Labor Productivity Changes.** The four parameters of the fitted Lévy alpha-stable distributions for labour productivity growth are plotted by industry. The data shown here only uses samples for France, Germany, Italy, Spain, and the United Kingdom. γ and δ are denominated in €1,000/employee, α and β are dimensionless. The horizontal bar represents the ± 1 bootstrapped standard errors. Parameter estimates for certain sub-samples are not reported if sample size does not meet the required threshold of 1000 observations.

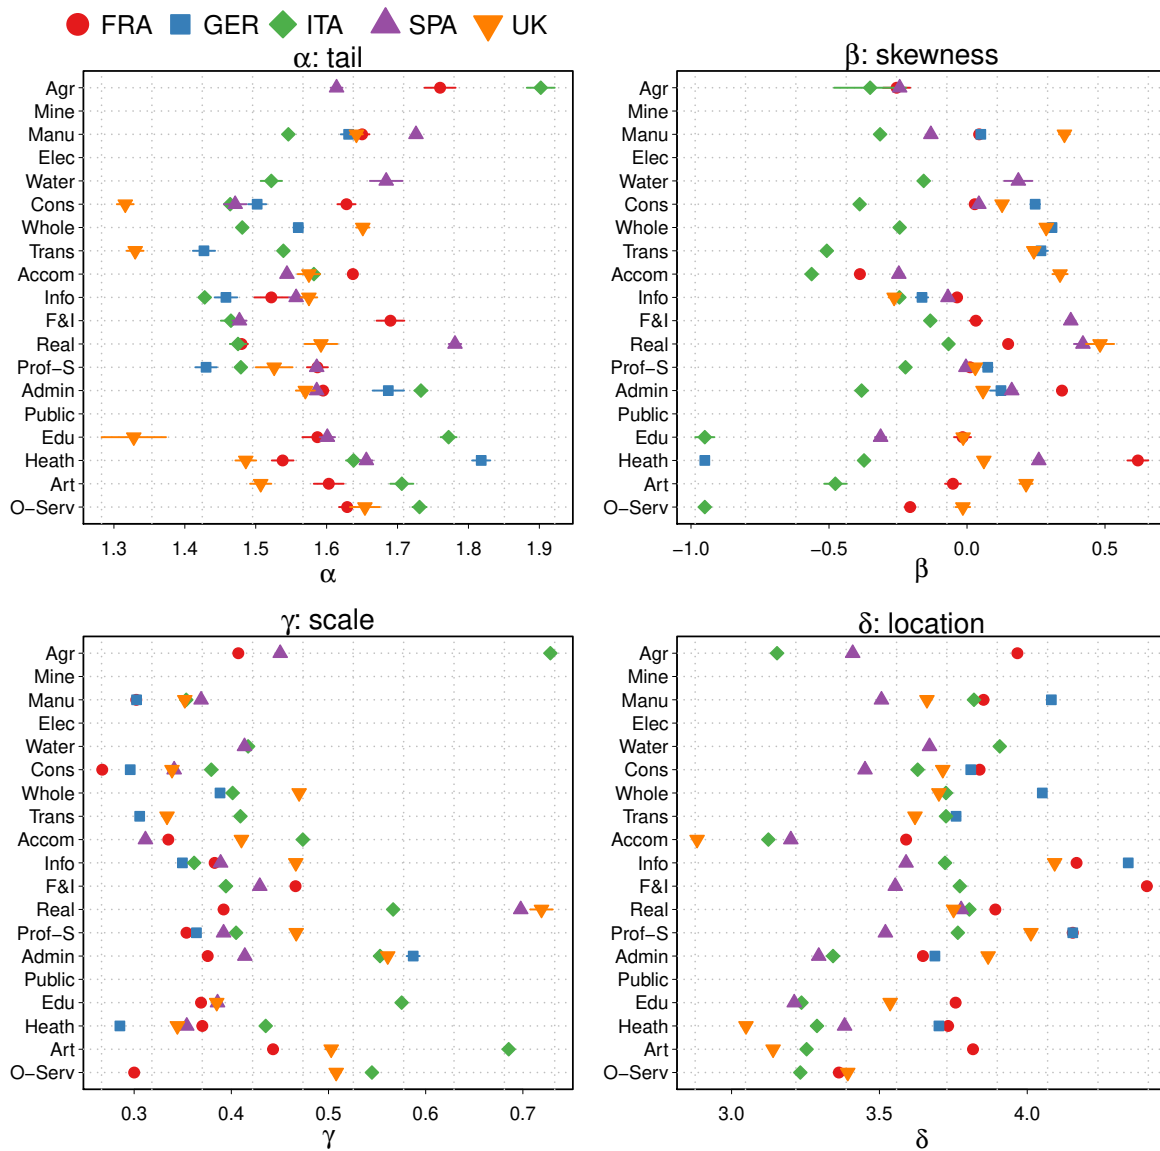


Figure 30: **Estimated parameters for country-industry sub-samples, Log-Labour Productivity.** The four parameters of the fitted Lévy alpha-stable distributions for log labour productivity are plotted by industry. The data shown here only uses samples for France, Germany, Italy, Spain, and the United Kingdom. γ and δ are denominated in $\log(\text{€}1,000/\text{employee})$, α and β are dimensionless. The horizontal bar represents the ± 1 bootstrapped standard errors. Parameter estimates for certain sub-samples are not reported if sample size does not meet the required threshold of 1000 observations.

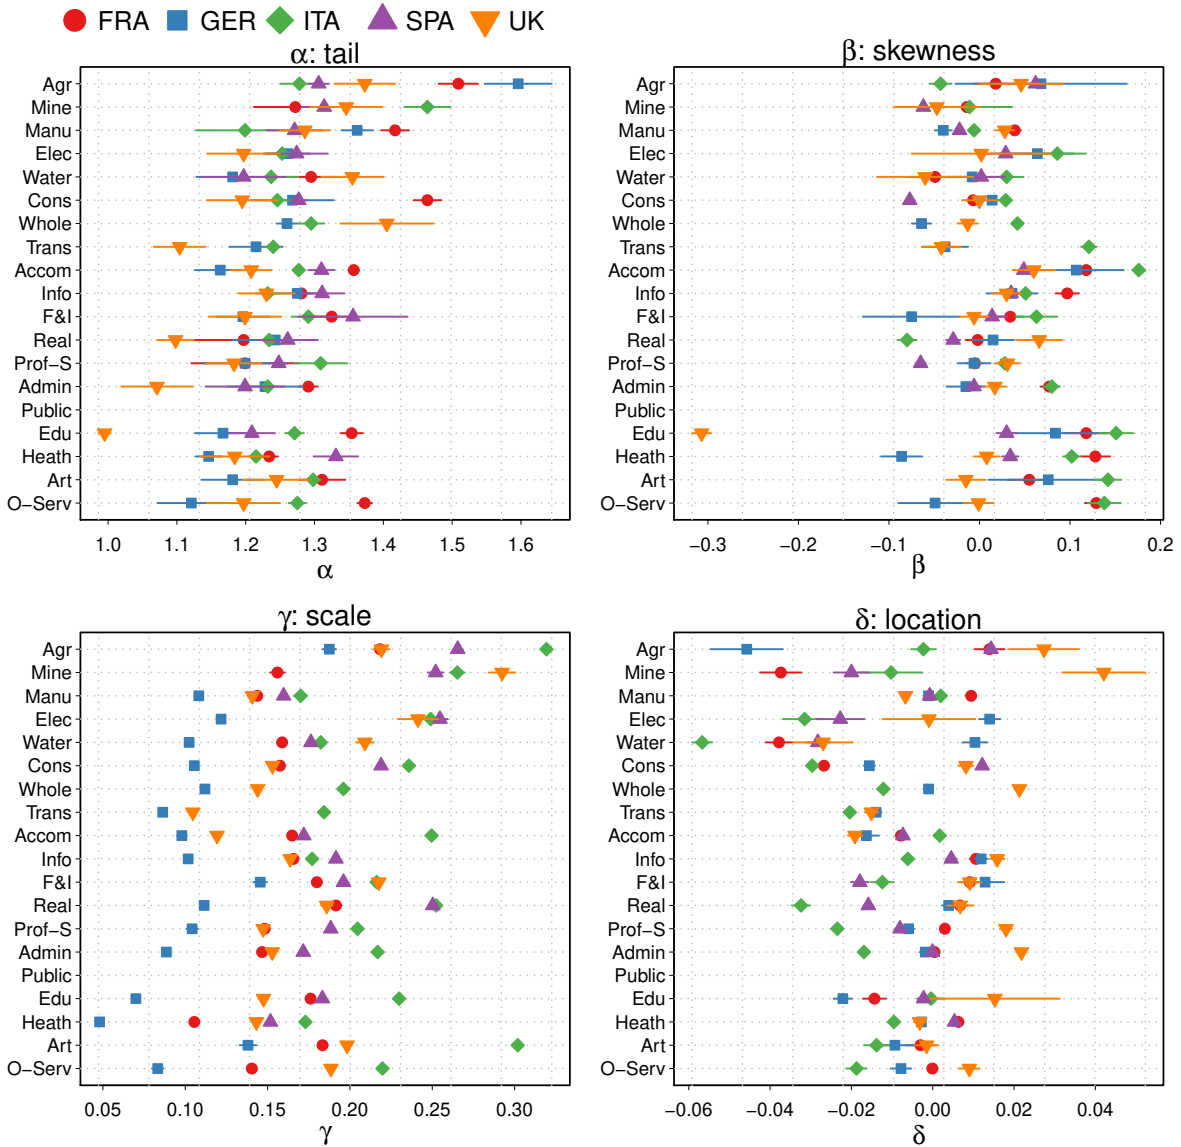


Figure 31: **Estimated Parameters with Bootstrapping Errors for Country-Industry Sub-Samples, Log Labor Productivity Growth.** The four parameters of the fitted Lévy alpha-stable distributions for labour productivity growth are plotted by industry. The data shown here only uses samples for France, Germany, Italy, Spain, and the United Kingdom. γ and δ are denominated in log (€1,000/employee), α and β are dimensionless. The horizontal bar represents the ± 1 bootstrapped standard errors. Parameter estimates for certain sub-samples are not reported if sample size does not meet the required threshold of 1000 observations.

I Further details on the effect of the logarithmic transformation on IQRs

As an extreme example, when the distribution simply shifts to the left, the dispersion measured by the 90-10 IQR increases even when the underlying distribution stays the same. Figure 32 graphically demonstrates this effect. The left panel shows two same Gamma distributions with the shape parameter

3 and the rate parameter 1. The Gamma distribution in the bold line is shifted from the same Gamma distribution in the bold line by 1. The 90-10 IQR is 4.2 for these two distributions. The log-transformed distributions are in the right panel. Even when the underlying distributions are the same, the IQR for the left distribution in the dotted line is 1.58, while that for the right distribution is 1.10. Standardization of the underlying distributions by some location measures, e.g. the mean, will solve this issue to some degree. But, it is sometimes not easy to find the common location measures for the different distributions and the results might be also quite sensitive to the choice of the particular location measure.




Figure 32: **Density of the original and logged Gamma random variable** Variable x follows a Gamma distribution with a shape parameter of three and a rate parameter of one on the dotted line. The bold line represents a transform of x by shifting it by one.

Table 8: Model Comparison: Country Average

Variable	Country	10-Fold Cross-Validation			Soofi-ID			AIC		
		Lévy	AEP	\mathcal{L} Ratio	Lévy	AEP	Diff	Lévy	AEP	\mathcal{L} Ratio
LP	Bulgaria	-257,825 (-3.98)	-259,160 (-4)	0.98	97.2	89.4	7.8	515648 (7.96)	518320 (8)	0.98
	Czech Republic	-388,911 (-7.42)	-392,263 (-7.49)	0.94	99.1	94.5	4.6	777,831 (14.85)	784,526 (14.98)	0.94
	Estonia	-41293 (-3.9)	-41617 (-3.93)	0.97	97.9	95.6	2.3	82589 (7.81)	83234 (7.87)	0.97
	Finland	-141,086 (-4.8)	-142,290 (-4.84)	0.96	99.5	94.5	5.0	282,171 (9.6)	284,579 (9.68)	0.96
	France	-508,009 (-4.82)	-513,626 (-4.87)	0.95	99.8	95.8	4.0	1,016,015 (9.64)	1,027,252 (9.75)	0.95
	Germany	-193,706 (-5.2)	-197,977 (-5.32)	0.89	98.9	92.9	6.1	387,414 (10.41)	395,953 (10.64)	0.89
	Hungary	-442,090 (-9.34)	-444,518 (-9.39)	0.95	99.3	93.9	5.4	884,172 (18.68)	889,035 (18.78)	0.95
	Italy	-1,155,033 (-4.79)	-1,166,870 (-4.84)	0.95	99.3	94.0	5.3	2,310,066 (9.58)	2,333,738 (9.68)	0.95
	Portugal	-523,010 (-4.01)	-526,223 (-4.03)	0.98	99.5	95.6	3.9	1,046,022 (8.01)	1,052,445 (8.06)	0.98
	Romania	-346,314 (-4.87)	-348,409 (-4.9)	0.97	98.8	96.8	2.0	692,630 (9.74)	696,818 (9.8)	0.97
	Slovakia	-124,891 (-4.29)	-126,252 (-4.34)	0.95	98.9	94.8	4.1	249,784 (8.59)	252,503 (8.68)	0.95
	Slovenia	-48,084 (-4.09)	-48,303 (-4.11)	0.98	97.8	95.8	2.0	96,169 (8.19)	96,604 (8.22)	0.98
	Spain	-845,671 (-4.61)	-860,433 (-4.69)	0.92	99.9	84.1	15.7	1,691,341 (9.22)	1,720,866 (9.38)	0.92
	Sweden	-625,595 (-7.12)	-627,063 (-7.13)	0.98	99.8	97.6	2.2	1,251,190 (14.24)	1,254,124 (14.27)	0.98
	United Kingdom	-285,557 (-5.24)	-291,995 (-5.36)	0.89	98.7	92.9	5.8	571,105 (10.49)	583,990 (10.73)	0.89
		Average	-395,138 (-5.23)	-399,133 (-5.28)	0.95	99.00	93.90	5.1	790,276 (10.47)	798266 (10.57)
LP Change	Bulgaria	-213117 (-3.76)	-214410 (-3.78)	0.98	99.2	96.6	2.6	426190 (7.51)	428820 (7.56)	0.98
	Czech Republic	-328,372 (-7.05)	-328,504 (-7.06)	1.00	99.2	97.1	2.1	656,666 (14.11)	657,008 (14.11)	1.00
	Estonia									
	Finland	-111,548 (-4.38)	-111,789 (-4.39)	0.99	99.4	98.0	1.4	223,101 (8.77)	223,576 (8.78)	0.99
	France	-384,476 (-4.34)	-386,071 (-4.36)	0.98	99.5	98.8	0.8	768,951 (8.68)	772,139 (8.71)	0.98
	Germany	-136,038 (-4.29)	-138,113 (-4.36)	0.94	98.4	93.6	4.8	272,105 (8.59)	276,225 (8.72)	0.94
	Hungary	-394,104 (-9.18)	-394,740 (-9.2)	0.99	99.5	98.8	0.7	788,207 (18.37)	789,479 (18.4)	0.99
	Italy	-955,516 (-4.41)	-959,238 (-4.43)	0.98	99.5	98.5	1.0	1,911,032 (8.82)	1,918,475 (8.85)	0.98
	Portugal	-438,372 (-3.56)	-440,800 (-3.58)	0.98	99.6	96.3	3.3	876,775 (7.13)	881,597 (7.17)	0.98
	Romania	-309,205 (-4.76)	-309,902 (-4.77)	0.99	99.4	94.4	5.0	618,398 (9.51)	619,803 (9.54)	0.99
	Slovakia	-119,661 (-4.14)	-119,751 (-4.15)	1.00	98.7	96.5	2.2	239,317 (8.28)	239,500 (8.29)	1.00
	Slovenia	-40,919 (-3.66)	-40,990 (-3.66)	0.99	99.0	98.0	1.0	81,838 (7.32)	81,976 (7.33)	0.99
	Spain	-726,804 (-4.27)	-735,065 (-4.31)	0.95	99.4	95.1	4.3	1,453,592 (8.53)	1,470,129 (8.63)	0.95
	Sweden	-548,430 (-6.64)	-548,561 (-6.64)	1.00	99.5	98.7	0.7	1,096,819 (13.28)	1,097,120 (13.29)	1.00
	United Kingdom	-223,179 (-4.51)	-226,210 (-4.57)	0.94	97.9	96.8	1.1	446,097 (9.01)	452,417 (9.14)	0.94
		Average	-352,124 (-4.93)	-353,867 (-4.95)	0.98	99.10	96.9	2.2	704,221 (9.85)	707,733 (9.89)

Notes: The table shows the model comparison results using two in-sample prediction metrics (Soofi ID and AIC), and one out-of-sample prediction metric (10-fold cross validation). The numbers in the parenthesis in AIC and CV are their average value calculated by dividing AIC and CV by the number of data points. The result shows that the Lévy alpha-stable distribution is always preferred over the asymmetric exponential power distribution. The former has a higher Soofi-ID, a lower AIC, and a higher predictive power. Note that the higher the Soofi-ID and the lower the AIC, the better the fit. The \mathcal{L} Ratio for AIC is $\exp((AIC_{Lévy} - AIC_{AEP})/2N) \leq 0$ while \mathcal{L} Ratio for CV is $\exp((CV_{AEP} - CV_{Lévy})/N)$.

Table 10: Model Comparison: Country Average for Log Variables

Variable	Country	10-Fold Cross-Validation			Soofi-ID			AIC		
		Lévy	AEP	\mathcal{L} Ratio	Lévy	AEP	Diff	Lévy	AEP	\mathcal{L} Ratio
Log LP	Bulgaria	-93748 (-1.54)	-85720 (-1.41)	1.14	89.1	99.2	-10.1	188097 (3.08)	171440 (2.81)	1.15
	Czech Republic	-67534 (-1.36)	-66816 (-1.34)	1.01	99.6	99.5	0.0	135065 (2.71)	133632 (2.68)	1.01
	Estonia	-13379 (-1.31)	-13344 (-1.31)	1.00	99.4	99.1	0.3	26754 (2.63)	26686 (2.62)	1.00
	Finland	-27713 (-0.96)	-27393 (-0.95)	1.01	99.6	99.2	0.4	55422 (1.92)	54786 (1.89)	1.01
	France	-97052 (-0.93)	-96083 (-0.92)	1.01	99.8	99.5	0.3	194104 (1.86)	192166 (1.84)	1.01
	Germany	-39273 (-1.07)	-38637 (-1.05)	1.02	99.3	99.6	-0.2	78546 (2.14)	77274 (2.11)	1.02
	Hungary	-60699 (-1.35)	-60227 (-1.34)	1.01	99.6	99.3	0.4	121397 (2.71)	120454 (2.69)	1.01
	Italy	-266625 (-1.14)	-262964 (-1.12)	1.02	99.6	99.6	0.0	533254 (2.27)	525929 (2.24)	1.02
	Portugal	-143178 (-1.15)	-142290 (-1.14)	1.01	99.5	98.9	0.6	286362 (2.3)	284579 (2.28)	1.01
	Romania	-94643 (-1.47)	-93858 (-1.45)	1.01	99.6	99.4	0.2	189283 (2.93)	187716 (2.91)	1.01
	Slovakia	-38139 (-1.38)	-37836 (-1.37)	1.01	99.6	99.5	0.1	76278 (2.77)	75671 (2.75)	1.01
	Slovenia	-10257 (-0.89)	-10138 (-0.88)	1.01	97.9	98.3	-0.4	20521 (1.77)	20275 (1.75)	1.01
	Spain	-185753 (-1.04)	-184457 (-1.04)	1.01	99.8	99.3	0.5	371503 (2.09)	368912 (2.07)	1.01
	Sweden	-76135 (-0.88)	-75515 (-0.88)	1.01	99.7	99.1	0.6	152280 (1.77)	151030 (1.75)	1.01
	United Kingdom	-64751 (-1.29)	-63965 (-1.28)	1.02	99.6	99.6	-0.1	129500 (2.59)	127929 (2.56)	1.02
	Average	-85258 (-1.18)	-83950 (-1.17)	1.02	98.8	99.3	-0.5	170558 (2.37)	167899 (2.33)	1.02
Log LP Growth	Bulgaria	-48521 (-0.94)	-47506 (-0.92)	1.02	99.3	99.4	-0.1	97051 (1.88)	95013 (1.84)	1.02
	Czech Republic	-36889 (-0.84)	-35650 (-0.81)	1.03	98.6	99.5	-0.9	73766 (1.68)	71299 (1.62)	1.03
	Estonia									
	Finland	-10391 (-0.42)	-10152 (-0.41)	1.01	99.3	98.8	0.6	20749 (0.85)	20303 (0.83)	1.01
	France	-23202 (-0.27)	-22525 (-0.27)	1.01	99.4	99.2	0.2	46382 (0.55)	45050 (0.53)	1.01
	Germany	1790 (0.06)	1616 (0.05)	0.99	99.6	97.7	1.8	-3591 (-0.12)	-3234 (-0.1)	0.99
	Hungary	-34898 (-0.93)	-34154 (-0.91)	1.02	99.1	99.5	-0.5	69780 (1.87)	68308 (1.83)	1.02
	Italy	-119589 (-0.59)	-117000 (-0.57)	1.01	99.4	98.8	0.6	239230 (1.18)	233999 (1.15)	1.01
	Portugal	-69168 (-0.61)	-68067 (-0.6)	1.01	99.5	98.9	0.6	138348 (1.22)	136133 (1.2)	1.01
	Romania	-63695 (-1.16)	-62537 (-1.14)	1.02	99.2	99.5	-0.3	127407 (2.32)	125073 (2.27)	1.02
	Slovakia	-25240 (-0.97)	-24576 (-0.95)	1.03	97.9	98.7	-0.8	50516 (1.95)	49153 (1.9)	1.03
	Slovenia	-2892 (-0.26)	-2780 (-0.25)	1.01	99.0	99.0	-0.0	5778 (0.53)	5558 (0.51)	1.01
	Spain	-86277 (-0.54)	-83685 (-0.52)	1.02	99.3	99.4	-0.1	172519 (1.07)	167369 (1.04)	1.02
	Sweden	-25043 (-0.32)	-24519 (-0.31)	1.01	99.6	99.0	0.6	50039 (0.64)	49037 (0.62)	1.01
	United Kingdom	-17346 (-0.39)	-17174 (-0.38)	1.00	99.5	98.4	1.1	34695 (0.77)	34349 (0.77)	1.00
	Average	-40097 (-0.58)	-39194 (-0.57)	1.01	99.2	99.0	0.2	80191 (1.17)	78386 (1.14)	1.01

Notes: The table shows the model comparison results for log labor productivity and its growth rate using two in-sample prediction metrics (Soofi ID and AIC), and one out-of-sample prediction metric (10-fold cross validation). The numbers in the parenthesis in AIC and CV are their average value calculated by dividing AIC and CV by the number of data points. The result shows that the AEP is generally preferred over Lévy alpha-stable distribution for both of the variables. However, the difference of the model performance is very small compared to the result for non-log variables.