On the modeling of size distributions when technologies are complex☆

Jakub Growiec *

Warsaw School of Economics, Department of Quantitative Economics, and Narodowy Bank Polski, Poland

A B S T R A C T

The study considers a stochastic R&D process where the invented production technologies consist of a large number n of complementary components. The degree of complementarity is captured by the elasticity of substitution of the CES aggregator function. Drawing from the Central Limit Theorem and the Extreme Value Theory we find, under very general assumptions, that the cross-sectional distributions of technological productivity are well-approximated either by the lognormal, Weibull, or a novel "CES/Normal" distribution, depending on the underlying elasticity of substitution between technology components. We find the tail of the "CES/Normal" distribution to be fatter than the Weibull tail but qualitatively thinner than the Pareto (power law) one. We also numerically assess the rate of convergence of the true technological productivity distribution to the theoretical limit with n as fast in the body but slow in the tail.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Most technologies used nowadays are complex in the sense that the production processes (and products themselves) consist of a large number of components which might interact with each other in complementary ways (e.g. Kremer, 1993; Blanchard and Kremer, 1997; Jones, 2011). Based on this insight, the current paper assumes that the total productivity of any given technology is functionally dependent on the individual productivities of its n components as well as the elasticity of substitution between them, σ. This functional relationship is captured by the CES aggregator function. The stochastic R&D process which invents new complex technologies is in turn assumed to consist in drawing productivities of the components from certain predefined probability distributions (Jones, 2005; Growiec, 2008a,b, 2013).

Based on this set of assumptions, we obtain surprisingly general results regarding the implied cross-sectional distributions of technological productivity. Namely, drawing from the Central Limit Theorem and the Extreme Value Theory, we find that if the number of components of a technology, n, is sufficiently large, these distributions should be well approximated either by:

(i) the lognormal distribution — in the case of unitary elasticity of substitution between the components ($\sigma = 1$ as in Kremer, 1993);
(ii) the Weibull distribution — in the case of perfect complementarity between the components (the "weakest link" assumption, $\sigma = 0$ as in Growiec, 2013),
(iii) the Gaussian distribution — in the (empirically very unlikely) case of perfect substitutability between the components ($\sigma \to \infty$),
(iv) a novel "CES/Normal" distribution — in any intermediate CES case, parametrized by the elasticity of substitution between the components ($\sigma > 0$, $\sigma \neq 1$).

We proceed to investigate the properties of the right tail of the "CES/Normal" distribution. Computing its Pareto as well as Weibull tail index confirms that, if technology components are gross complements but are not perfectly complementary ($\sigma \in (0, 1)$), the tail of this distribution decays faster than the tail of any Pareto distribution (i.e., it does not follow a power law) but slower than the tail of any Weibull distribution.

This tail result is interesting because although the prevalence of fat-tailed distributions has been documented and thoroughly discussed for firm sizes, along with a wide array of other

* Correspondence to: Katedra Ekonomii Ilościowej SGH, al. Niepodległości 162, 02-554 Warszawa, Poland. Fax: +48 22 564 9322.
    E-mail address: jakub.growiec@sgh.waw.pl.

phenomena in economics and finance,[1] clearly *most* economic variables do not have this property.[2] The distribution of technological productivity, with which we deal here and which has not (to our knowledge) been studied in the empirical literature so far, is in turn one of the important primitives for the firm size distribution. Hence our finding that, from the theoretical point of view, technological productivity distributions should *not* be expected to be fat-tailed, indicates that the apparent emergence of power law tails in firm size distributions[3] must be driven by other phenomena, such as, e.g., endogenous technology choice (Jones, 2005; Growiec, 2008a,b), resource misallocation (Hsieh and Klenow, 2009; Jones, 2011), or aggregation across multi-product firms (Fu et al., 2005; Growiec et al., 2008).

Our aforementioned theoretical contribution to the literature is supplemented with a series of numerical simulations, allowing us to approximate the rate of convergence of the true technological productivity distribution to the theoretical limit with n. We identify this rate to be fast in the body of the distribution but slow in the tails which capture rare events. We also numerically assess the dependence of the limiting "CES/Normal" distribution on the degree of complementarity between the technology components, σ.

Potential empirical applications of the theoretical result, reaching beyond the scope of the current paper, include providing answers to the following research questions:

- Does the "CES/Normal" distribution derived here (Eq. (15)) fit the data on firm sizes, sales, R&D spending, etc.? What is the implied value of σ?
- Do industries differ in terms of their technology complexity as captured by n?
- Do industries differ in terms of the complementarity of technology components as captured by σ?
- How do firms' optimal technology choices and production function aggregation enter the picture? What are the implications for the shape of the aggregate production function?

The remainder of this paper is structured as follows. Section 2 sets up the model and provides the principal analytical results. Section 3 presents the numerical results. Section 4 concludes.

2. The model

2.1. Distributions of complex technologies

The point of departure of the current model is the assumption that technologies, invented within the R&D process, are inherently complex and consist of a large number of complementary components. Formally, this can be written down in the following way.

Assumption 1. The R&D process determines the productivity of any newly invented technology Y as a constant elasticity of substitution (CES) aggregate over $n \in \mathbb{N}$ independent draws X_i, $i =$ 1, ..., n, from the elementary idea distribution $\mathcal{F}$:

$$Y = \begin{cases} \min\{X_i\}_{i=1}^n, & \theta = -\infty, \\ \left(\dfrac{1}{n} \sum_{i=1}^n X_i^\theta \right)^{1/\theta}, & \theta \in (-\infty, 0) \cup (0, 1], \\ \prod_{i=1}^n X_i^{1/n}, & \theta = 0. \end{cases} \quad (1)$$

The elementary distribution $\mathcal{F}$ is assumed to have a positive density on $[w, v]$ and zero density otherwise (where $w \geq 0$ and $v > w$ can be infinite). For the case $\theta = -\infty$, it is also assumed to satisfy the condition of a regularly varying lower tail (Leadbetter et al., 1983):

$$\lim_{p \to 0_+} \frac{\mathcal{F}(w + px)}{\mathcal{F}(w + p)} = x^\alpha \quad (2)$$

for all $x > 0$ and a certain $\alpha > 0$. For the cases $\theta \in (-\infty, 0) \cup (0, 1]$, it is assumed that $EX_i^\theta < \infty$ and $D^2(X_i^\theta) < \infty$. For the case $\theta = 0$, it is assumed that $E \ln X_i < \infty$ and $D^2(\ln X_i) < \infty$.

The parameter n in the above assumption captures the number of constituent components of any given (composite) technology, and thus measures the complexity of any state-of-the-art technology. The substitutability parameter θ is related to the elasticity of substitution σ via $\theta = \frac{\sigma-1}{\sigma}$, or $\sigma = \frac{1}{1-\theta}$. The case $\theta < 0$ captures the case where the components of technologies are gross complements ($\sigma \in [0, 1)$), whereas $\theta \in (0, 1]$ implies that they are gross substitutes ($\sigma > 1$).

It should be noted at this point that, as argued repeatedly by Kremer (1993), Jones (2011) and Growiec (2013), the gross complementarity case is much more likely to provide an adequate description of real-world production processes than the gross subsitutability case. The example of the explosion of the space shuttle *Challenger* due to a failure of an inexpensive O-ring, put forward by Kremer (1993), is perhaps the best possible illustration of the potentially complementary character of components of complex technologies.

More precisely, the *minimum* case (a Leontief function) reflects the extreme case where technology components are *perfectly* complementary, and thus the actual productivity of a complex idea is determined by the productivity of its "weakest link" (or "bottleneck"). This case was assumed in the related contribution by Growiec (2013). Although likely, this case need not hold exactly in reality, since certain deficiencies of design can often be covered by advantages in different respects. The more general CES case captures exactly this possibility (see also Klump et al., 2012).

The limiting Cobb–Douglas case ($\theta = 0$) is the threshold case delineating gross complementarity from gross substitutability. As shown by Kremer (1993), this case is already quite illustrative of effects of complementarity between components of technologies.

Although technical in nature, restriction (2) imposed on elementary probability distributions $\mathcal{F}$ can also be interpreted in economic terms. First, the support of the distribution must be bounded from below by w, which means researchers are not allowed to draw infinitely "bad" technologies (zero is a natural lower bound). This rules out distributions defined on the whole $\mathbb{R}$ such as the Gaussian. Second, the pdf of the distribution $\mathcal{F}$ cannot increase smoothly from zero at w; there must be a jump. This means that the probability of getting a draw which is "as bad as it gets" cannot be negligible, and this rules out a few more candidate distributions such as the lognormal or the Fréchet. Third, the lowest possible value of the random variable cannot be an isolated atom, which rules out all discrete distributions such as the two-point distribution, the binomial, negative binomial, and Poisson. Yet, the set of distributions satisfying (2) is still reasonably

---

[1] Including firm sales, firms' R&D spending, asset returns, and city sizes (Sutton, 1998; Gabaix, 1999; Axtell, 2001; Eeckhout, 2004; Clauset et al., 2009; Gabaix, 2009).

[2] In line with the ubiquitous assumption of Gaussian error terms in econometrics.

[3] The existence of power-law tails in empirical distributions has also been contested by some authors, see e.g. Stanley et al. (1995) or Bee et al. (2013).

large. It includes, among others, the frequently assumed Pareto, uniform, truncated Gaussian, and Weibull distributions (Growiec, 2013).

2.2. Normalized productivity

Based on the aforementioned setup, our objective is to derive the limiting distribution of Y when the number of technology components $n \to \infty$. For analytical convenience, we denote $\mu_\theta = EX_i^\theta$ and $\sigma_\theta = \sqrt{D^2(X_i^\theta)} = \sqrt{EX_i^{2\theta} - (EX_i^\theta)^2}$, for any $\theta \in (-\infty, 0) \cup (0, 1]$. For the limiting case of $\theta = 0$, we define $\mu_0 = E(\ln X_i)$ and $\sigma_0 = E(\ln X_i)^2 - (E \ln X_i)^2$.

The first observation is that by the Law of Large Numbers, Y itself will almost surely converge to a degenerate distribution with $n \to \infty$. Indeed, it is straightforward to see that in the perfect complementarity case ($\theta = -\infty$), $\min\{X_i\}_{i=1}^n \xrightarrow{a.s.} w$ where w is the lower bound of the underlying distribution $\mathcal{F}$. Accordingly, in the CES case ($\theta \in (-\infty, 0) \cup (0, 1]$), $\left(\frac{1}{n}\sum_{i=1}^n X_i^\theta\right)^{1/\theta} \xrightarrow{a.s.} \mu_\theta^{1/\theta}$. Finally, in the multiplicative case ($\theta = 0$), $\prod_{i=1}^n X_i^{1/n} \xrightarrow{a.s.} e^{\mu_0}$.

Hence, to obtain meaningful results related to productivity *distributions*, more subtlety is needed. One ought to normalize the distributions with finite n appropriately in order to make use of the Central Limit Theorem as well as Extreme Value Theory. In this regard, we shall base our considerations upon the following definition:

Definition 1. The *normalized productivity* of any technology Y of complexity $n \in \mathbb{N}$ and component complementarity $\theta = -\infty$ or $\theta = 0$ is given by $\tilde{Y}(n)$:

$$
\tilde{Y}(n) = \begin{cases} \min\left\{\dfrac{X_i - w}{\mathcal{F}^{-1}\left(\frac{1}{n}\right) - w}\right\}_{i=1}^n, & \theta = -\infty, \\[3mm] \left(\displaystyle\prod_{i=1}^n \dfrac{X_i}{e^{\mu_0}}\right)^{\frac{1}{\sigma_0\sqrt{n}}}, & \theta = 0. \end{cases} \tag{3}
$$

In the case $\theta \in (-\infty, 0) \cup (0, 1]$ we define the *shifted normalized productivity* as

$$
\tilde{Y}(n, \omega) = \left(\frac{1}{\sqrt{n}}\sum_{i=1}^n \left(\frac{X_i^\theta - \mu_\theta}{\sigma_\theta}\right) + \omega\right)^{1/\theta}, \tag{4}
$$

where $\omega > 0$ is a shift parameter. The normalized productivity $\tilde{Y}(n)$ is defined as

$$
\tilde{Y}(n) = \lim_{\omega \to \infty} \frac{\tilde{Y}(n, \omega) - E\tilde{Y}(n, \omega)}{D\tilde{Y}(n, \omega)}. \tag{5}
$$

The shift parameter $\omega > 0$ does not play any role in the economic interpretation of our results, but considering it "sufficiently large" will be a useful tool in the proofs of our propositions.

Note that this normalization procedure, while necessary in the subsequent derivations, does not have a direct economic interpretation because the underlying stochastic R&D process is assumed to consist in inventing technologies with a fixed number of components n.

2.3. Limiting productivity distributions

Letting the technology complexity n be arbitrarily large, we obtain the following results:

Proposition 1. *If Assumption 1 holds with $\theta = -\infty$ ($\sigma = 0$), then as $n \to \infty$, $\tilde{Y}(n)$ converges in distribution to the standard Weibull distribution with the shape parameter α:*

$$
P(\tilde{Y}(n) \geq x) = [1 - \mathcal{F}(xp_n + w)]^n \xrightarrow{d} e^{-x^\alpha}, \tag{6}
$$

where $w = \inf\{x \in \mathbb{R} : \mathcal{F}(x) > 0\}$ and $p_n = \mathcal{F}^{-1}\left(\frac{1}{n}\right) - w$.

Proposition 2. *If Assumption 1 holds with $\theta = 0$ ($\sigma = 1$), then as $n \to \infty$, $\tilde{Y}(n)$ converges in distribution to the standard lognormal distribution:*

$$
P(\tilde{Y}(n) \geq x) \xrightarrow{d} 1 - \Phi(\ln x). \tag{7}
$$

Proposition 3. *If Assumption 1 holds with $\theta \in (-\infty, 0) \cup (0, 1]$ ($\sigma \in (0, 1) \cup (1, +\infty)$), then as $n \to \infty$, $\tilde{Y}(n, \omega)$ converges in distribution to the "CES/Normal" distribution with complementary cdf:*

$$
P(\tilde{Y}(n, \omega) \geq x) \xrightarrow{d} \Phi(x^\theta - \omega), \quad \theta \in (-\infty, 0), \tag{8}
$$

$$
P(\tilde{Y}(n, \omega) \geq x) \xrightarrow{d} 1 - \Phi(x^\theta - \omega), \quad \theta \in (0, 1], \tag{9}
$$

and thus the following pdf:

$$
g(x; \omega) = \frac{|\theta|}{\sqrt{2\pi}} x^{\theta-1} e^{-\frac{(x^\theta - \omega)^2}{2}}, \quad x > 0. \tag{10}
$$

Hence, the class of "CES/Normal" distributions encompasses the Gaussian distribution for the limiting case $\theta = 1$ where the technology components are perfectly substitutable.

Proof of Propositions 1–3. For the case $\theta = -\infty$ the proposition follows directly from the Fisher–Tippett–Gnedenko extreme value theorem applied to the distribution $\mathcal{F}$ (Theorem 1.1.3 in de Haan and Ferreira, 2006), rephrased so that it captures the minimum instead of maximum). From the theorem specifying the domain of attraction of the Weibull distribution (Theorem 1.2.1 in de Haan and Ferreira, 2006; Section 1.3 in Kotz and Nadarajah, 2000), we obtain the necessary and sufficient conditions for the complementarity mechanism to work. The implied parameter α is found to be unitary for a wide range of distributions $\mathcal{F}$ (Growiec, 2013), including the Pareto, uniform, and truncated Gaussian distribution.

With $\theta = 0$, by the Central Limit Theorem the distribution of $\tilde{Y}(n)$ converges in distribution with $n \to \infty$ to the lognormal distribution:

$$
\ln \tilde{Y}(n) = \sqrt{n}\left(\frac{\frac{1}{n}\sum_{i=1}^n \ln X_i - \mu_0}{\sigma_0}\right) \xrightarrow{d} N(0, 1) \tag{11}
$$

and hence

$$
\tilde{Y}(n) = \left(\prod_{i=1}^n \frac{X_i}{e^{\mu_0}}\right)^{\frac{1}{\sigma_0\sqrt{n}}} \xrightarrow{d} \log N(0, 1). \tag{12}
$$

For the case $\theta \in (-\infty, 0) \cup (0, 1]$, we may use the Central Limit Theorem again, obtaining:

$$
\left(\frac{1}{\sqrt{n}}\sum_{i=1}^n \left(\frac{X_i^\theta - \mu_\theta}{\sigma_\theta}\right) + \omega\right) \xrightarrow{d} N(\omega, 1). \tag{13}
$$

Note that by assumption, all $X_i \geq 0$ and thus $\frac{1}{\sqrt{n}} \sum_{i=1}^{n} \left(\frac{X_i^\theta - \mu_\theta}{\sigma_\theta} \right)$ $\geq -\frac{\sqrt{n} \mu_\theta}{\sigma_\theta}$ for every finite n. There does not exist any $\omega > 0$ such that $\frac{1}{\sqrt{n}} \sum_{i=1}^{n} \left(\frac{X_i^\theta - \mu_\theta}{\sigma_\theta} \right) \geq -\omega$ uniformly for all $n \in \mathbb{N}$, though.[4]

The limiting distribution takes the following form. For $y > 0$ and with $\theta < 0$, $\tilde{Y}(n, \omega)$ converges in distribution to $\bar{Y}(\omega)$ with cdf:

$$\mathcal{G}(y; \omega) = \lim_{n \to \infty} P(\tilde{Y}(n, \omega) \leq y) = 1 - \bar{\Phi}(y^\theta - \omega), \qquad (14)$$

where the cdf of the truncated normal distribution $\bar{\Phi}(y^\theta - \omega) = \frac{\Phi(y^\theta - \omega) - \Phi(-\omega)}{1 - \Phi(-\omega)} \approx \Phi(y^\theta - \omega)$ if ω is sufficiently large and thus the truncation hurts a negligible part of the distribution.

Upon differentiation, we obtain the following pdf of the limiting "CES/Normal" distribution, parametrized by $\theta < 0$, $\mu_\theta > 0$ and $\sigma_\theta > 0$:

$$g(y; \omega) = \frac{-\theta}{\sqrt{2\pi}} \frac{y^{\theta-1}}{1 - \Phi(-\omega)} e^{-\frac{(y^\theta - \omega)^2}{2}} \approx \frac{-\theta}{\sqrt{2\pi}} y^{\theta-1} e^{-\frac{(y^\theta - \omega)^2}{2}}. \quad (15)$$

The final approximation becomes arbitrarily good for sufficiently large ω.

Conversely, for $y > 0$ and with $\theta \in (0, 1]$, $\tilde{Y}(n, \omega)$ converges in distribution to $\bar{Y}(\omega)$ with cdf:

$$\mathcal{G}(y; \omega) = \lim_{n \to \infty} P(\tilde{Y}(n, \omega) \leq y) = \bar{\Phi}(y^\theta - \omega) \approx \Phi(y^\theta - \omega). \quad (16)$$

Upon differentiation, we obtain the following pdf of the limiting "CES/Normal" distribution, parametrized by $\theta \in (0, 1]$, $\mu_\theta > 0$ and $\sigma_\theta > 0$:

$$g(y; \omega) = \frac{\theta}{\sqrt{2\pi}} \frac{y^{\theta-1}}{1 - \Phi(-\omega)} e^{-\frac{(y^\theta - \omega)^2}{2}} \approx \frac{\theta}{\sqrt{2\pi}} y^{\theta-1} e^{-\frac{(y^\theta - \omega)^2}{2}}. \quad (17)$$

Taking the limit of $\omega \to \infty$, we obtain the "CES/Normal distribution" of normalized technological productivity:

$$\tilde{Y} = \lim_{\omega \to \infty} \frac{\bar{Y}(\omega) - E\bar{Y}(\omega)}{D\bar{Y}(\omega)}. \quad \blacksquare \qquad (18)$$

2.4. Tail properties

To assess the right tail properties of the distributions obtained above, we shall compare their tail decay to the Pareto and Weibull distributions. Such comparisons will allow us to understand if, under any parameterization, the "CES/Normal" distribution allows for approximately power-law decay (Pareto distribution; fat tails), or Weibull-type exponential decay (thin tails).

Formally, the Pareto tail index ξ (capturing power law decay) is defined according to the following formula:

$$\lim_{x \to +\infty} \frac{1 - \mathcal{F}(\lambda x)}{1 - \mathcal{F}(x)} = \lambda^{-\frac{1}{\xi}}, \quad \lambda > 0. \qquad (19)$$

Table 1
Tail index of the technological productivity distribution.
Source: Own computations.

Case	Pareto ξ	Weibull ψ
Weibull ($\theta = -\infty$)	0	$1/\alpha$
CES/Normal (gross complements, $\theta < 0$)	0	$+\infty$
Lognormal ($\theta = 0$)	0	$+\infty$
CES/Normal (gross substitutes, $\theta \in (0, 1]$)	0	$1/2\theta$
Pareto(ϕ)	$1/\phi$	$+\infty$

By construction, for any Pareto(ϕ) distribution, the tail index is equal to $\xi = 1/\phi$. In particular, for the celebrated Zipf's law (e.g. Gabaix, 1999; Axtell, 2001), $\phi = 1$ and thus $\xi = 1$.

The Weibull tail index ψ (capturing exponential decay), on the other hand, is obtained as follows:

$$\lim_{x \to +\infty} \frac{\ln(1 - \mathcal{F}(\lambda x))}{\ln(1 - \mathcal{F}(x))} = \lambda^{\frac{1}{\psi}}, \quad \lambda > 0. \qquad (20)$$

Again by construction, for any Weibull(α) distribution, this tail index is equal to $\psi = 1/\alpha$. In particular, the exponential distribution has $\alpha = 1$ and thus $\psi = 1$.

We find that for $\theta \in (0, 1]$ (when components of technologies are gross substitutes) the "CES/Normal" distribution decays asymptotically to the thin-tailed Weibull distribution, whereas for $\theta < 0$ (when these components are gross complements), it decays faster than the Weibull distribution but slower than a power law (see Table 1).[5]

Please note that, regarding both tail indexes, a zero result indicates that a given distribution decays faster (has a thinner tail) than any Pareto or Weibull distribution, respectively. Conversely, an infinite limit implies that a given distribution decays slower. For example, since the Pareto(ϕ) decays slower than the Weibull, the Weibull distribution has a zero Pareto tail index and the Pareto distribution has an infinite Weibull tail index. The lognormal distribution decays faster than the Pareto but slower than the Weibull.

For the novel "CES/Normal" distribution we find that gross complementarity of components of technologies gives rise to relatively fatter tails than gross substitutability, but these tails are nevertheless never sufficiently fat to generate a power law. Moreover, in the empirically less likely case of gross substitutability of technological components, the tail of the resulting distribution is just as thin as the Weibull one.

Hence, if technologies are sufficiently complex ($n \to \infty$), technological productivity distributions should not be expected to be fat-tailed. This implies that fat tails in firm size distributions are not inherited from the underlying distributions of technological productivity and their apparent emergence must be due to other reasons.

3. Numerical results

The most important advantage of above analytical results is that they provide theoretical limits for the distributions of complex technologies, regardless of the underlying distribution of technology components $\mathcal{F}$. Unfortunately, these limits are exactly correct only if the technologies are *infinitely complex*, though. It is therefore of great importance to assess the pace of convergence of

---

[4] In the unlikely case where the latter inequality fails, we replace $\frac{1}{\sqrt{n}} \sum_{i=1}^{n} \left(\frac{X_i^\theta - \mu_\theta}{\sigma_\theta} \right) = 1$ to keep our results in the space of real numbers. Observe, however, that if ω is sufficiently large, (a) the probability of such failure is very low for all n because $X_i \geq 0$ are independently drawn from a distribution which has positive density over a non-degenerate interval, and (b) for any $n \in \mathbb{N}$, this probability tends to 0 as $\omega \to \infty$. Hence such a replacement does not affect the subsequently derived shape of the limiting distribution.

[5] Curiously, for any finite truncation point ω with $\theta < 0$, the sheer presence of truncation generates power law decay of the limiting distribution, with an exponent $-1/\theta$. The mechanism driving this result is akin to the "reflective lower bound" model of power laws, cf. Gabaix (1999). Since the truncation is an artifact of the method of proof and not the underlying economic model, we set this case aside. Details are available upon request.
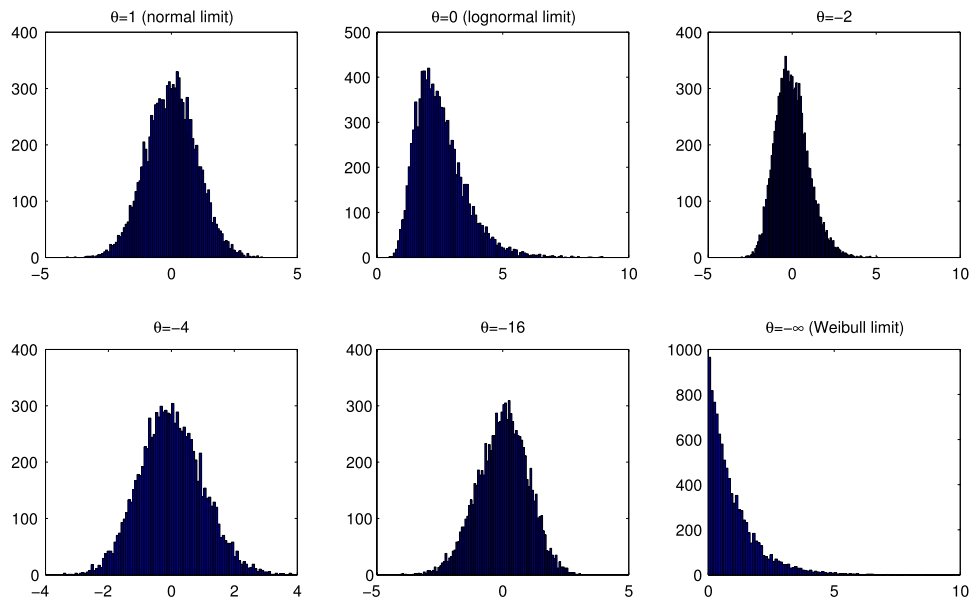
Fig. 1. Histograms of simulated $\tilde{Y}(n)$.
Notes. Parameter values used to produce this figure: $n = 1000$, $a = 0.5$, $b = 2$. When computing Eq. (5), the shift parameter ω was set at $\omega = 8 \cdot \mu_\theta$.
Source: Own computations.

true distributions of technological productivity to the limiting ones with the number of components, n, and to estimate the magnitude of departures from the theoretical limit which could be expected if n is in fact finite. To this end, we have carried out a series of numerical computations.

At this point, observe that not only the number of components n but also the assumed shape of the sampling distribution $\mathcal{F}$ has an impact of the shape of the resultant aggregate technological productivity. Moreover, assumptions on $\mathcal{F}$ can also affect the pace of convergence of the resulting distribution to its theoretical limit with n. Assessing these impacts quantitatively is left for further research.

Another interesting issue to be handled here is the dependence of the limiting "CES/Normal" distribution on the complementarity parameter θ (or equivalently, the elasticity of substitution, σ). This will be assessed numerically as well. The current section will first describe our numerical framework and then move on to an outline of the results.

3.1. Generating the distribution of Y

The preliminary step of our numerical exercises consists in generating a sample of n units, randomly and independently drawn from the sampling distribution $\mathcal{F}$. For the sake of simplicity, we assume it to be a *uniform* distribution defined on an interval in the positive domain (which is a particular instance of a distribution satisfying Assumption 1):

$$X_i \sim U[a, b], \quad b > a > 0, \ i = 1, 2, \ldots, n. \tag{21}$$

Next, we compute the normalized CES aggregate of these random draws according to Eq. (3) if $\theta = 0$ or $\theta = -\infty$, and Eq. (5) otherwise. Apart from the two aforementioned limiting cases, we consider four arbitrary values of the complementarity parameter θ. For a fixed sample size of n, we repeat this procedure $m = 10000$ times and plot the empirical histogram of $\tilde{Y}(n)$. The histograms of the generated variables are presented in Fig. 1. These histograms (with $B = 100$ bins) are then transformed into empirical pdfs.

Fig. 1 does not yet confirm the trend of increasing skewness of the empirical distribution when θ declines towards $-\infty$, a result which can be inferred from looking at the pdf of the limiting

"CES/Normal" distribution. This is because the pace of convergence of the tail of the distribution with n to the asymptotical limit tends to fall as θ declines towards $-\infty$.

3.2. The lognormal and Weibull limits

The second step of our numerical exercise consists in assessing the convergence to the theoretical lognormal limit for the case $\theta = 0$ as well as the theoretical Weibull limit obtained when $\theta = -\infty$. We see that the theoretical distributions indeed align with the simulated data almost perfectly when $n = 1000$.[6] These results are contained in Figs. 2–3.

3.3. The "CES/Normal" limit

Having sorted out the two limiting cases, we shall now address the pace of convergence to the theoretical formula for the limiting "CES/Normal" distribution under intermediate values of the complementarity parameter, $\theta \in (-\infty, 0) \cup (0, 1]$. As it can be seen on Fig. 4, the theoretical "CES/Normal" limit is almost entirely converged upon if $n = 1000$; it is clear that none of the seemingly similar (and more generously parametrized) distributions can be fitted to the simulated data equally well.

This numerical exercise confirms that (a) the Weibull distribution misses the shape of the pdf completely when θ is finite, and (b) all other considered distributions tend to underestimate the probabilities of tail events. In fact, the actual limiting distribution is much more skewed than the estimated pdfs.

Regarding Fig. 4, please note the following difference between the "CES/Normal" distribution and the "CES/Normal Free" case. The first one takes the (known) theoretical values of mean and variance $(\mu_\theta, \sigma_\theta)$ as well as θ itself as given, whereas the latter takes them

---

[6] Nonlinear least squares fit of the Generalized Gamma distribution (a three-parameter class of distributions whose parametric form is somewhat similar to the "CES/Normal" distribution) is included for comparison. Clearly, the fit of this distribution to our simulated data is quite good, but not as good as of the theoretically derived "CES/Normal" distribution.
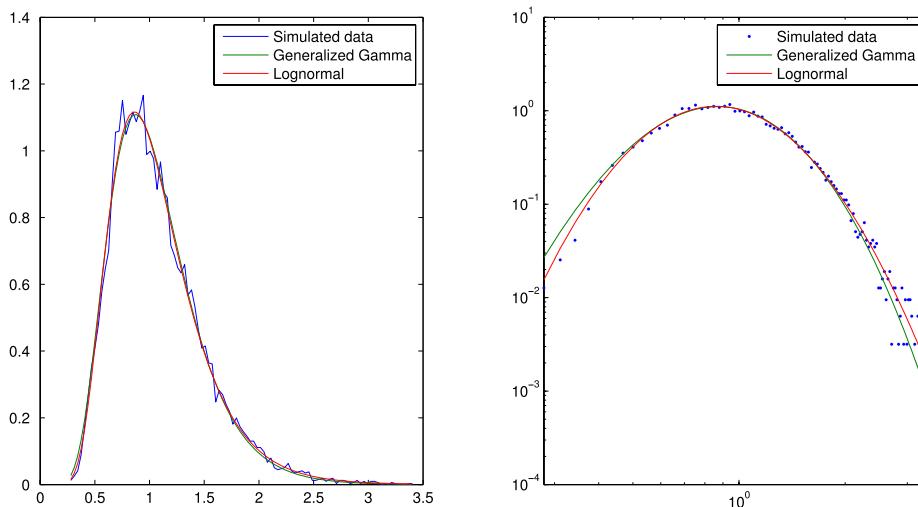
Fig. 2. Simulated pdf of $\bar{Y}(n)$ vs. the lognormal limit for $n \rightarrow \infty$. The right panel zooms on the right tail in log–log coordinates.
Note. Parameter values used to produce this figure: $n = 1000$, $a = 0.5$, $b = 2$.
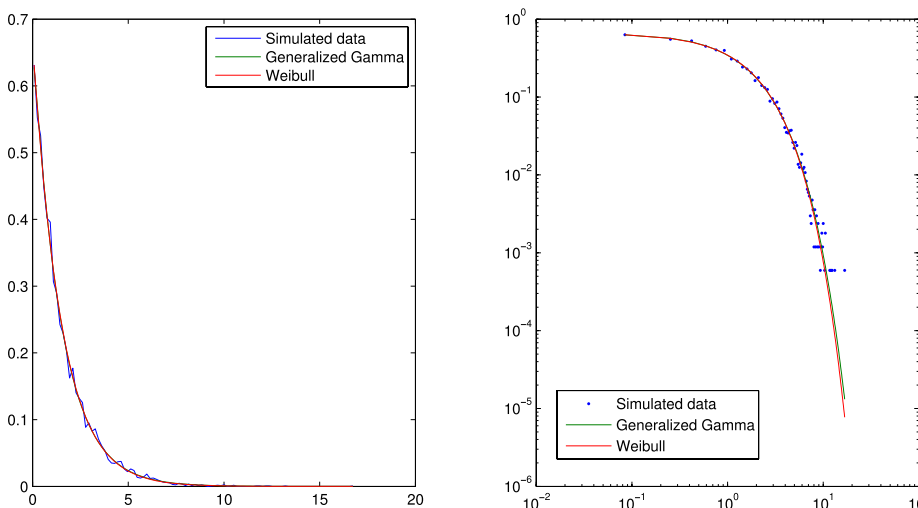*Source:* Own computations.

Fig. 3. Simulated pdf of $\bar{Y}(n)$ vs. the Weibull limit for $n \rightarrow \infty$. The right panel zooms on the right tail in log–log coordinates.
Note. Parameter values used to produce this figure: $n = 1000$, $a = 0.5$, $b = 2$.
Source: Own computations.

as free parameters to be estimated from the (simulated) data by nonlinear least squares. We see that improving the fit in the body of the distribution of the finite CES aggregate ($n = 1000 < \infty$) heavily compromises the quality of fit in the right tail.

3.4. Pace of convergence as $n \rightarrow \infty$

The next step consists in repeating the numerical experiment for a fixed value of $\theta \in (-\infty, 0) \cup (0, 1]$ but various sample sizes n to assess the pace of convergence of the resultant distribution to the theoretical "CES/Normal" limiting distribution (Eq. (15)).

In Fig. 5 we observe that as n increases, the resulting distribution gradually evolves from the assumed uniform distribution $\mathcal{F}$ of X_i to the limiting "CES/Normal" distribution, derived in the previous section of the current study. In the body of the distribution, convergence is fast and is largely done already for $n = 16$. The tails of the distribution are however much thinner for small n than in the limiting distribution. Even for $n = 2000$, although the fit in the body is already perfect, no observations have been found for tail

events exceeding 6.[7] This mirrors the known fact that tails of a distribution need much more time to take their final shape, because they are by definition capturing rare events. Therefore as long as the tails of the distribution $\mathcal{F}$ do not exactly coincide with the limiting distribution, convergence will be much faster in the body of the distribution than in its tails.

3.5. Dependence of the "CES/Normal" distribution on θ

The final step of the numerical exercise is to illustrate the dependence of the limiting "CES/Normal" distribution on the complementarity parameter θ, with special reference to the right tail. As illustrated by Fig. 6 which concentrates on the case where technology components are gross complements ($\theta < 0$), we find that the greater is the degree of complementarity (lower θ), the

---

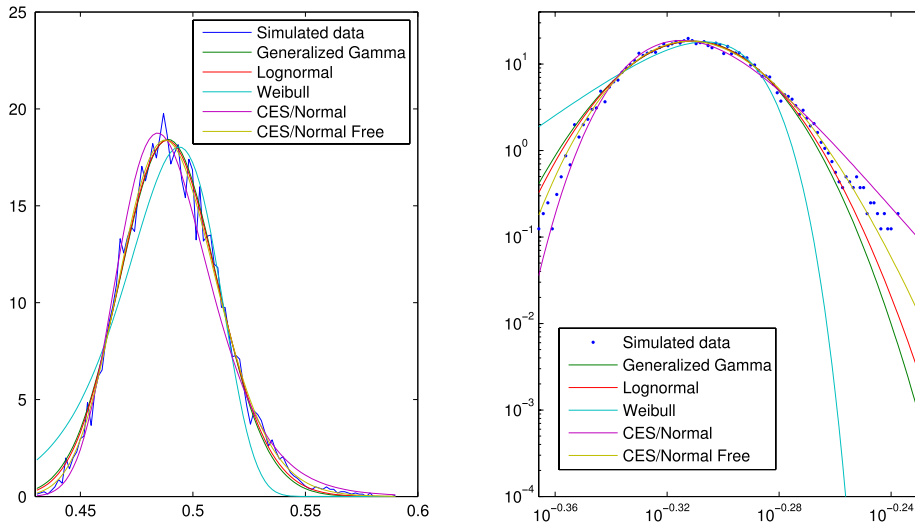[7] Note that the support of the underlying distribution $\mathcal{F}$ is bounded in this numerical exercise.

Fig. 4. Simulated pdf of $\tilde{Y}(n)$ vs. the "CES/Normal" limit (Eq. (15)) for $n \to \infty$. The right panel zooms on the right tail in log–log coordinates.
Note. Parameter values used to produce this figure: $n = 1000$, $a = 0.5$, $b = 2$, $\theta = -4$.
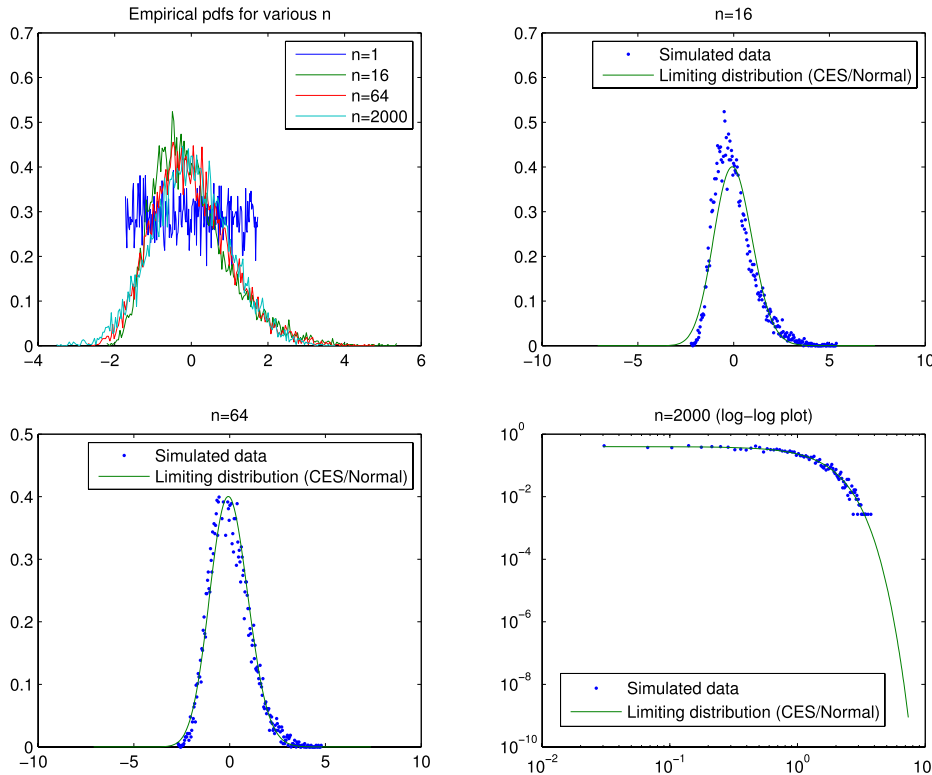*Source:* Own computations.

Fig. 5. Convergence of $\tilde{Y}(n)$ to the CES/Normal limit (Eq. (15)) for different values of n. Assumed parameter values: $a = 0.5$, $b = 2$, $\theta = -4$.
Note. Parameter values used to produce this figure: $a = 0.5$, $b = 2$, $\theta = -4$.
Source: Own computations.

fatter the tail of the limiting distribution.[8] Yet, in line with our theoretical results, power-law decay is generally ruled out. None of the lines in the right panel of Fig. 6 is asymptotic to a straight line.

---

[8] Please note that the tail of the distribution with greatest complementarity is probably misspecified due to numerical error, likely to bias the results particularly strongly when we take numbers to so large negative powers as -100.

4. Conclusion

The objective of the current paper has been to identify the shape of cross-sectional distributions of technological productivity in a world where technologies are complex and consist of a large number of complementary components. Drawing from the Central Limit Theorem and Extreme Value Theory, we find that if the number of components of a technology, n, is sufficiently large, these distributions should be well approximated either by:

Fig. 6. Convergence of $\tilde{Y}(n)$ to the CES/Normal limit (Eq. (15)) for different values of θ. The right panel zooms on the right tail in log–log coordinates.
Note. Parameter values used to produce this figure: $a = 0.5$, $b = 2$.
Source: Own computations.

(i) the lognormal distribution − in the case of unitary elasticity of substitution between the components ($\sigma = 1, \theta = 0$);

(ii) the Weibull distribution − in the case of perfect complementarity between the components (the "weakest link" assumption, $\sigma = 0, \theta = -\infty$),

(iii) the Gaussian distribution − in the (empirically very unlikely) case of perfect substitutability between the components ($\sigma = +\infty, \theta = 1$),

(iv) a novel "CES/Normal" distribution − in any intermediate CES case, parametrized by the elasticity of substitution between the components ($\sigma > 0, \sigma \neq 1$ or equivalently $\theta < 1, \theta \neq 0$).

We also find that, as long as technology components are gross complements but not perfectly complementary ($\sigma \in (0, 1)$), the "CES/Normal" distribution has a tail which is fatter than the Weibull one but is still not a *fat* tail: it decays unambiguously faster than the power law. This implies that fat tails in firm size distributions are not inherited from the underlying distributions of technological productivity and their apparent emergence must be due to other reasons.

Our theoretical contribution to the literature has been supplemented with a series of numerical simulations, allowing us to approximate the rate of convergence of the true distribution to the theoretical limit with n. We have also numerically assessed the dependence of the limiting "CES/Normal" distribution on the degree of complementarity between the technology components, σ (or equivalently, θ).

What could still be done as an extension to the current study, is to:

• discuss the implied moments of the limiting distributions and check if there is convergence in distribution when $\theta \to 0$ or $\theta \to -\infty$,

• provide approximate theoretical results on the pace of convergence of n-unit technologies to the "CES/Normal", Weibull or lognormal limit, and most importantly to

• verify the empirical relevance of "CES/Normal" distributions. Do we find it in data on firm sizes, sales, R&D spending, etc.?

What is the implied degree of complementarity between the technology components, σ?

We leave these interesting questions for further research.

References

Axtell, R., 2001. Zipf distribution of US firm sizes. Science 293 (5536), 1818–1820.

Bee, M., Riccaboni, M., Schiavo, S., 2013. The size distribution of US cities: Not Pareto, even in the tail. Econom. Lett. 120 (2), 232–237.

Blanchard, O., Kremer, M., 1997. Disorganization. Q. J. Econ. 112 (4), 1091–1126.

Clauset, A., Shalizi, C., Newman, M., 2009. Power-law distributions in empirical data. SIAM Rev. 51 (4), 661–703.

de Haan, L., Ferreira, A., 2006. Extreme Value Theory. An Introduction. Springer.

Eeckhout, J., 2004. Gibrat's law for (all) cities. Am. Econ. Rev. 94 (5), 1429–1451.

Fu, D., Pammolli, F., Buldyrev, S.V., Riccaboni, M., Matia, K., Yamasaki, K., Stanley, H.E., 2005. The growth of business firms: Theoretical framework and empirical evidence. Proc. Natl. Acad. Sci. 102 (52), 18801–18806.

Gabaix, X., 1999. Zipf's law for cities: An explanation. Q. J. Econ. 114, 739–767.

Gabaix, X., 2009. Power laws in economics and finance. Ann. Rev. Econ. 1, 255–293.

Growiec, J., 2008a. A new class of production functions and an argument against purely labor-augmenting technical change. Int. J. Econ. Theory 4 (4), 483–502.

Growiec, J., 2008b. Production functions and distributions of unit factor productivities: Uncovering the link. Econom. Lett. 101 (1), 87–90.

Growiec, J., 2013. A microfoundation for normalized ces production functions with factor-augmenting technical change. J. Econ. Dyn. Control 37 (11), 2336–2350.

Growiec, J., Pammolli, F., Riccaboni, M., Stanley, H.E., 2008. On the size distribution of business firms. Econom. Lett. 98, 207–212.

Hsieh, C.-T., Klenow, P.J., 2009. Misallocation and manufacturing TFP in China and India. Q. J. Econ. 124 (4), 1403–1448.

Jones, C.I., 2005. The shape of aggregate production functions and the direction of technical change. Q. J. Econ. 120 (2), 517–549.

Jones, C.I., 2011. Intermediate goods and weak links in the theory of economic development. Am. Econ. J. Macroecon. 3 (2), 1–28.

Klump, R., McAdam, P., Willman, A., 2012. The normalized CES production function: Theory and empirics. J. Econ. Surv. 26 (5), 769–799.

Kotz, S., Nadarajah, S., 2000. Extreme Value Distributions. Theory and Applications. Imperial College Press, London.

Kremer, M., 1993. The O-ring theory of economic development. Q. J. Econ. 108 (3), 551–575.

Leadbetter, M.R., Lindgren, G., Rootzén, H., 1983. Extremes and Related Properties of Random Sequences and Processes. Springer.

Stanley, M.H.R., Buldyrev, S.V., Havlin, S., Mantegna, R.N., Salinger, M.A., Stanley, H.E., 1995. Zipf plots and the size distribution of firms. Econom. Lett. 49, 453–457.

Sutton, J., 1998. Technology and Market Structure: Theory and History. MIT Press, Cambridge MA.