

# Taking Melitz Seriously: A Simple Approach for Identifying Structural Parameters of Trade Models with Firm Heterogeneity

Saad Ahmad  
Zeynep Akgul

ECONOMICS WORKING PAPER SERIES  
Working Paper 2018-08-D

U.S. INTERNATIONAL TRADE COMMISSION  
500 E Street SW  
Washington, DC 20436  
August 2018

Office of Economics working papers are the result of ongoing professional research of USITC Staff and are solely meant to represent the opinions and professional research of individual authors. These papers are not meant to represent in any way the views of the U.S. International Trade Commission or any of its individual Commissioners. Working papers are circulated to promote the active exchange of ideas between USITC Staff and recognized experts outside the USITC and to promote professional development of Office Staff by encouraging outside professional critique of staff research.

# Taking Melitz Seriously: A Simple Approach for Identifying Structural Parameters of Trade Models with Firm Heterogeneity

Saad Ahmad and Zeynep Akgul  
Office of Economics Working Paper 2018-08-D  
August 2018

## **Abstract**

Quantitative trade models based on the theory of firm heterogeneity generally assume that distributions of firm productivity as well as firm size are characterized by the power-law, especially in the upper tail of the distribution where rare events such as exporting occur. The most frequent distribution used in such models is the Pareto distribution. While power-laws are widely used in the literature, it may not be an accurate fit for the entire distribution due to fluctuations in the tails. Therefore, it is paramount to identify a minimum threshold above which the power-law provides a good fit for the data. This is especially important in estimating the structural parameters of the firm heterogeneity model for use in policy analysis as biased estimates may distort trade volume and welfare responses. In this paper, as in Clauset et al. (2009), we combine maximum-likelihood and the Kolmogorov-Smirnov (KS) statistic to estimate both the minimum threshold for truncating the data as well as the shape parameter, under a power-law, of the firm size and productivity distributions. We then impute the elasticity of substitution across varieties that are appropriate to use in firm heterogeneity models of trade.

Saad Ahmad  
Office of Economics, Research Division  
[Saad.Ahmad@usitc.gov](mailto:Saad.Ahmad@usitc.gov)

Zeynep Akgul  
Center for Global Trade Analysis, Department of Agricultural Economics, Purdue University  
Office of Economics, Research Division  
[zakgul@purdue.edu](mailto:zakgul@purdue.edu)

## Introduction

Structural parameter values play a critical role in determining trade volumes and welfare responses to policy changes in trade models. However, it has always been a challenge to obtain the appropriate parameter values such as trade elasticities from empirical methods such as the gravity model. This is especially true in firm heterogeneity models of trade since there are two key parameters to consider, the shape parameter of productivity distribution across firms and the elasticity of substitution across varieties. Since there are more parameters to consider and the theoretical model is different than traditional gravity models, both the interpretation and the estimation of firm heterogeneity parameters need to be reconsidered.

While there are a number of methodologies prevalent in the current literature for obtaining these parameters, they often lack consistency with the underlying firm heterogeneity theory, indicating a clear need for continued efforts towards theory-consistent parameterization of firm heterogeneity models (Akgul et al., 2015; Ahmad and Akgul, 2017). Indeed, the lack of a general and theoretically-sound approach for obtaining parameters in firm heterogeneity models remain one of the main challenges in advancing their widespread adoption for policy analysis.

A recent approach in Ahmad and Akgul (2017) proposes estimating the structural parameters of firm heterogeneity models by using the theoretical relationship between the distribution of firm size and the distribution of firm productivity. In firm heterogeneity models, the common assumption is that firm productivity follows the Pareto distribution, which is a power-law model with an exponent  $\gamma$ , equal to the shape parameter of productivity distribution. If firm productivity follows a power-law model, then firm size also follows a power-law model; however, with a different exponent  $\alpha = \frac{\gamma}{\sigma-1}$ .<sup>1</sup> Since  $\alpha$  can be estimated directly from firm-level data, it provides a useful way to infer the ratio of structural parameters in the firm heterogeneity model. Based on this methodology, Ahmad and Akgul (2017) use firm-level ORBIS data for motor vehicles and parts sector (MVH) to fit the total factor productivity and firm size distribution to power law models and estimate the structural parameters of the firm heterogeneity model.

We start with this approach and improve the methodology in two key respects. First, there are a limited number of observations for US firms in the MVH sector, and this limits the efficiency of the estimates. We address sample size concerns by using this methodology on all US manufacturing firms in ORBIS as well as considering other countries that have more firm observations in ORBIS, such as Japan and the EU. These changes allow us to vastly increase our sample for subsequent estimations.

---

<sup>1</sup> See Ahmad and Akgul (2017) for the theoretical link between these two distributions.

Second, one of the principal criticisms of assuming a power-law model for a given empirical dataset is the fact that the real world variables do not follow the power-law model over the entire range. In fact, the data typically only follows the power-law model above a lower bound. In particular, the power-law function  $p(x) = Cx^{-\alpha}$  diverges as  $x$  approaches 0, given a positive value of  $\alpha$ . This means that there is a minimum value  $x_{min}$  (lower bound) below which the distribution deviates from the power-law. In order to obtain an unbiased estimate of the shape parameter of productivity distribution, it is important to focus on the upper tail of the sample, rather than the entire range of firm observations. This in turn requires that we are able to obtain an accurate estimate of the lower bound. If the lower bound value is too low, the sample does not follow a power-law and we are trying to fit a non-power-law data to a power-law distribution, and the result is a biased estimate of the shape parameter. If the lower bound value is too high, then we are omitting relevant information from the sample that may increase the statistical error on the parameter estimates as well as the bias from finite size effects (Clauset et al., 2009).<sup>2</sup>

A common method to obtain a lower bound is visualization. There are two ways to estimate visually. The first one is to plot the shape parameter as a function of the lower bound, i.e. truncation point, and identify where the value fluctuates and where it becomes stable. Then choose the point where the relationship becomes stable. The second method is to depict a log-log plot and identify the point where the PDF or the CDF of the distribution becomes relatively straight. The second approach is adopted in Ahmad and Akgul (2017).

In order to avoid the subjectivity of this visualization method, Clauset et al. (2009) offer a more robust and methodical approach in choosing a lower bound. In this paper, we adopt their approach, which is based on minimizing the “distance” between the power-law model and the empirical data. They suggest choosing the value of  $x_{min}$  that makes the model and the distribution of the empirical data as similar as possible. If chosen  $x_{min}$  is higher than the true value, then the sample size is reduced, which may cause statistical fluctuation and makes the probability distribution a poor match. If instead, the chosen  $x_{min}$  is lower than the true value, the data and the model will be fundamentally different causing the distribution to differ. To quantify the distance between the two distributions, they use the Kolmogorov-Smirnov (KS) statistic, which is defined as the maximum distance between the CDFs of the data and the fitted model. Clauset et al. (2009) state that  $n=1000$  observations or more is sufficient to obtain good results with this approach. Lastly, they estimate the power-law exponent on the truncated sample using maximum likelihood.

---

<sup>2</sup> While both cases are not desirable, it may be more better to err with a higher minimum value because the error will be less severe compared to the case with higher sample size and a lower  $x_{min}$  (Clauset et al., 2009).

To summarize our approach in this paper, we combine the methodology in Ahmad and Akgul (2017) and Clauset et al. (2009) to estimate the structural parameters of firm heterogeneity, namely the shape parameter of productivity distribution and the elasticity of substitution across varieties. We analyze the power-law model in three steps following Clauset et al. (2009). Then in the fourth step we impute the value of elasticity of substitution across varieties based on the power-law model.

- I. We estimate the lower bound parameter  $x_{min}$  by minimizing the distance between the empirical distribution and model distribution based on the Kolmogorov-Smirnov (KS) statistic.
- II. Based on the value of  $x_{min}$  we estimate the power-law exponent of firm size and firm productivity using the method of Maximum Likelihood.
- III. We compare the fit of the power-law model with alternative distributions such as the exponential and log-normal distributions using a Likelihood Ratio test.
- IV. We use the estimates of power-law exponents in firm size and firm productivity to impute the elasticity of substitution across varieties in the sector with heterogeneous firms.

We use the ORBIS firm-level database and focus on the manufacturing sector of the US, Japan, and the EU for the years 2012-2016.<sup>3</sup> We then compare the results of this aggregated sector with that of a more disaggregated sector, MVH, in the same regions. For firm size, we use two variables: firm operating revenue and number of employees. In order to calculate the firm productivity levels, we use labor productivity where the firm's operating revenue is divided by the number of employees.

The results show that the power-law provides reasonably good fits for the empirical data and returns exponents above unity, satisfying the theoretical constraint that  $\gamma > \sigma - 1$ . The likelihood ratio tests suggest that the power-law model is a better fit for firm size and firm productivity in manufacturing and the MVH sector than the exponential distribution. However, the likelihood ratio tests against the log-normal distribution are mostly inconclusive, with associated p-values larger than the target value.

The resulting values of elasticity of substitution for the manufacturing sector are found to be in the range of 2.28 – 2.71 when operating revenue is used as a proxy for firm size. There is a slight variation in elasticity values across regions, which may result in variation in trade volume and welfare responses to trade policies. Elasticity values are relatively lower when the number of employees is used for firm size (2.26 – 2.55), reflecting slight variation across regions.

---

<sup>3</sup> The pooled sample consists of all manufacturing firms in ORBIS for the US, EU and Japan.

When we focus on the MVH sector alone, we observe that  $\hat{\sigma}$  values vary slightly across regions and are in the range of 2.88-2.92 when operating revenue is used for firm size, and in the range of 2.72-2.97 when number of employees is used for firm size. Overall, these values are slightly lower than what is found in the literature, suggesting that the estimation strategy is important in obtaining the appropriate values for the theoretical model in question.

## Empirical Methodology

The probability density function (PDF) of a continuous power-law model,  $p(x)$ , can be described as

$$p(x)dx = \Pr(x \leq X < x + dx) = Cx^{-\alpha}dx$$

where  $C$  is a constant,  $X$  is the observed value,  $x$  is the data that are modeled by the distribution and  $\alpha$  is the corresponding power-law exponent, i.e. the shape parameter of the distribution. As discussed in Clauset et al. (2009), this PDF does not hold for all  $x$ . In fact, it may diverge as  $x \rightarrow 0$ . Therefore, the power-law model applies only above a lower bound, which is denoted as  $x_{min}$ .<sup>4</sup> The resulting PDF of a continuous power-law model is given as

$$p(x) = \frac{\alpha - 1}{x_{min}} \left( \frac{x}{x_{min}} \right)^{-\alpha}$$

where  $x_{min}$  is the lower bound for the power-law model, data follows a power-law for  $x \geq x_{min}$ , and  $\alpha$  is the corresponding power-law exponent. The associated complementary cumulative distribution function (CCDF, i.e. 1-CDF) is given as

$$P(x) = \Pr(X \geq x) = \int_x^{\infty} p(x)dx = \left( \frac{x}{x_{min}} \right)^{-\alpha+1}$$

We analyze the power-law model in three steps following Clauset et al. (2009). For the implementation of this methodology, we rely on a power-law fitting library in Python that was developed by Alstott et al. (2014). We now turn to the description of the methodology for each of these three steps.

### Estimating the Lower Bound Parameter:

Following Clauset et al. (2009), a numerical method is used to select the  $x_{min}$  that yields the best power-law model for the data. Specifically, for each  $x_{min}$  over some reasonable range, the Kolmogorov-Smirnov (KS) statistic is utilized to quantify the distance between the empirical distribution and model distribution. While other measures can also be used for quantifying distance, the KS statistic has been shown by Clauset et al. (2009) to perform well in these

---

<sup>4</sup> We note that in this empirical framework, the power-law does not apply to firms below this lower-bound, and so may not be a perfect fit for theoretical models that assume a power-law distribution for all firms.

estimations. The KS statistic is computed as the maximum distance between the CDFs of the data and the power-law model:

$$D = \max |S(x) - P(x)|$$

where  $S(x)$  is the CDF of the data and  $P(x)$  is the CDF for the theoretical power-law model, both for observations with value  $x \geq x_{min}$ . The estimated  $x_{min}$  is the one that provides the best fit to the data by minimizing the distance  $D$ .

**Estimating the Power-Law Exponent:**

The first step of estimating the exponent in a power-law model requires the correct identification and estimation of the lower bound parameter,  $x_{min}$ . Once the value of  $x_{min}$  is estimated based on the methodology described in (i) above, the power-law exponent is estimated using the method of Maximum Likelihood.

The Log likelihood function  $\mathcal{L}$  is given as

$$\mathcal{L} = \ln p(x|\alpha) = \ln \prod_{i=1}^n \frac{\alpha - 1}{x_{min}} \left( \frac{x}{x_{min}} \right)^{-\alpha}$$

When we maximize  $\mathcal{L}$  with respect to  $\alpha$  such that  $\frac{\partial \mathcal{L}}{\partial \alpha} = 0$ , the Maximum Likelihood Estimator is:

$$\hat{\alpha} = 1 + n \left[ \sum_{i=1}^n \ln \frac{x_i}{x_{min}} \right]^{-1}$$

where  $x_i, i = 1, 2, \dots, n$  are the observed values of  $x$  such that  $x \geq x_{min}$ .

**Likelihood Ratio Tests for Alternative Distributions:**

If the power-law is a good fit for the dataset, we should also investigate whether alternative distributions provide a better fit than the power-law. In order to make that evaluation, we follow Clauset et al. (2009) and use the likelihood ratio test, which computes the logarithm of the ratio of the likelihoods of the data between two distributions. We compare the power-law model to the exponential and log-normal distributions. A positive value of the likelihood ratio indicates that the power-law model is a better fit compared to the alternative, while a negative value indicates that the power-law model is a worse fit compared to the alternative.

**Computing the Elasticity of Substitution:**

In the firm heterogeneity literature, when firm productivity follows a power-law model, in this case Pareto distribution, firm size also follows a power-law model with a Pareto distribution, but with a different power-law exponent. di Giovanni et al. (2011) shows that these two

exponents are connected such that we can infer the value of elasticity of substitution across varieties by using these two estimated power-law exponents. More specifically, if firm productivity is from a power-law model with exponent  $\gamma$ , then the firm size also follows a power-law model with exponent  $\zeta = \frac{\gamma}{\sigma-1}$  where  $\sigma$  is the elasticity of substitution. Since  $\zeta$  along with  $\gamma$  can be estimated directly from firm-level data, this theoretical relationship can be used to infer a value for  $\sigma$ .

## Firm-Level Data

In this paper, we use the ORBIS database to obtain annual firm-level financial data on the manufacturing sector in the US, Japan, and the EU. We restrict the time frame of our study to the 2012-2016 period. While we consider an aggregated manufacturing sector, we also recognize that a more disaggregated sectoral analysis may result in different power-law exponents and higher dispersion. Thus, for comparison, we also conduct our power-law estimation on a more disaggregated sector. We select the MVH sector in GTAP for this analysis, and we use the 4-digit NACE codes are used for sectoral identification. Sector codes include: (i) 2910 Manufacture of Motor Vehicles, (ii) 2920 Manufacture of bodies (coachwork) for motor vehicles, manufacture of trailers and semi-trailers, and (iii) 2930 Manufacture of parts and accessories for motor vehicles.

ORBIS uses both administrative and public data to provide firm-level information for over 200 million companies worldwide. Several procedures have been undertaken in ORBIS to verify the quality of reported data, including an indexation strategy to ensure the uniqueness of individual firms as well as an analysis to detect unusual variations in financial values between years.

## Results

We estimate both the lower bound on the power-law behavior and the power-law exponent in firm size as well as for productivity. As a proxy for firm size, we consider firm-level information on operating revenue and the number of employees. As a proxy for productivity, we use the ratio of operating revenue to number of employees, a standard measure of labor productivity. We also use goodness of fit tests to compare the power-law distribution to alternative distributions such as the exponential and log-normal distributions.

In the first section of the results, we focus on an aggregated analysis with the complete manufacturing sector data, where we group the data into four regions: (i) firm-level data for the manufacturing sector in the US, (ii) firm-level data for the manufacturing sector in Japan, (iii) firm-level data for the manufacturing sector in Europe and (iv) a pooled version of all firms in the manufacturing sector of the US, Japan and Europe. In the second section of the results, we focus on a more disaggregated analysis where only firms in the MVH sector are considered. As



in the first section of the results, we again group the data for these firms in four regions and see if this results in different estimates.

### Results for the Manufacturing Sector

We first discuss the power-law fits and estimates of power-law exponents in the manufacturing sector. For firm size, we use two proxies: operating revenue and the number of employees.

Table 1 presents the estimates of lower bound ( $\hat{x}_{min}$ ) of the power-law distribution for operating revenue in the manufacturing sector of the US, Japan, the EU, and the pooled data. Once the value of  $\hat{x}_{min}$  is calculated based on the KS statistic in each region, the data are truncated at that lower bound and the power-law exponents are estimated. Table 1 also presents the number of observations in each region before and after truncation. We observe that the value of lower bound for operating revenue in the US manufacturing sector is dramatically higher than that for the rest of the regions. While  $\hat{x}_{min}$  values for Japan, the EU and pooled data are close (\$63,204, \$79,146, and \$60,393, respectively), it is \$4,200,836 for the US. This difference stems from the fact that there are fewer observations in ORBIS for the US (4852 before truncation). According to our estimates, 19% of all available observations in US manufacturing are above the  $\hat{x}_{min}$  value (926 after truncation). This percentage drops to 11% for Japan, 2% in the EU and around 3% in the pooled region. Overall, a small range of observations in all regions follow a power-law, yet the corresponding firms are much larger in the US compared to other regions.

**Table 1: Lower Bound for Pareto Law in the Manufacturing Sector for Operating Revenue**

	X <sub>min</sub> (dollars)	Observations (before truncation)	Observations (after truncation)
US	4,200,836	4852	926
JPN	63,204	103664	11463
EU	79,146	1977461	44394
All	60,393	2085977	72253

The fit of the power-law to our data sets for operating revenue in the manufacturing sector is shown in Figure 1. Complementary cumulative distribution function (1-CDF;  $p(X_i \geq x)$ ) in each region are reported based on the  $\hat{x}_{min}$  values on Table 1. Firm operating revenue seems to follow a power-law model up until the tail of the distribution in each region. At the tail, the power-law fit diverges slightly from the empirical fit. This may result from the fact that we use all the firms in the database and cannot distinguish between exporters and non-exporters. As discussed in di Giovanni et al. (2011), there is a systematic effect of international trade on firm size distribution such that power-law exponents differ between exporters and non-exporters in their French firm-level data. Since the right-tail of the firm size distribution is often associated with exporting firms, the divergence observed in Figure 1 can be explained by the fact that the

right-tail may have a different exponent due to differences between exporters and non-exporters. Since the ORBIS database does not provide the exporting status of the firm this information, we pool together all of the firm-level information in the database without taking into account exporting activity.

**Figure 1: Complementary CDF of firm operating revenues and the fitted power-law distributions for Manufacturing Sector**

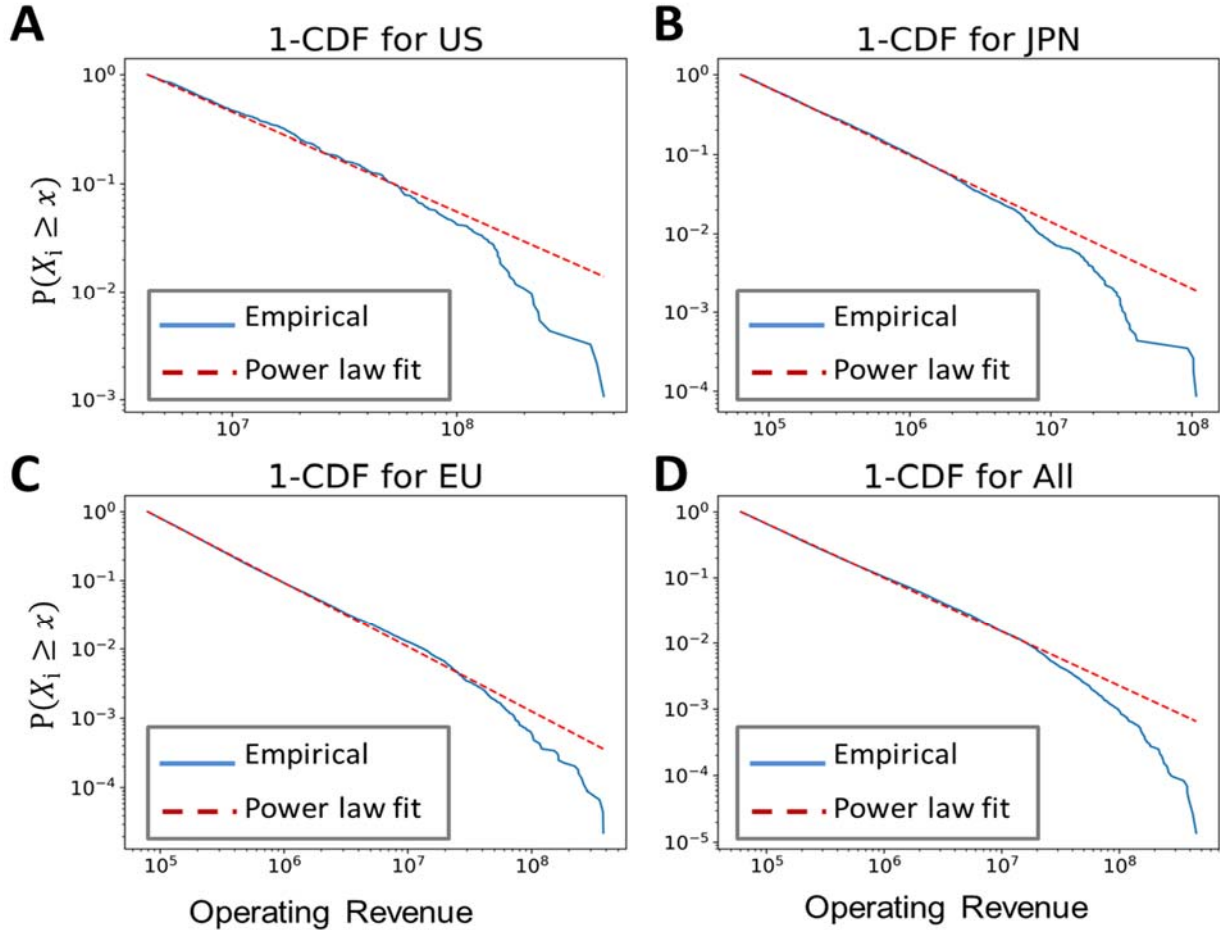


Table 2 reports estimates of the power-law exponents as well as the likelihood ratio tests in each region. We follow Clauset et al. (2009) in choosing exponential and log-normal as the alternative distributions. The power-law exponent is estimated based on the  $\hat{x}_{min}$  values reported in Table 1. The resulting  $\hat{\alpha}$  values are found to be above one for each region, which satisfies the mathematical constraint for firm heterogeneity parameters,  $\gamma > \sigma - 1$ . Specifically,  $\hat{\alpha}$  is 1.913 for the US, 1.841 for Japan, 1.936 for the EU, and 1.822 for all the regions. The results show that the power-law exponents are similar across the regions for this particular sector.

**Table 2: Power-Law Estimates and LR Tests for Manufacturing Sector (Operating Revenue)**

	Alpha	Exponential LR	Exponential p	Log-normal LR	Log-normal p
US	1.913	7.458	0.000	-3.204	0.001
JPN	1.841	21.328	0.000	-4.077	0.000
EU	1.936	32.469	0.000	0.417	0.677
All	1.822	48.928	0.000	-1.442	0.149

*Notes: The first column gives estimates of the power-law exponent. For each non-nested alternative distribution, we report the normalized log-likelihood ratio (LR, where LR is normalized by its standard deviation), and the corresponding p-values (p) for the significance of LR. LR will be positive if the data are more likely in the first distribution, while it will be negative if the data are more likely in the alternative distribution. In this paper, the first distribution is the power-law.*

We also test the fits against alternative distributions and report the log-likelihood ratio as well as the corresponding p-value in Table 2. Positive values of the LR mean that the power-law model provides a better fit compared to the alternative distribution, while negative values mean that the power-law model provides a worse fit compared to the alternative distribution. The reported p-values indicate the significance of the test. Small p-values indicate that the alternate model has a worse fit and should be rejected in favor of the power-law model. In this paper we choose a p-value of 0.1 following Brzezinski (2014) such that if the reported p-value is larger than 0.1, it is not possible to choose between the two models.

Positive LR values in Column 2 of Table 2 indicate that the power-law model is a better fit compared to the exponential distribution for all regions. The associated p-values are low such that the exponential distribution can be ruled out as a plausible model for the operating revenue data in manufacturing. However, the LR values for log-normal distribution are negative in the US, Japan and All regions. This suggests that power-law is not a good fit against log-normal in this dataset. While for EU, the LR value is positive, the test is inconclusive since the corresponding p-value is large. Therefore, for the EU and also for the pooled sample, power-law and log-normal are not distinguishable.

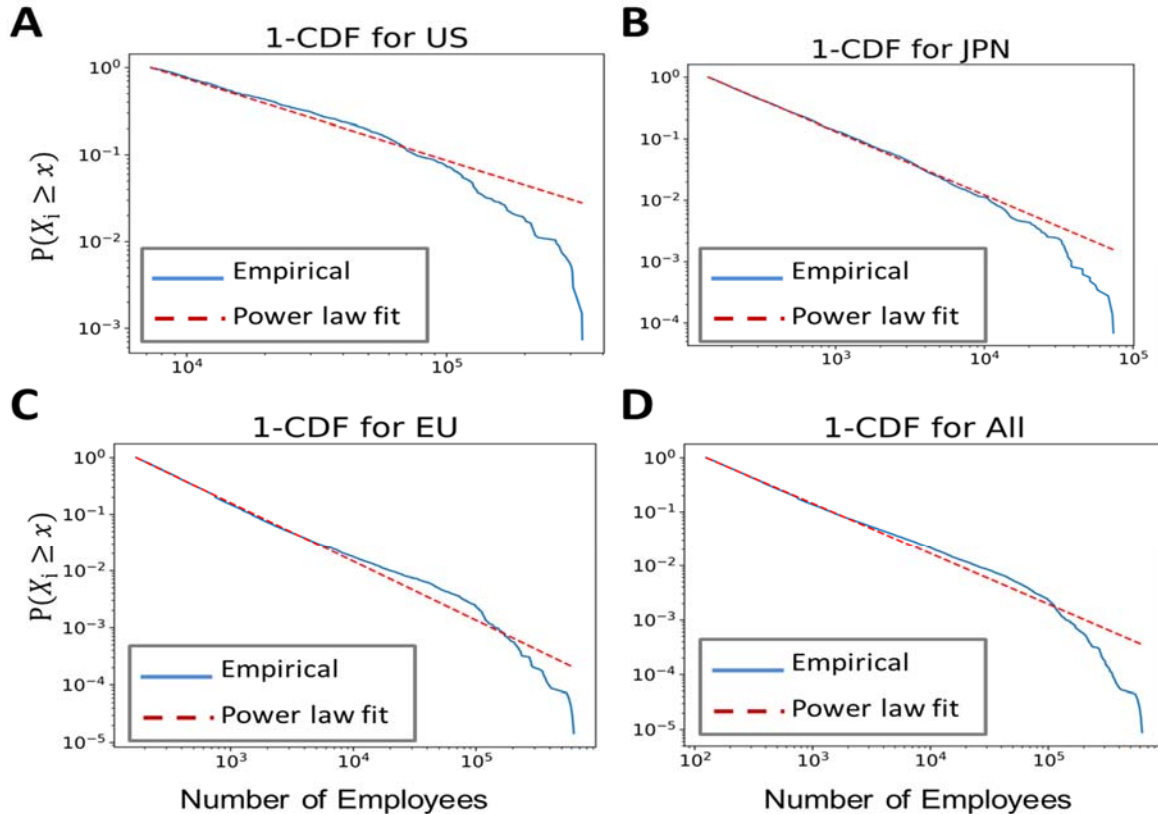
A similar analysis for firm size is conducted with the number of employees in manufacturing for the same four regions. Table 3 reports the estimates of lower bound ( $\hat{x}_{min}$ ) of power-law and the number of observations for number of employees data. The average number of observations above  $\hat{x}_{min}$  for the number of employees in the manufacturing sector is 1345 (27%), 14459 (14%), 68442 (4%), and 110844 (5%) for the US, Japan, the EU, and All, respectively. Similar to the operating revenue data, we observe that a relatively larger fraction of observations for number of employees are above the lower bound in the US compared to the other regions.

**Table 3: Lower Bound for Pareto Law in the Manufacturing Sector for Number of Employees**

	$X_{min}$ (workers)	Observations (before truncation)	Observations (after truncation)
US	7,274	4999	1345
JPN	139	101916	14459
EU	170	1947108	68442
All	125	2054023	110844

Figure 2 shows the power-law fit for number of employees with the CCDF in each region based on the  $\hat{x}_{min}$  values in Table 3. Similar to the operating revenue data, the power-law fit diverges slightly from the empirical fit at the right tail.

**Figure 2: Complementary CDF of firm employees and the fitted power-law distributions for Manufacturing Sector**



The power-law exponent, based on the  $\hat{x}_{min}$  values, is reported in Table 4. The resulting  $\hat{\alpha}$  values are found to be around 2 for each region. These results satisfy the mathematical constraint for firm heterogeneity parameters,  $\gamma > \sigma - 1$ . Specifically,  $\hat{\alpha}$  is 1.936 for the US, 2.030 for Japan, 2.036 for the EU, and 1.933 for all regions. Similar to the operating revenue data, we observe little variation in the power-law exponents across regions.

**Table 4: Power-Law Estimates and LR Tests for Manufacturing Sector (Number of Employees)**

	Alpha	Exponential LR	Exponential p	Log-normal LR	Log-normal p
US	1.936	7.197	0.000	-5.137	0.000
JPN	2.030	24.892	0.000	-1.405	0.160
EU	2.036	40.454	0.000	3.490	0.000
All	1.933	64.026	0.000	-0.962	0.336

Notes: The first column gives estimates of power-law exponent. For each non-nested alternative distribution, we report the normalized log-likelihood ratio (LR, where LR is normalized by its standard deviation), and the corresponding p-values (p) for the significance of LR. LR will be positive if the data are more likely in the first distribution, while it will be negative if the data are more likely in the alternative distribution. In this paper, the first distribution is the power-law.

Comparison against alternative distributions for the number of employees draws a similar conclusion to the operating revenue case. Positive LR values are observed when the power-law is compared against the exponential distribution, which suggests that the power-law model is a better fit for every region. On the other hand, comparison against log-normal distribution does not produce a systematic conclusion. LR value is negative for the US with a low p-value, which suggests that power-law is a worse fit against log-normal. For Japan and All, LR value is still negative; however, the p-values are large such that the test is inconclusive. In comparison, for the EU, LR value is positive with a low p-value suggesting that power-law model is a better fit for these data against log-normal.

In order to impute the elasticity of substitution, we also require the power-law exponent for firm productivity. In this paper, we use a standardized measure of firm productivity by calculating labor productivity, as operating revenue divided by the number of employees.

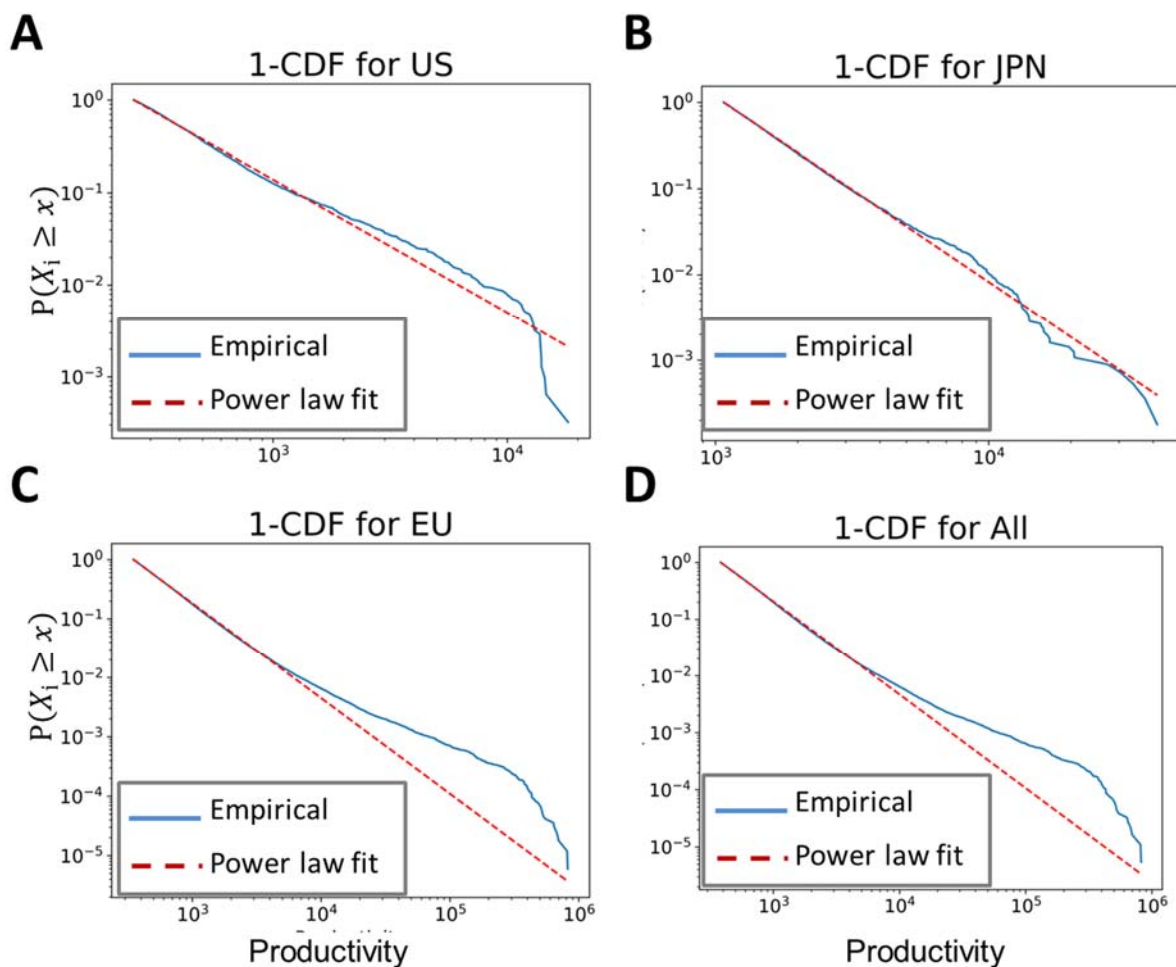
Table 5 reports the estimates of lower bound ( $\hat{x}_{min}$ ) of power-law and the number of observations of firm productivity. We observe that the average number of observations above  $\hat{x}_{min}$  of firm productivity in the manufacturing sector is 3103 (65%), 5598 (6%), 167709 (9%), and 183138 (9%) for the US, Japan, the EU, and All, respectively. A substantially larger fraction of observations for firm productivity is above the lower bound in the US compared to the other regions. This stems from the fact that  $\hat{x}_{min}$  for the US is lower than the rest of the regions. Firm productivity in Japanese manufacturing sector seems to have the largest  $\hat{x}_{min}$  value.

**Table 5: Lower Bound for Pareto Law in the Manufacturing Sector for Productivity**

	Xmin (\$ per worker)	Observations (before truncation)	Observations (after truncation)
US	254	4781	3103
JPN	1,071	101890	5598
EU	349	1924283	167709
All	381	2030954	183138

The fit of the power-law to firm productivity in the manufacturing sector is shown in Figure 3. While the power-law model is a relatively good fit in the US and Japan based on Figure 3, it is not as good for the EU and the pooled region.

**Figure 3: Complementary CDF of firm productivity and the fitted power-law distributions for the Manufacturing Sector**



The resulting  $\hat{\gamma}$  values, shown in Table 6, are in the range of 2.442 – 3.146. These are the values that we use for the shape parameter of Pareto distribution in the firm heterogeneity model.

When we compare the power-law model against alternative distributions in Table 6, LR values are found to be positive for both the exponential and log-normal distributions suggesting that power-law model is a better fit than the alternatives. P-values are low for all regions in the exponential case and also for the EU and All in the log-normal case. More generally, we can rule out exponential distribution as a plausible model for productivity in all regions. We can also rule out log-normal for the EU and All regions.

**Table 6: Power-Law Estimates and Tests for Manufacturing Sector (Productivity)**

	Gamma	Exponential LR	Exponential p	Log-normal LR	Log-normal p
US	2.442	13.344	0.000	1.237	0.216
JPN	3.146	11.076	0.000	0.927	0.354
EU	2.611	24.540	0.000	6.224	0.000
All	2.642	24.241	0.000	6.276	0.000

Notes: The first column gives estimates of power-law exponent. For each non-nested alternative distribution, we report the normalized log-likelihood ratio (LR, where LR is normalized by its standard deviation), and the corresponding p-values (p) for the significance of LR. LR will be positive if the data are more likely in the first distribution, while it will be negative if the data are more likely in the alternative distribution. In this paper, the first distribution is the power-law.

Table 7 reports the imputed values of the elasticity of substitution across manufacturing varieties. We use the power-law exponent in firm size ( $\hat{\alpha} = \frac{\hat{\gamma}}{\hat{\sigma}-1}$ ) and in productivity ( $\hat{\gamma}$ ) to impute the elasticity of substitution.

**Table 7: Imputed Values of Elasticity of Substitution for the Manufacturing Sector**

	Sigma (Operating Revenue)	Sigma (Employees)
US	2.28	2.26
JPN	2.71	2.55
EU	2.35	2.28
All	2.45	2.37

When operating revenue is used for firm size, the resulting  $\hat{\sigma}$  values are found to be in the range of 2.28 – 2.71. There is a slight variation in the elasticity values across regions, which may result in variation in trade volume and welfare responses to trade policies. The elasticity values are relatively low when the number of employees is used as a proxy for firm size. The range is 2.26 – 2.55, reflecting a slight variation across regions.

#### Results for the Motor Vehicles and Parts (MVH) Sector

It is important to note that manufacturing is a highly aggregated sector. In order to test whether the elasticity values change when the sector is disaggregated, we conduct a similar analysis by focusing on the MVH sector of the US, Japan, the EU, and All regions (pooled data).

Table 8 presents the estimates of lower bound ( $\hat{x}_{min}$ ) of power-law for operating revenue in the MVH sector. We observe that the average number of observations above  $\hat{x}_{min}$  for operating revenue in the MVH sector is 5 (28%), 341 (91%), 2423 (25%), and 2522 (25%) for the US, Japan, the EU, and All, respectively. For the US, the estimated  $\hat{x}_{min}$  value is extremely large compared to the rest of the regions due to the low number of observations in the dataset. The

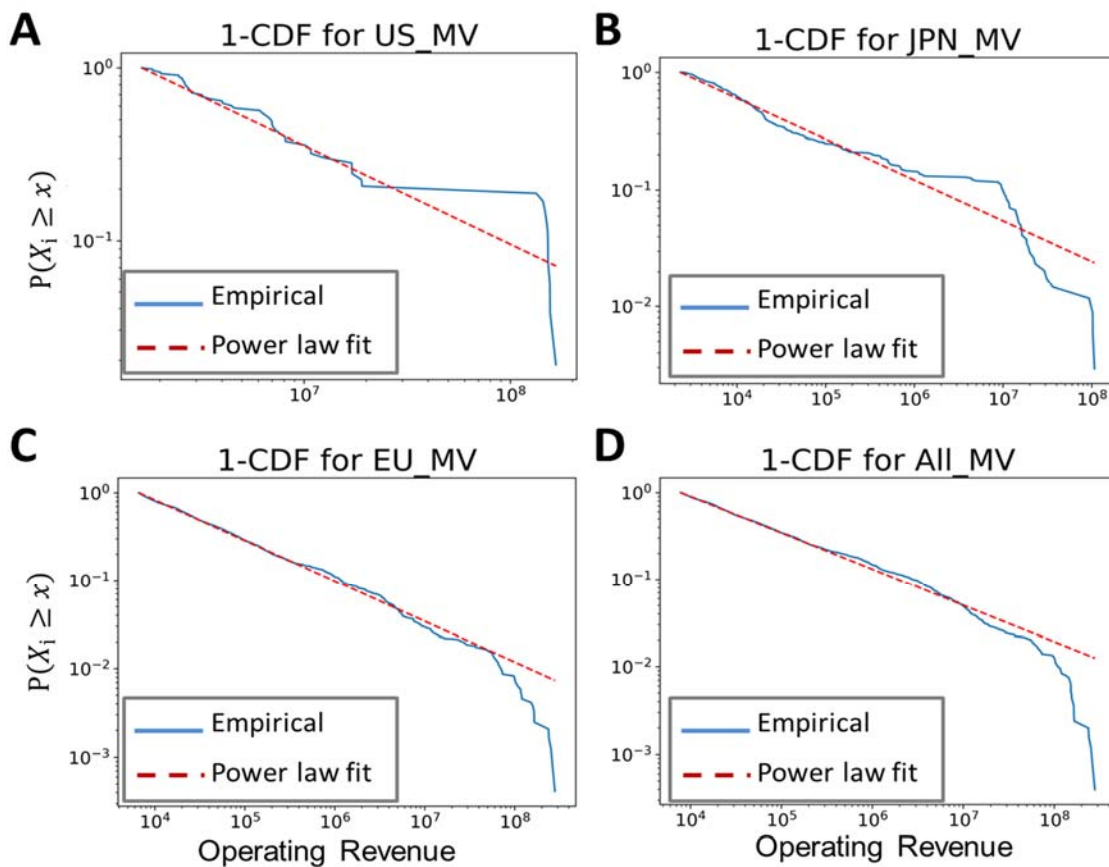
firms in the tail are significantly larger. The case of Japan is also striking in the sense that majority of the observations is above  $\hat{x}_{min}$ .

**Table 8: Lower Bound for Pareto Law in the MVH Sector for Operating Revenue**

	Xmin (dollars)	Observations (before truncation)	Observations (after truncation)
US	1,635,686	18	5
JPN	2,315	376	341
EU	6,651	9724	2423
All	7,817	10190	2522

Power-law fits to firm productivity in the MVH sector are shown in Figure 4. The power-law fits are slightly better in the EU and All regions compared to the US and Japan. Nonetheless, they all diverge in the right tail.

**Figure 4: Complementary CDF of firm operating revenues and the fitted power-law distributions for the MVH Sector**





Similar to the manufacturing sector, Table 9 shows that the exponential distribution can be ruled out as a plausible fit for operating revenue in the MVH sector. LR values are positive for all regions with low p-values. On the other hand, log-normal distribution cannot be ruled out. While the LR values are negative, indicating that power-law model is a worse fit compared to log-normal for the operating revenue data, the p-values are large which makes the test inconclusive. Therefore, power-law and log-normal are not distinguishable in these data. The resulting estimates of power-law exponents are above 1 for all regions, within the range of 1.348 – 1.571. This range satisfies the mathematical constraint.

**Table 9: Power-Law Estimates and LR tests for the MVH Sector (Operating Revenue)**

	Alpha	Exponential LR	Exponential p	Log-normal LR	Log-normal p
US	1.571	5.103	0.000	-1.042	0.298
JPN	1.348	17.114	0.000	-1.176	0.240
EU	1.461	24.326	0.000	-1.517	0.129
All	1.418	29.761	0.000	-2.619	0.009

*Notes: The first column gives estimates of the power-law exponent. For each non-nested alternative distribution, we report the normalized log-likelihood ratio (LR, where LR is normalized by its standard deviation), and the corresponding p-values (p) for the significance of LR. LR will be positive if the data is more likely in the first distribution, while it will be negative if the data is more likely in the alternative distribution. In this paper the first distribution is chosen to be the power-law fit.*

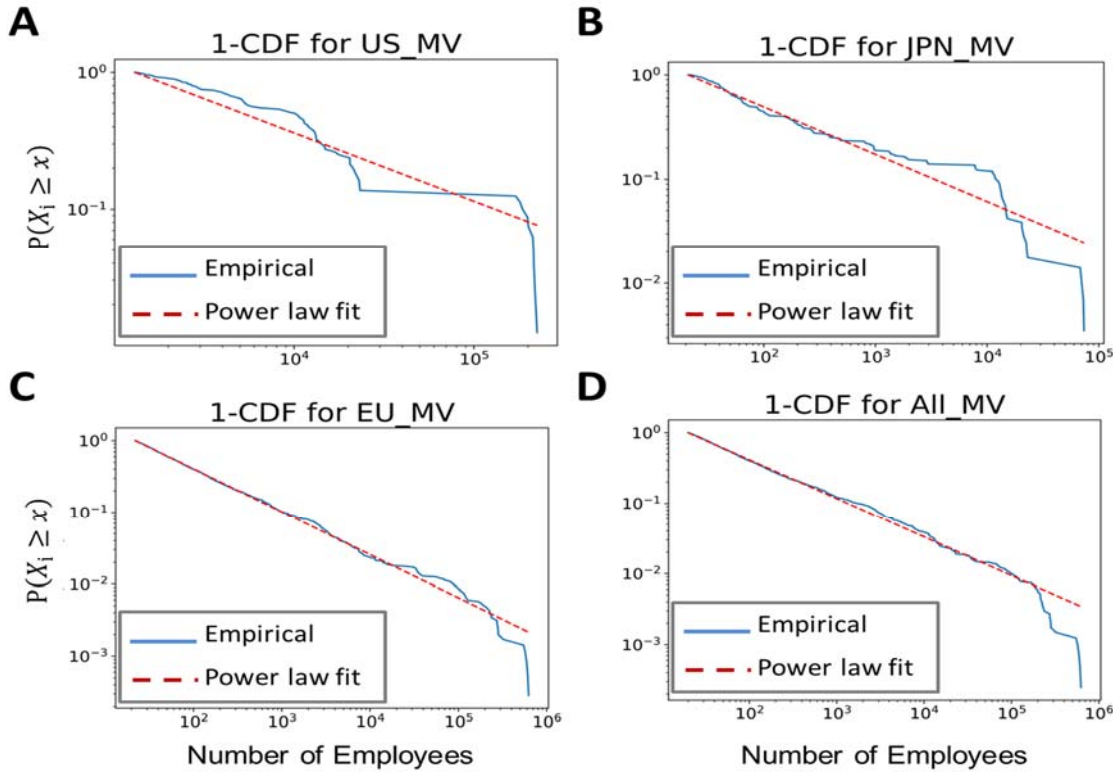
Table 10 presents the estimates of lower bound ( $\hat{x}_{min}$ ) of power-law for the number of employees in the motor vehicles and parts sector. The size of the available data for the US is quite limited with only 9 observations before truncation, 6 of which is above the lower bound 1300 for number of employees.

**Table 10: Lower Bound for Pareto Law in the MVH Sector (Number of Employees)**

	Xmin (workers)	Observations (before truncation)	Observations (after truncation)
US	1,300	9	6
JPN	21	367	284
EU	22	9586	3423
All	20	10043	3952

Power-law fits for number of employees are shown in Figure 5. They are similar to the operating revenue data, where power-law model seems to provide a better fit in the EU and All regions, compared to the US and Japan. However, the number of available observations in the MVH sector of the US is only 18 and so may not be enough to do a meaningful analysis of power-law fit.

**Figure 1: Complementary CDF of firm employees and the fitted power-law distribution for the MVH sector**



Corresponding likelihood ratio tests are presented in Table 11. Again, we can rule out exponential distribution as an alternative since the LR values are positive and p-values are small. On the other hand, the comparison of the power-law fit against log-normal distribution is inconclusive. While the LR values are negative, indicating power-law is a worse fit than log-normal, the corresponding p-values are large.

**Table 11: Power-Law Estimates and LR Tests for the MVH Sector (Number of Employees)**

	Alpha	Exponential LR	Exponential p	Log-normal LR	Log-normal p
US	1.499	4.057	0.000	-2.346	0.019
JPN	1.453	13.231	0.000	-0.744	0.457
EU	1.599	19.807	0.000	0.326	0.744
All	1.548	24.589	0.000	-0.081	0.936

Notes: The first column gives estimates of power-law exponent. For each non-nested alternative distribution, we report the normalized log-likelihood ratio (LR, where LR is normalized by its standard deviation), and the corresponding p-values (p) for the significance of LR. LR will be positive if the data are more likely in the first distribution, while it will be negative if the data are more likely in the alternative distribution. In this paper, the first distribution is the power-law.

The resulting power-law exponents are in the range of 1.453-1.599 when the number of employees is used as a measure of firm size. These values are slightly lower than the power-law exponents in the aggregated manufacturing sector, which are around the range of 2, suggesting that firms in the MVH sector are more dispersed in size than in the larger manufacturing sector.

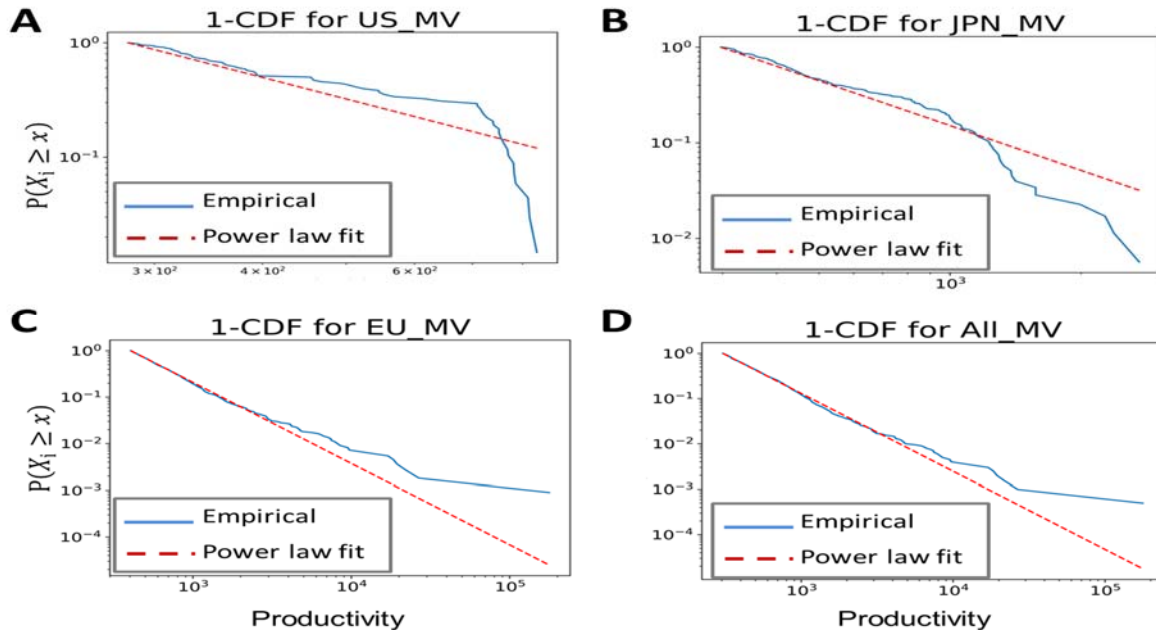
Table 12 presents the estimates of lower bound ( $\hat{x}_{min}$ ) of power-law for firm productivity in the MVH sector. We observe that the average number of observations above  $\hat{x}_{min}$  for productivity in the MVH sector is 67 (74%), 175 (48%), 1105 (12%), and 2015 (20%) for the US, Japan, the EU, and All, respectively. Unlike in firm size,  $\hat{x}_{min}$  values are close amongst the different regions.

**Table 12: Lower Bound for Pareto Law in the MVH Sector for Productivity**

	Xmin (\$ per worker)	Observations (before truncation)	Observations (after truncation)
US	280	90	67
JPN	297	367	175
EU	405	9461	1105
All	305	9918	2015

Figure 6 plots the power-law fits for firm productivity in the MVH sector. These plots suggest that the power-law is not as good a fit for productivity as in the case of firm size at least for this dataset. There is a systematic divergence of the variable in question in the right tail for all regions and may indicate need of more than one power-law fit with two different exponents.

**Figure 2: Complementary CDF of firm productivity and the fitted power-law distribution for the MVH sector**



Estimates of power-law exponents and likelihood ratio tests are presented in Table 13. LR values are negative for both the exponential and log-normal distributions in the US and Japan indicating power-law is not a good fit against either of the alternative distributions. Corresponding p-values are not small enough; therefore, the test is not conclusive. For the cases of the EU and All regions, LR values are positive for both the exponential and log-normal distributions suggesting that power-law mode may provide a better fit than the two alternatives. Corresponding p-values for log-normal are again large making the test inconclusive. The resulting power-law exponent for firm productivity is found to be in the range of 2.557-2.947 for the four regions in our data.

**Table 13: Power-Law Estimates and LR Tests for the MVH Sector (Productivity)**

	Gamma	Exponential LR	Exponential p	Log-normal LR	Log-normal p
US	2.947	-2.762	0.006	-2.027	0.043
JPN	2.557	-0.619	0.536	-2.269	0.023
EU	2.743	2.666	0.008	1.717	0.086
All	2.720	2.708	0.007	0.478	0.633

*Notes: The first column gives estimates of power-law exponent. For each non-nested alternative distribution, we report the normalized log-likelihood ratio (LR, where LR is normalized by its standard deviation), and the corresponding p-values (p) for the significance of LR. LR will be positive if the data are more likely in the first distribution, while it will be negative if the data are more likely in the alternative distribution. In this paper, the first distribution is the power-law.*

Based on the values  $\hat{\gamma}$  in Table 13 and the  $\hat{\alpha}$  values in Tables 9 and 11, we calculate the elasticity of substitution values for MVH. The results are presented in Table 14. We observe that  $\hat{\sigma}$  values slightly vary across regions and are in the range of 2.88-2.92 when operating revenue is used as firm size and in the range of 2.72-2.97 when the number of employees is used for firm size.

**Table 14: Imputed Values of Elasticity of Substitution for MVH Sector**

	Sigma (Operating Revenue)	Sigma (Employees)
US	2.88	2.97
JPN	2.90	2.76
EU	2.88	2.72
All	2.92	2.76

These values are higher than the elasticity values for the manufacturing sector. This suggests that the MVH sector products are more homogeneous compared to the manufacturing sector. Since the manufacturing sector is highly aggregated, it contains a variety of products that are much more differentiated compared to the MVH sector.

## Conclusion

In this paper, we estimate the structural parameters of the model of trade with firm heterogeneity, namely the shape parameter of the productivity distribution and the elasticity of substitution across varieties. We use ORBIS firm-level data and focus on an aggregated manufacturing sector and a disaggregated MVH sector. We combine the methodology in Ahmad and Akgul (2017) and Clauset et al. (2009). We first estimate the lower-bound for the power-law minimizing the KS distance statistic between CDFs of the data and the model. We then estimate the power-law exponent of firm size and firm productivity based on Maximum Likelihood. The resulting power-law exponents are used to impute the elasticity of substitution. In addition, we compare the power-law fit to alternative distributions such as the exponential and log-normal distributions using likelihood ratio tests.

It is important to note that identification of optimal  $x_{min}$  for isolation of the upper tail of distribution also depends on the distance metric used to represent the differences between the empirical data and the fitted power-law. There are several alternative metrics that are often considered for this purpose. For this study, we restricted our attention to KS distance; however, it is worth investigating other metrics such as Kuiper or Anderson-Darling (Alstott et al. 2014).

It is also quite interesting to see that the empirical distributions considered in this study have a concave downward trend in log-log plots, therefore showing systematic deviations from power-law behavior particularly at the upper tail of the fit interval. This suggests that a single power-law fit may not be adequate for fully capturing the distribution behavior. An improved approach for fitting can involve two power-law distributions with different exponents but a common intersection point. Such double power-law distributions have been considered for explaining various phenomenological distributions in computer science and economics (Mitzenmacher, 2003).

## References

- Ahmad, S. & Akgul, Z., 2017. Using Power Laws to Identify the Structural Parameters of Trade Models with Firm Heterogeneity. USITC Working Paper, 2017-5-B. Available at: [https://www.usitc.gov/research\\_and\\_analysis/staff\\_products.htm](https://www.usitc.gov/research_and_analysis/staff_products.htm)
- Alstott, J., Bullmore, E. & Plenz, D., 2014. powerlaw: A Python Package for Analysis of Heavy-Tailed Distributions. PLOS ONE, 9(1), p.e85777. Available at: <http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0085777>.
- Brzezinski, M., 2014. Do wealth distributions follow power laws? Evidence from 'rich lists. Physica A: Statistical Mechanics and its Applications, 406, pp.155–162. Available at: <https://www.sciencedirect.com/science/article/pii/S0378437114002544>.
- Clauset, A., Shalizi, C. & Newman, M., 2009. Power-Law Distributions in Empirical Data. SIAM Review, 51(4), pp.661–703. Available at: <https://epubs.siam.org/doi/abs/10.1137/070710111>.
- Mitzenmacher, M., 2003. A Brief History of Generative Models for Power Law and Lognormal Distributions. Internet Mathematics, 1(2), pp.226–251. Available at: <https://projecteuclid.org/euclid.im/1089229510>.
- di Giovanni, J. & Levchenko, A., 2013. Firm entry, trade, and welfare in zipf's world. Journal of International Economics, 89 (2), pp. 283-296. Available at: <https://www.sciencedirect.com/science/article/pii/S0022199612001365>
- di Giovanni, J., Levchenko, A. & Ranciere, R., 2011. Power laws in firm size and openness to trade: Measurement and implications. Journal of International Economics 85 (1), pp. 42-52. Available at: <https://www.sciencedirect.com/science/article/pii/S0022199611000535>