European Economic Review 172 (2025) 104912

Contents lists available at ScienceDirect

European Economic Review

journal homepage: www.elsevier.com/locate/eer

Firms and economic performance: A view from trade[☆]

Alessandra Bonfiglioli [a,b,c,*], Rosario Crinò [a,c,d], Gino Gancia [e,c]

[a] University of Bergamo - Department of Economics, Via dei Caniana 2, 24127, Bergamo, Italy
[b] Queen Mary University of London, United Kingdom
[c] CEPR, Paris, France
[d] CESifo, Germany
[e] University of Milano-Bicocca, Italy

ARTICLE INFO

Dataset link: EEREV-D-23-01191_Replication files.zip(Original data)

JEL classification:
F12
F14

Keywords:
US imports
Firm heterogeneity
Superstar firms
International trade

ABSTRACT

We use transaction-level US import data to compare firms from virtually all countries in the world competing in a single destination market. First, we decompose countries' sales into the contribution of the number of firm-products, their average appeal and its dispersion. Then, by making distributional assumptions consistent with the data, we identify new structural parameters that are useful in understanding the role of firm heterogeneity for trade and economic performance. We find that differences in the dispersion of appeal are quantitatively important in explaining exports, even after controlling for selection, average appeal and other determinants of trade, and that they are relevant for welfare. We also find that countries with a higher GDP per capita export more per firm largely because they have a higher dispersion of appeal, hence more heterogeneous firms.

1. Introduction

Understanding differences in economic performance across countries has always been a great challenge. Until recently, the main focus was on measuring aggregate productivity from national accounts. The availability of firm-level data revolutionized the field by showing that productivity varies enormously even across firms within countries. One striking fact emerging from this literature is the increasingly dominant role of top firms both domestically (e.g., Autor et al., 2020) and in international markets (e.g., Freund and Pierola, 2015).[1] However, some important questions on the role of firm heterogeneity still remain unanswered: How relevant quantitatively are firms that deviate from the country-industry average to explain aggregate performance? How does

☆ We thank Evi Pappa (the Editor), two anonymous Referees, Sergi Basco, Andy Bernard, Paula Bustos, Ed Glaeser, Kiminori Matsuyama, Peter Neary, Laura Ogliari, Josep Pijoan-Mas, Veronica Rappoport, Steve Redding, Mathias Thoenig and seminar participants at the 2020 AEA Annual Meeting (San Diego), the 2018 Barcelona GSE Summer Forum (Barcelona), the CAGE International Research Day (Warwick), the DICE Workshop (Tuscany), the Italian Trade Study Group (Milan), the 1st QMUL-SUFE Workshop (London), the 2019 WIEN (Vienna), ESSIM (Oslo), CEMFI, ECB, ESADE, ESSEC Business School (Cergy-Pontoise), IMT (Lucca), LSE, OECD (Paris), University of Ancona, Universitat Autonoma de Barcelona, University of Barcelona, University of Edinburgh, University of Kent, University of Nottingham, University of Oxford, University of Siena and Università della Svizzera Italiana for valuable comments.

* Corresponding author at: University of Bergamo - Department of Economics, Via dei Caniana 2, 24127, Bergamo, Italy.
E-mail addresses: alessandra.bonfiglioli@unibg.it (A. Bonfiglioli), rosario.crino@unibg.it (R. Crinò), ginoalessandro.gancia@unimib.it (G. Gancia).

[1] For instance, according to The Economist (17 September 2016), 10% of the world's public companies generate 80% of all profits. In a sample of 32 mostly developing countries, the top five firms account on average for 30% of a county's total exports (Freund and Pierola, 2015).

https://doi.org/10.1016/j.euroecorev.2024.104912
Received 12 December 2023; Received in revised form 14 November 2024; Accepted 17 November 2024
Available online 25 November 2024

firm heterogeneity vary across countries, and with the level of development? What are the implications of this variation for trade flows and welfare?[2]

To answer these questions, in this paper we use transaction-level data on US seaborne imports, containing information on unit values, volumes and the identity of exporting firms for 6-digit products from over 100 countries in 2002 and 2012, to compare firms from virtually all countries in the world. This unique dataset allows us to provide a comprehensive account of how the distribution of firm-level characteristics explains aggregate sales in a single destination market. Moreover, by making distributional assumptions consistent with the data, we identify new structural parameters illustrating the importance of cross-country variation in firm heterogeneity, even within detailed sectors, for trade and economic performance.

As a preliminary step, we follow recent methodological advances in trade theory to quantify the role of various firm-level margins.[3] First, we decompose countries' sales to the US within a given 4-digit industry and year into an extensive margin—the number of firm-products per country—and an intensive margin—the average sales per firm-product from a given country. This exercise shows that each margin accounts on average for half of the overall variation in sales. Second, we decompose the intensive margin, into two parts: the average "appeal" of the firm-products, which we can extract from the data, and deviations from its average. More precisely, we define appeal as quality-to-price ratios, where quality is a demand shifter. Intuitively, countries with more appealing firm-products sell more. However, heterogeneity in appeal also affects total sales because consumers can substitute high-appeal firm-products for low-appeal firm-products. We show that when the elasticity of substitution is higher than two, we are in a "superstar economy" in the sense that more dispersion in appeal, conditional on its average, implies larger sales per firm-product. We find that heterogeneity explains roughly half of the cross-country variation in average sales per firm-product within an industry. We also find that countries with a higher GDP per capita sell more per firm largely because they have more heterogeneous, and hence top, firms.

These results suggest that countries differ in economic performance not just because they host firms that are on average better or worse. Equally important is the ability to breed firms that are exceptional, in that they deviate from the country-industry mean. The main focus of the paper is to study the latter source of economic performance, i.e., differences in firm heterogeneity across countries. In doing so, however, we face three challenges. First, since our sample only includes firms that export to the US, one concern is that differences in dispersion might be driven by selection, especially if progressively less appealing firms export from origins selling more to the US. Second, we would like to know if heterogeneity in a country-industry-year triplet is well approximated by continuous distributions, despite the granular nature of the data. If so, this may suggest that exceptional firms are a manifestation of economic performance rather than outliers. Finally, our preliminary decompositions rely on values of the elasticity of substitution, which are however difficult to estimate.

To make progress, we impose more structural restrictions. Several papers argue that sales are well approximated by a log-normal distribution (see, for instance, Cabral and Mata, 2003; Head et al., 2014; Bas et al., 2017). We confirm this in our sample. In particular, we find that log-normal distributions, with parameters that vary across country-industry-year triplets, provide a good fit of the data.[4] Under the assumption of log-normality, we can use Quantile-Quantile (QQ) regressions in each triplet to obtain estimates of the shape parameters of the distributions of sales that are independent of truncation. This allows us to separately identify differences in heterogeneity and in selection. Moreover, under our assumptions, there is a simple mapping between the distribution of sales and the distribution of appeal. This mapping allows us to flexibly separate cross-country differences in the heterogeneity of appeal from cross-industry differences in the elasticity of substitution, without relying on any estimates of the latter parameter. Using this approach, we confirm that differences in heterogeneity across countries are large; that they are strongly correlated with the volume of exports, even after controlling for selection, average appeal and other determinants of trade; and that heterogeneity correlates positively with per-capita income across countries.

Finally, we argue that firm heterogeneity is important not just to explain observed trade flows, but also for welfare. In particular, we implement an accounting exercise showing that the price index of the basket of imported varieties is significantly lower from origins with higher dispersion. Other things equal, we find that the price index of imports from a country-industry with a variance of log sales one standard deviation higher than the sample mean, roughly corresponding to the observed average difference between Germany and China, would be around 14% lower. Moreover, it is known that both welfare and the gains from trade depend on the extent of heterogeneity (e.g., Bas et al., 2017). Yet, structural models usually allow for cross-country differences in average productivity, but not in its dispersion (e.g., Costinot and Rodríguez-Clare, 2014). Our results imply that considering differences in firm heterogeneity is also important.

This paper is related to the literature on the role of firms for explaining trade flows. Some papers have studied the contribution of the extensive and intensive margin (e.g., Fernandes et al., 2023; Bernard et al., 2018; Fernandes et al., 2016; Bernard et al., 2009; Chaney, 2008; Hummels and Klenow, 2005). This strand of literature has shown that larger and richer countries have both more and bigger exporters. We further decompose the intensive margin into the contribution of average appeal and its dispersion.[5] Freund

[2] International comparisons of firm-level characteristics are limited to a few countries (Gennaioli et al., 2013; Bartelsman et al., 2013; Bloom et al., 2016; Poschke, 2018) and use national data that are not suited for the structural decomposition implemented in this paper. Quantitative models of international trade do not let the degree of firm heterogeneity vary across countries.

[3] In particular, see Hottman et al. (2016) and Redding and Weinstein (2024).

[4] Existing studies usually do not allow the parameters to vary across country-industry pairs, this resulting in a slightly worse fit. To further improve the fit of the log-normal distribution, some papers use a convolution with a second random variable (e.g., Kondo et al., 2023 or Sager and Timoshenko, 2019).

[5] Other papers have documented the importance of quality as a determinant of appeal and hence trade (e.g., Crinò and Ogliari, 2017; Feenstra and Romalis, 2014; Hallak and Schott, 2011; Khandelwal, 2010). The focus of this paper is the distribution of appeal and not its components.

and Pierola (2015) and Fernandes et al. (2016) have documented a number of important facts about the trade shares captured by top firms. Our framework clarifies that looking at market shares is not enough to quantify the contribution of heterogeneity, which also depends on structural parameters. In terms of results, the finding that richer countries have bigger exporters partly because their firms are more unequal is, to our knowledge, new.[6]

The paper is closely related to Redding and Weinstein (2024), who develop a log-linear decomposition to flexibly quantify the contribution of various firm-level margins in explaining variation of US imports across countries.[7] We complement their analysis in several ways. First, our main goal is to move beyond an accounting exercise by making distributional assumptions that allow us to identify new structural parameters. To illustrate the practical relevance of this contribution, we relate these structural parameters to measures of economic performance across countries and discuss their implications. Second, our focus is primarily on the role of firm heterogeneity. For this reason, our decomposition is designed to fully separate the effect of average and dispersion in the level of appeal. Due to the concavity of the log operator, the log-linear decomposition in Redding and Weinstein (2024) overstates the contribution of heterogeneity conditional on the mean. Finally, we use different data.[8]

The remainder of the paper is organized as follows. Section 2 introduces a theoretical framework that illustrates the role of firm heterogeneity when decomposing countries' sales. Section 3 describes the firm-level data on US imports that we use in the empirical analysis. Section 4 reports the results from a preliminary decomposition of US imports into firm-level margins. In Section 5 , we make distributional assumptions that allow us to identify new structural parameters and study how they relate to trade and economic performance. Section 6 discusses the welfare implication of the results in models of trade with heterogeneous firms and selection. Section 7 concludes.

2. Firm heterogeneity and trade accounting with CES preferences

We now show how to map countries' sales to a given destination in a given industry into firm-level characteristics. Our approach builds on previous work by Redding and Weinstein (2024). However, we depart from their log-linear decompositions to fully separate the role of average characteristics (in levels) relative to dispersion. The only restriction we impose in this section is a Constant Elasticity of Substitution (CES) demand system.[9] We assume that firms produce differentiated varieties and we identify each variety as a different technology. Since we are interested in studying how these technologies affect sales, in the empirical section, we will take the firm-product pair as the basic unit of analysis and we will refer to it simply as a "firm" or "variety".[10]

2.1. Preferences and demand

Consider consumers located in a destination d. In the empirical section, the destination will be the US market. Preferences over consumption of goods produced in I industries are:

$$U_d = \sum_{i=1}^{I} \beta_i \ln C_{di}, \quad \beta_i > 0, \quad \sum_{i=1}^{I} \beta_i = 1. \tag{1}$$

Each industry $i \in \{1, \dots, I\}$ produces differentiated varieties and consumers have CES preferences over these varieties:

$$C_{di} = \left\{ \sum_{\omega \in \Omega_{di}} \left[\gamma_d(\omega) c_d(\omega) \right]^{\frac{\sigma_i - 1}{\sigma_i}} \right\}^{\frac{\sigma_i}{\sigma_i - 1}}, \quad \sigma_i > 1, \tag{2}$$

where $c_d(\omega)$ is quantity consumed of variety ω, $\gamma_d(\omega)$ is a demand shifter, Ω_{di} denotes the set of varieties available for consumption in market d in industry i, and σ_i is the elasticity of substitution between varieties. In general, we use lowercase letters for variables referring to a single variety and uppercase letters for more aggregate variables. The demand shifter $\gamma_d(\omega)$ is often interpreted as "quality", because it measures the value of a certain product for a given quantity consumed. It captures both the intrinsic quality of the variety and its specific appeal in the destination market considered. Since we have data on one destination market only, we will not be able to distinguish between them. With this *caveat* in mind, from now on, we may refer to γ as quality.

We denote by $p_d(\omega)$ the price of variety ω and by P_{di} the minimum cost of one unit of the consumption basket C_{di}:

$$P_{di} = \left\{ \sum_{\omega \in \Omega_{di}} \left[\frac{p_d(\omega)}{\gamma_d(\omega)} \right]^{1-\sigma_i} \right\}^{\frac{1}{1-\sigma_i}}. \tag{3}$$

---

[6] Other papers have studied export patterns at the firm level by destination (e.g., Eaton et al., 2011; Mayer et al., 2014). We instead exploit variation across exporters serving the same destination market. In doing so, we also differ fundamentally from the literature studying how export affects the sales distribution in the country of origin (e.g., di Giovanni et al., 2011).

[7] In turn, Redding and Weinstein (2024) build on Hottman et al. (2016), who use barcode data from grocery stores to decompose sales across multi-product firms.

[8] In Bonfiglioli et al. (2021), we instead use the same data to study how concentration has changed among firms exporting to the US.

[9] CES preferences are a dominant paradigm in the literature and allow us to derive an exact decomposition of firms' sales. However, the qualitative results derived here do not depend on any specific preferences. Moreover, as discussed in Appendix C, the analysis in Section 5 applies to a set of demand systems more general than CES (see Mrázová and Neary (2017), Mrázová et al. (2021)).

[10] In this way, we do not impose any exogenous nesting structure between varieties produced by the same firm and across different firms. Similarly, we do not impose any restriction on the technology of multi-product firms. While studying product scope is also an interesting question, we feel that our data are not sufficiently disaggregated to do full justice to it. In untabulated results, we find, however, that using firms, rather than firm-products, as the basic unit of analysis has no bearing on our main conclusions.

Then, demand for a variety ω can be expressed as:

$$c_d(\omega) = p_d(\omega)^{-\sigma_i} \gamma_d(\omega)^{\sigma_i - 1} P_{di}^{\sigma_i} C_{di}. \tag{4}$$

As usual, demand is a negative function of the price, with elasticity σ_i. Conditional on prices, demand is increasing in quality, with elasticity $\sigma_i - 1$.

2.2. Decomposing sales

Consider an industry i, and define N_{doi} as the number of varieties sold from country o to destination d in industry i and $\bar{r}_{doi}$ as their average sales. Then, total sales originating from o are:

$$R_{doi} = N_{doi} \cdot \bar{r}_{doi}. \tag{5}$$

This equation decomposes countries' sales into the contribution of an extensive margin (the number of varieties) and an intensive margin (average sales of each variety).

Next, consider the intensive margin, $\bar{r}_{doi}$. From (4), we can write sales of variety ω as:

$$r_d(\omega) = \tilde{\gamma}_d(\omega)^{\sigma_i - 1} A_{di}, \tag{6}$$

where $\tilde{\gamma}_d(\omega) \equiv \gamma_d(\omega)/p_d(\omega)$ is defined as the quality-to-price ratio and $A_{di} \equiv P_{di}^{\sigma_i - 1} E_{di}$ summarizes demand conditions where E_{di} is expenditure allocated to sector i in destination d. Note that, within a given industry-destination, differences in sales depend solely on differences in quality-to-price ratios. For this reason, from now on, we refer to $\tilde{\gamma}_d(\omega)$ as a synthetic measure of "appeal". Average sales per variety can then be expressed as a function of the distribution of appeal:

$$\bar{r}_{doi} = \frac{1}{N_{doi}} \sum_{\omega \in \Omega_{doi}} \tilde{\gamma}_d(\omega)^{\sigma_i - 1} A_{di}. \tag{7}$$

We can now decompose the intensive margin, $\bar{r}_{doi}$, into two key statistics of the distribution of appeal: its average and an appropriate measure of dispersion. To this end, define the arithmetic mean of $\tilde{\gamma}_d(\omega)$ from a single origin o as:

$$\mathbb{E}\left[\tilde{\gamma}_{doi}\right] \equiv \frac{1}{N_{doi}} \sum_{\omega \in \Omega_{doi}} \tilde{\gamma}_d(\omega). \tag{8}$$

Then, sales of the firm with average appeal are $\mathbb{E}\left[\tilde{\gamma}_{doi}\right]^{\sigma_i - 1} A_{di}$. Heterogeneity matters when average sales differ from the sales of the average firm. To quantify its role, we add and subtract $\mathbb{E}\left[\tilde{\gamma}_{doi}\right]^{\sigma_i - 1} A_{di}$ to (7) so as to obtain:

$$\bar{r}_{doi} = \left[\mathbb{E}\left[\tilde{\gamma}_{doi}\right]^{\sigma_i - 1} + \frac{1}{N_{doi}} \sum_{\omega \in \Omega_{doi}} \left\{ \tilde{\gamma}_d(\omega)^{\sigma_i - 1} - \mathbb{E}\left[\tilde{\gamma}_{doi}\right]^{\sigma_i - 1} \right\} \right] A_{di}. \tag{9}$$

This equation shows that average sales per variety can be decomposed into two terms. The first term in square brackets captures the contribution of the average appeal of varieties from a given country. The second term captures the importance of heterogeneity in appeal from that origin. Clearly, (9) shows that the contribution of heterogeneity is zero if all quality-to-price ratios from a given country are identical. But what is the sign of this term if appeal does vary across varieties? The answer to this question depends on the value of σ_i, because the latter captures the willingness to reallocate demand towards better varieties.

To see how, note from (6) that sales are a convex function of appeal when $\sigma_i > 2$. In this case, by Jensen's inequality, the contribution of heterogeneity in (9) is positive. When $\sigma_i = 2$, instead, sales are linear in appeal, so that only its average, and not its distribution, matters. Finally, when $\sigma_i < 2$, sales are a concave function of appeal, so that more heterogeneity has a negative contribution to average sales. In other words, if varieties are sufficiently substitutable, we are in a "superstar economy", in the sense that the possibility to reallocate expenditure from less to more appealing varieties increases average sales when holding constant the mean quality-to-price ratio. Note that, while $\sigma_i > 1$ is sufficient for demand to increase more than proportionally with $\tilde{\gamma}$, the elasticity of revenues is $\sigma_i - 1$ because of the negative effect of quantity on prices. This is why an elasticity greater than two is needed for revenues to increase more than proportionally with $\tilde{\gamma}$.

Interestingly, we can also rewrite (9) in terms of the distribution of observed sales and σ_i. In particular, substituting $\tilde{\gamma}_d(\omega)$ from (6), yields:

$$\bar{r}_{doi} = \left\{ \mathbb{E}\left[r_{doi}^{1/(\sigma_i - 1)} \right] \right\}^{\sigma_i - 1} + \frac{1}{N_{doi}} \sum_{\omega \in \Omega_{doi}} \left\{ r_d(\omega) - \left\{ \mathbb{E}\left[r_{doi}^{1/(\sigma_i - 1)} \right] \right\}^{\sigma_i - 1} \right\}, \tag{10}$$

where $\mathbb{E}\left[r_{doi}^{1/(\sigma_i - 1)} \right] \equiv \sum_{\omega \in \Omega_{doi}} r_{doi}^{1/(\sigma_i - 1)} / N_{doi}$.

Eqs. (5) and (10) can be used to decompose the sales originating from any country or group of countries. All is needed is estimates of σ_i and the observed firm-level sales from any country or group of countries. Note also that while documenting the properties of the distribution of sales, such as the variance or the role of top firms, is an interesting exercise, it is not sufficient to assess the importance of firm heterogeneity, i.e., of having firms that differ from the country's average. For instance, (10) shows that if $\sigma_i = 2$ heterogeneity is irrelevant, regardless of how sales are distributed. However, this also highlights the main limitation of the exercise, which hinges on having reliable estimates of the parameter σ_i. Finally, although the second term on the right-hand side of (10) captures the effect of reallocation among heterogeneous firms, which depends both on the distribution of sales and σ_i, in what follows we will often refer to it simply as a measure of "heterogeneity".

Table 1
Descriptive statistics.

	Mean	Median	Std. Dev.
(a) Sample coverage			
Share of Piers exports in total exports to the US (based on customs data)	0.83	0.77	0.55
(b) Sample composition			
N. of firms	44	8	249
N. of firm-product pairs (varieties)	55	9	316
Total exports ($1000)	60 347	2360	536 000
Average exports per variety ($1000)	1273	230	11 058
(c) Moments			
Variance of log sales	3.54	3.07	3.12
Variance of log prices	0.41	0.19	0.67

Notes. The variable in panel (a) is computed for each country in the years 2002 and 2012. Reported statistics are the mean, median and standard deviation of the variable across all countries and years. The variables in panels (b) and (c) are computed for each country-industry-year triplet. Reported statistics are the mean, median and standard deviation of each variable across all triplets.

3. The data

To perform the empirical analysis, we need data on the sales of individual products in a single destination market by firms from different origin countries. We obtain this information using transaction-level data on US seaborne imports from Piers, a database administered by IHS Markit.[11] Piers contains the complete detail of the bill of lading of any shipment that is imported into the US by sea. IHS markit collects the bills of lading filed with the US Customs, verifies and standardizes their information, and makes the resulting data available for sale. We purchased from IHS Markit information on the universe of seaborne manufacturing import transactions of the US, by exporting firm and product, in two years, 2002 and 2012. For each transaction, we know the complete name of the exporting firm, its country of origin, the exported product (according to the 6-digit level of the HS classification), the value (inclusive of freight, insurance and other transportation charges, in US dollars) and the quantity (in kilograms) of the transaction; we compute the unit value of each transaction as the ratio of value to quantity.[12] We use a correspondence table developed by the World Integrated Trade Solutions to map each HS6 product exported by a firm into a 4-digit SIC industry.[13] The final data set comprises 1,350,574 observations at the firm-product-year level. Firms belong to 366 4-digit SIC manufacturing industries and 104 origin countries spanning the five continents.

Three features of Piers are especially relevant for our analysis. First, access to Piers is not restricted and can be obtained by anyone, albeit at a fee. Second, all firms in Piers use the same export mode (by sea), which favors comparability. Third, Piers contains the full name of each firm, which allows us to precisely identify companies exporting to the US, thereby reducing the risk of over-counting them. These characteristics enable us to compare export flows, and their firm-level determinants, from virtually all countries in the world to the same destination market.

Since our decompositions are valid for any subset of firms and sales, the fact that Piers only contains seaborne shipments implies that our results should be interpreted, strictly speaking, as applying to firms selling in the US by sea. At the same time, maritime trade is the principal trade mode worldwide, accounting for 70% of world trade values and 80% of world trade volumes (UNCTAD, 2018). It also represents the bulk of US trade in manufacturing and the main mode of selling goods in the American market for most countries: as shown in panel (a) of Table 1, Piers accounts for 83% of overall manufacturing exports to the US for the average country in our sample, and for 77% for the median country.[14] In Appendix A.1, we discuss the countries for which Piers provides a less extensive coverage of exports to the US. In Section 4.1, we show that excluding these countries has no bearing on the results.

Panel (b) of Table 1 provides further details on the composition of our sample. All variables in this panel are computed separately for each country-industry-year triplet, and reported statistics are calculated across all available triplets. The average triplet has 44 firms and 55 varieties, a value of total exports to the US exceeding $60 million and average exports per firm-product slightly above $1 million. Appendix A.1 provides more details on data construction, discusses the advantages and limitations of Piers relative to the (restricted-access) transaction-level data collected by the US Customs, and compares a number of moments obtained from Piers with their counterparts based on aggregate trade data from various sources.

We now present some new stylized facts about how sales vary across countries and industries. In panel (c) of Table 1, we report summary statistics on two moments, the variance of log sales and the variance of log prices in each country-industry-year triplet.

[11] Since 2022, Piers is administered by S&P Global, as a result of its merger with IHS Markit.

[12] In the case of firms with multiple shipments (bills of lading) of the same product in a given year, we purchased from IHS Markit information on the total value and quantity of these shipments across all bills of lading, but not the detailed information on each bill of lading, which would have been prohibitively expensive.

[13] The 4-digit level of industry aggregation strikes a balance between number and comparability of products within each industry, as required by the methods we use for estimating the elasticities of substitution in Section 4.1.

[14] These numbers are similar to figures reported by Feenstra and Weinstein (2017) for an earlier and more limited vintage of the Piers database. See also Flaaen et al. (2023) for further discussion on the main features of bill of lading databases like Piers.

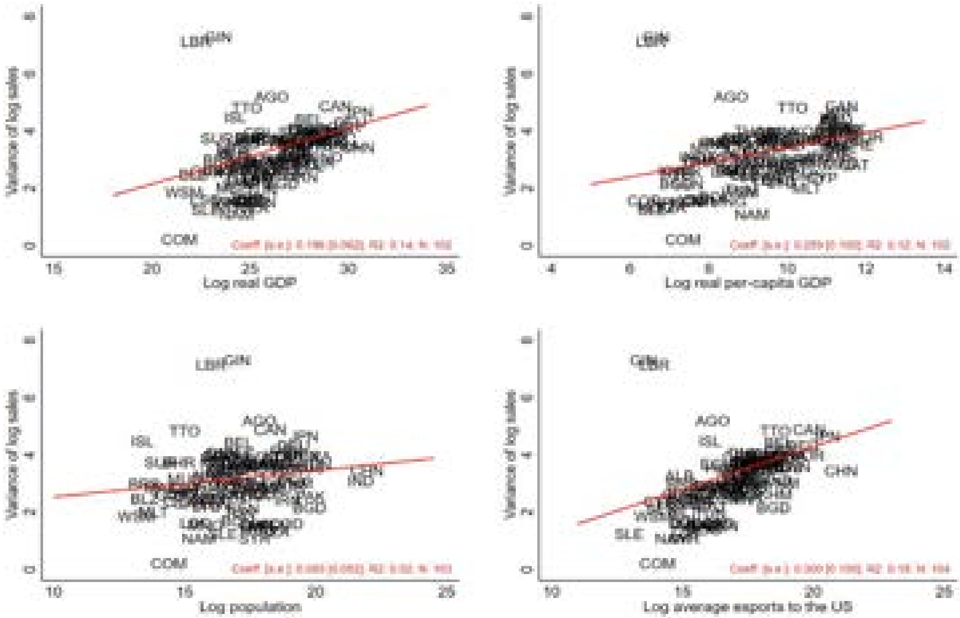

Fig. 1. Sales Dispersion and Country Characteristics.
Notes. Variance of log sales and average exports to the US are computed for each country-industry-year triplet and then averaged across all industries and years for each country. The other country characteristics are simple averages between the years 2002 and 2012.

Reported statistics are the mean, median and standard deviation of each variable across all triplets. Sales dispersion is high and varies markedly across triplets. It has also significantly increased over the sample period, by 10% on average across countries and industries (untabulated).[15] Given that we know the identity of firms, we can also compute the change in sales dispersion driven by reallocation among firms active in both years. In the subsample of continuing varieties, sales dispersion has increased by 29% while in the rest of the sample it has increased by approximately 8%.[16] Conversely, price dispersion is relatively small and exhibits low cross-sectional variation. Unreported figures also indicate that price dispersion has remained stable over time.

To have a first sense of how firm heterogeneity varies across countries and correlates with trade and economic performance, Fig. 1 shows the cross-country relationships between sales dispersion and the log of four country characteristics: real GDP, real per-capita GDP, population and average exports to the US.[17] The variance of log sales is computed for each country-industry-year triplet and is then averaged across all industry-years within a country to neutralize compositional effects due to differences in the industrial structure of production. The first graph shows that sales dispersion is strongly positively correlated with real GDP. The second and third graphs dig into this relationship by dividing real GDP in its two components: real per-capita GDP and population. The plots highlight that sales are significantly more dispersed in richer and, to a lesser extent, larger countries. Finally, the fourth graph studies the relationship between sales dispersion and average exports to the US, computed as the mean value of exports across all industries and years for each country. The plot shows that countries in which sales are more dispersed across firms export more to the US on average.

4. Decomposing US imports: Preliminary evidence

4.1. Firms, average appeal and heterogeneity

We now implement the decompositions presented in Section 2.2. For this we need an estimate of the elasticity of substitution between varieties in each industry, σ_i. For our baseline decomposition, we source these parameters from Broda and Weinstein (2006), who extend the estimation approach proposed by Feenstra (1994) to accommodate measurement error and deal with infeasible values.[18] We have estimates of σ_i for 356 out of 366 industries; all estimates satisfy the theoretical restriction that $\sigma_i > 1$. In the

[15] These results are in line, both qualitatively and quantitatively, with evidence based on US firm-level sales data and on cross-country product-level export data (Bonfiglioli et al., 2018, 2019).

[16] Continuing varieties account for 28% of total exports to the US in the average country-industry pair in 2012.

[17] GDP and population data are sourced from the World Development Indicators.

[18] Feenstra (1994) proposes to transform the supply and demand functions of each variety by taking differences both over time and relative to a reference good, and then to identify the elasticity of substitution using moment conditions arising from the assumption of disturbance independence across the two transformed equations. Broda and Weinstein (2006) complement this approach both with a grid search to deal with imaginary or wrongly signed estimates and with a

Table 2
Decomposition of countries' exports.

	First Step		Second Step	
	Extensive Margin	Intensive Margin	Average	Dispersion
	(1)	(2)	(3)	(4)
(a) Baseline	0.466***	0.534***	0.403***	0.597***
	(0.003)	(0.003)	(0.119)	(0.119)
(b) No small countries	0.478***	0.522***	0.378***	0.622***
	(0.003)	(0.003)	(0.101)	(0.101)
(c) No countries with small market shares	0.505***	0.495***	0.403***	0.597***
	(0.006)	(0.006)	(0.121)	(0.121)
(d) No countries with large market shares	0.464***	0.536***	0.374***	0.626***
	(0.003)	(0.003)	(0.094)	(0.094)
(e) No industries with high shares of imported inputs	0.461***	0.539***	0.448***	0.552***
	(0.003)	(0.003)	(0.150)	(0.150)
(f) No trading companies	0.456***	0.544***	0.382***	0.618***
	(0.003)	(0.003)	(0.118)	(0.118)
(g) No top-1 firm in each triplet	0.506***	0.494***	0.319***	0.681***
	(0.003)	(0.003)	(0.122)	(0.122)
(h) No firms with sales 2SD above triplet average	0.498***	0.502***	0.549***	0.451***
	(0.003)	(0.003)	(0.155)	(0.155)
(i) Alternative elasticity of substitution - Reverse-Weighting	0.471***	0.529***	0.470***	0.530***
	(0.003)	(0.003)	(0.068)	(0.068)
(j) Alternative elasticity of substitution - Hummels	0.480***	0.520***	0.571***	0.429***
	(0.004)	(0.004)	(0.160)	(0.160)

Notes. Columns (1) and (2) perform the decomposition in Eq. (5). Each coefficient is obtained from a separate regression, run across country-industry-year triplets, of the log of the corresponding margin on the log of total exports to the US in the triplet, controlling for industry-year fixed effects. Columns (3) and (4) perform the decomposition in Eq. (10). Each coefficient is obtained from a separate regression, run across country-industry-year triplets, of the corresponding margin on average exports to the US in the triplet, controlling for industry-year fixed effects. Panel (a) uses the whole sample of triplets (24454 observations). Panel (b) uses the sample that excludes countries for which the share of Piers exports in total exports to the US is smaller than 45%, i.e., the first group of countries in map (b) of Figure A1 (21192 observations). Panel (c) uses the sample that excludes countries whose market shares fall below the 5th percentile of the distribution in a given industry and year (15206 observations). Panel (d) uses the sample that excludes countries whose market shares fall above the 95th percentile of the distribution in a given industry and year (23964 observations). Panel (e) uses the sample that excludes industries for which the average share of imports of parts and components in total US imports over 1972–2001 is above 25% (19765 observations). Panel (f) uses the sample that excludes firms whose name contains words starting with "trad", "logist", "transp", "export" or "import" (23998 observations). Panel (g) uses the sample that excludes the top-1 firm in each triplet (20632 observations). Panel (h) uses the sample that excludes firms whose total exports to the US are at least two standard deviations above the average exports for their triplet (24454 observations). Panel (i) uses estimates of the elasticity of substitution obtained using the Reverse-Weighting estimator (Redding and Weinstein, 2024) on the transaction-level data from Piers (18880 observations). Panel (j) uses estimates of the elasticity of substitution obtained using Hummels' (2001) approach on the transaction-level data from Piers (12904 observations). The standard errors are corrected for heteroskedasticity. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

average industry, σ_i equals 4.21, while it is equal to 2.79 in the median industry; Appendix Table B1 reports the median values of σ_i in each 2-digit SIC sector. Regarding the distribution of σ_i across industries, the 10th, 25th, 75th and 90th percentiles are equal to 1.80, 2.18, 3.90 and 5.57, respectively. Importantly, the estimates of σ_i are larger than 2 for the vast majority of industries (294, i.e., 83% of the total). As explained in the previous section, this implies that greater heterogeneity in appeal should translate into higher average sales. For robustness, we also use two alternative sets of estimates of σ_i obtained on our transaction-level data, which however are available for fewer industries (more details below).

We start by decomposing countries' exports to the US into the contributions of the extensive and intensive margin. To this purpose, we take logs of (5) and run separate regressions, across all available country-industry-year triplets, of $\ln N_{doi,t}$ and $\ln \bar{r}_{doi,t}$ on $\ln R_{doi,t}$, where the subscript t stands for time (the year 2002 or 2012) and will henceforth be used in all equations that are taken to the data. In all regressions, we control for industry×year fixed effects, so that our decomposition focuses on variation in sales across countries within each industry and year.[19] The properties of OLS imply that the coefficients on $\ln R_{doi,t}$ obtained from these regressions add up to one, and thus provide the percentage contribution of each margin to explaining variation in countries' exports to the US. We similarly decompose the intensive margin $\bar{r}_{doi,t}$ into the contribution of average appeal and dispersion in appeal, by regressing each term on the right-hand side of (10) on $\bar{r}_{doi,t}$ and industry×year fixed effects.

The results are reported in Table 2. Panel (a) shows the decompositions performed on the full sample. The estimates in columns (1) and (2) refer to the contributions of the extensive and intensive margin. Each margin explains roughly half of the variation in

weighting of the GMM objective function to accommodate measurement error in unit values. The Feenstra (1994) and Broda and Weinstein (2006) approaches require more than two years of data for each variety and thus cannot be applied to our transaction-level data. We draw from Broda and Weinstein (2006) elasticity estimates at the 10-digit level of the HS classification based on US import data for the 1990–2001 period. Then, we compute the median elasticity in each 4-digit SIC industry across all HS10 products belonging to it.

[19] Controlling for industry×year fixed effects also neutralizes any change in product and industry classifications, or in their mapping, over time.

countries' exports to the US. Hence, countries that export more to the US within a given industry and year do so because they sell both a larger number of varieties and more of each variety, with the contributions of the two margins being roughly equivalent in our data. The estimates in columns (3) and (4) refer instead to the decomposition of the intensive margin into the contributions of average appeal and its heterogeneity. The results show that reallocation between heterogeneous firms contributes similarly to average appeal to explaining variation in average exports per variety. Hence, firm heterogeneity is an important factor for understanding countries' export performance.

These patterns are consistent across a large number of robustness checks, which are presented in the remaining panels of Table 2 . In panel (b), we exclude countries for which Piers covers less than 45% of total exports to the US and find that the coefficients hardly change. In panels (c) and (d), we exclude countries with small and large market shares, respectively. The former (latter) are market shares falling below the 5th (above the 95th) percentile of the distribution of market shares in a given industry and year. These exercises show that the decompositions are not driven by either small or large exporters. In panel (e), we exclude industries for which the share of imports of parts and components in total US imports is above 25%.[20] Also in this case, the results are largely unchanged, suggesting that the decompositions are not driven by industries in which US imports predominantly consist of intermediate inputs reflecting firms' participation in global value chains. In panel (f), we drop trading companies, which we identify as firms whose name contains words starting with "trad", "logist", "transp", "export" or "import".[21] The main patterns are unaffected, suggesting that the decompositions reflect heterogeneity among direct exporters rather than the role of shipping firms.

One reason for the importance of firm heterogeneity in explaining sales is the presence of top firms in each country. As long as these "superstar firms" are exceptional, i.e., they have significantly better appeal compared to the remaining firms in an industry, their presence would be associated with both high sales and high dispersion. It is known that top firms can define the export performance of a sector. As in previous studies, the top firm in each country plays a dominant role also in our sample, accounting for 25% of total exports to the US, on average.[22] But are these firms really outliers, in the sense that they significantly affect the quantitative role of dispersion in explaining exports? A simple way of answering this question is to remove top firms from each triplet and redo the decompositions using the remaining sample of firms. In panel (g), we exclude the top-1 firm, defined as the firm with the highest sales across products in the triplet. In panel (h), we instead remove all firms whose total exports (across all products) are at least two standard deviations above the average exports for their triplet. The decomposition of the extensive and intensive margin is essentially unchanged, while the contribution of heterogeneity in appeal decreases, at most, only marginally. This suggests that the importance of heterogeneity for explaining the intensive margin of countries' exports does not merely reflect the presence of outliers.

As already discussed, the importance of heterogeneity for the decomposition of average exports depends on the elasticity of substitution, σ_i. To study the sensitivity of our decompositions to this parameter, we now repeat the analysis using two alternative sets of values of σ_i, which we directly estimate on the transaction-level data from Piers using the Reverse-Weighting (RW) estimator introduced by Redding and Weinstein (2024) and the approach proposed by Hummels (2001).[23] We obtain elasticity estimates for 259 industries with the RW estimator and for 163 industries with Hummels' (2001) approach. The median values of these elasticities are larger than 2 in all 2-digit SIC sectors (see Appendix Table B1). Moreover, despite the use of different data and estimation methods, all estimates of σ_i are highly correlated among each other, with pairwise correlations among the three sets of estimates ranging from 0.6 to 0.99 across sectors.

The results of the decompositions using the alternative estimates of σ_i are reported in panels (i) and (j) of Table 2. The coefficients in columns (1) and (2) are slightly different from their counterparts in panel (a) due to the use of smaller samples; reassuringly, however, this has almost no bearing on the quantification of the contributions of the extensive and intensive margin. More importantly, the coefficients in columns (3) and (4) are also in line with those in panel (a), even though the elasticity of substitution influences these coefficients directly, and not just through sample size. This confirms that both average appeal and its heterogeneity explain roughly half of the variation in average exports per variety. Overall, these results suggest that the decompositions are robust to the sample and the empirical methods used to estimate the elasticity of substitution. Nevertheless, in Section 5, we present an approach that does not rely on estimates of σ_i.

We now pause to briefly discuss the relationship between these results and the existing literature. The fact that the extensive margin explains about half of the variation in trade flows is consistent with previous findings (e.g., Fernandes, Freund and Pierola, 2016, (Fernandes et al., 2023; Redding and Weinstein, 2024)). The contribution of firm heterogeneity in affecting the volume of trade has received less attention. Redding and Weinstein (2024), who develop an alternative decomposition of US imports, find that dispersion of firm attributes accounts for 36% of the variation in measures of revealed comparative advantage. However, their

[20] We use data on imports of parts and components from Schott (2004) for the pre-sample period 1972–2001.

[21] These companies represent 9% of firms in our sample. See Flaaen et al. (2023) for further discussion on the role of trading companies using bill of lading data on US imports for more recent years.

[22] These findings are consistent with results obtained by Freund and Pierola (2015) for a sample of developing countries.

[23] See Appendix E for a description of the RW estimator. As for Hummels' (2001) approach, the author shows that, in a multi-sector model of trade with monopolistically competitive firms, the elasticity of substitution can be estimated by regressing the log of imports of each variety on the log sum of its freight rate and import tariff, controlling for exporting country fixed effects, importing country fixed effects and standard gravity variables. To implement this approach using the Piers data, we regress, separately for each 4-digit SIC industry, the log of imports of each variety on the log of one plus the freight rate and import tariff of the corresponding HS6 product imported into the US from the firm's origin country in a given year (see Appendix A.2 for details on data sources and variables definitions). We control for exporting country×year fixed effects, which subsume all gravity variables given that our data refer to a single importing country. For each 4-digit SIC industry, the elasticity of substitution is then obtained as the coefficient on the log sum of freight rate and import tariff times minus one.

Table 3
Intensive margin of exports and country characteristics.

	Country-Level Averages			Country Fixed Effects		
	Intensive Marg. (1)	Average (2)	Dispersion (3)	Intensive Marg. (4)	Average (5)	Dispersion (6)
Real per-capita GDP	0.354***	0.187***	0.167***	0.298**	0.098	0.199***
	(0.117)	(0.061)	(0.057)	(0.117)	(0.061)	(0.059)
Population	0.094	0.037	0.057	0.117	−0.006	0.123***
	(0.084)	(0.044)	(0.043)	(0.082)	(0.043)	(0.043)
Obs.	102	102	102	102	102	102
R2	0.12	0.12	0.12	0.10	0.04	0.19

Notes. The table reports cross-country regressions of the three terms in Eq. (10) on country characteristics. In particular, in columns (1)–(3), the dependent variables indicated in the columns' headings are constructed separately for each country–industry–year triplet and then averaged at the country level. In columns (4)–(6), the dependent variables are the country fixed effects obtained by regressing each of the variables indicated in the columns' headings on country and industry-year fixed effects across triplets. Real per-capita GDP and population are simple averages of these variables between the years 2002 and 2012. All explanatory variables are in logs and all dependent variables are in US$ million. The standard errors are corrected for heteroskedasticity. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

log-linear decomposition holds constant the mean of the *log* of firm attributes, which is negatively affected by the dispersion of the *level* of attributes. Once the effect of the mean and dispersion of the level of attributes is fully separated, we find that the contribution of heterogeneity is reduced to around 25%.

4.2. Average exports and country characteristics

Our preliminary evidence suggests sales from larger and richer countries to be more dispersed. At the same time, these countries have been shown to have higher exports per firm (see, for instance, (Fernandes et al., 2016). We now use our decomposition to ask: are average exports per firm higher in larger and richer countries because the latter have better firms or because they have more heterogeneous firms? To tackle this question, we start by computing the arithmetic mean of each of the three terms in (10) across all industries and years for each country. Then, we regress each term on the log of countries' population and real per-capita GDP.[24] The coefficients obtained from the regressions for average appeal and dispersion in appeal add up to the coefficients obtained from the regression for average exports per variety. Hence, they provide an additive decomposition of the margins along which population and income per capita correlate with the intensive margin of countries' exports.

The results are reported in Table 3. Column (1) confirms that income per capita is strongly positively correlated with average exports per variety. The coefficient on population is also positive, albeit imprecisely estimated. Columns (2) and (3) decompose the coefficients in column (1) into the contributions of average appeal and its dispersion, respectively. The results suggest that both average appeal and heterogeneity have a quantitatively similar importance for explaining the correlations of exports with population and income per capita.

Next, to better isolate the role played by the distribution of firm characteristics within industries, we regress each of the three terms in (10) on country and industry×year fixed effects. Because the industry×year effects absorb average differences in a given term across industry-years, the country effects are identified from the comparison of countries exporting to the US in the same industry and year. As such, the country effects reflect the average value of a given term net of effects due to differences in the composition of export industries across countries. We then repeat the specifications in columns (1)-(3) using as dependent variables the country effects rather than the simple country-level averages of the three terms.

The results are reported in columns (4)-(6) of Table 3. The estimated coefficients confirm the positive correlation of average exports per variety with income per capita and population, with the coefficient on the latter variable being imprecisely estimated. However, both correlations are now mostly driven by heterogeneity in appeal. These results indicate that, once compositional effects are neutralized, the main reason why richer countries export more per firm is dispersion, i.e., the prevalence of above-average firms within industries.

5. Firm heterogeneity and trade with log-normal distributions

The results so far suggest that the degree of firm heterogeneity varies significantly across countries, even within narrowly-defined industries, and that this variation may represent an important source of economic performance: it correlates with exports and GDP. However, these results are subject to three limitations. First, the degree of firm heterogeneity computed on observed sales need not reflect just technological factors. In particular, since only a subset of domestic firms sell into a given foreign destination, selection into exporting will directly affect the variance of their sales. Second, another concern is that the results may be affected by the granularity of the data. The third limitation is that the quantifications rely on estimates of the elasticity of substitution, σ_i.

[24] Population and per-capita GDP are averaged between the years 2002 and 2012.

We now want to separate the effect of differences in the primitive distribution of appeal from that of truncation, granularity and σ_i. Since the effects of truncation and granularity depend on the shape of the distribution, to make progress, we need to make some additional structural assumptions. Following several papers in the literature, we start from the hypothesis that sales are well approximated by log-normal distributions. Under this assumption, we can estimate the parameters of the unconditional distribution of exports, and compare them to the moments of the data in a way that allows us to assess the importance of selection and granularity while controlling for cross-industry differences in σ_i.

5.1. Sales with log-normal distributions: Theory

We start by showing the mapping between the hypothetical (unconditional) distribution from which appeal is drawn and observed sales in the log-normal case. Assume that, in a given country and industry, the log of appeal, $\ln \tilde{\gamma}_d(\omega)$, is drawn from a normal distribution with mean $\widetilde{\mu}_{doi}$ and variance $\widetilde{\varsigma}^2_{doi}$. Since sales in destination d are a power function of appeal, $r_d(\omega) = \tilde{\gamma}_d(\omega)^{\sigma_i - 1} A_{di}$, they will be log-normal$\left(\mu_{doi}, \varsigma^2_{doi}\right)$, where:

$$\mu_{doi} = \left(\sigma_i - 1\right) \widetilde{\mu}_{doi} + \ln A_{di} \tag{11}$$

$$\varsigma^2_{doi} = \left(\sigma_i - 1\right)^2 \widetilde{\varsigma}^2_{doi}. \tag{12}$$

Note that sales dispersion depends both on the variance of log appeal and on the elasticity of substitution. This is because differences in appeal translate into larger differences in demand the more substitutable varieties are. However, the effect of σ_i on the *log* of ς^2_{doi} is industry specific, and can thus be fully absorbed by controlling for industry fixed effects. Note also that these results apply whenever sales are a power function of firm characteristics. As discussed in Appendix C, this can be the case under a class of demand functions more general than CES.

Assume next that, as in standard heterogeneous-firms trade models with a fixed export cost, sales are truncated from below at $r^{\min}_{doi}$. In this case, the observed mean, $\mathbb{E}\left(r_{doi}\right)$, is the conditional expectation for $\tilde{\gamma}(\omega)^{\sigma_i - 1} A_{di} \geq r^{\min}_{doi}$:

$$\mathbb{E}\left(\ln r_{doi}\right) = \mu_{doi} + \varsigma_{doi} \frac{\phi\left(\theta_{doi}\right)}{1 - \Phi\left(\theta_{doi}\right)}, \tag{13}$$

with $\theta_{doi} \equiv (\ln r^{\min}_{doi} - \mu_{doi})/\varsigma_{doi}$, and where ϕ and Φ are the density and the CDF of the standard normal distribution, respectively. Naturally, average observed sales are increasing in the minimum $r^{\min}_{doi}$.

While log sales are drawn from a normal distribution with variance ς^2_{doi}, the variance of observed sales also depends on truncation:

$$\mathbb{V}\left(\ln r_{doi}\right) = \varsigma^2_{doi} \left[1 + \frac{\theta_{doi}\phi\left(\theta_{doi}\right)}{1 - \Phi(\theta_{doi})} - \left(\frac{\phi\left(\theta_{doi}\right)}{1 - \Phi(\theta_{doi})}\right)^2\right]. \tag{14}$$

Since truncation is a mean-preserving contraction combined with a mean-changing rigid shift, the variance of the truncated distribution is less than the variance of the original normal distribution. Hence, the ratio $\varsigma^2_{doi}/\mathbb{V}\left(\ln r_{doi}\right)$ is increasing in the cutoff, $\ln r^{\min}_{doi}$: intuitively, as smaller sales are removed, the variance of the remaining distribution falls.

In what follows, we will use these results to argue that cross-country differences in $\widetilde{\varsigma}^2_{doi}$, i.e., differences in dispersion of the unconditional distribution of appeal, are important for explaining sales. To do so, we will first show that sales are well approximated by log-normal distributions. Second, we will use the properties of log-normal distributions to estimate the parameter ς^2_{doi}. Third, we will use industry fixed effects to net out cross-country differences in the variance of log sales from the influence of σ_i. Finally, we will use the ratio $\varsigma^2_{doi}/\mathbb{V}\left(\ln r_{doi}\right)$ to control for the cutoff, and hence to separate the effect of selection.

5.2. Sales with log-normal distributions: Estimation

We now provide formal statistical tests that our data are well approximated by the log-normal distribution.[25] Given that our unit of analysis is the country–industry–year triplet, we perform the tests on each triplet separately. The results are shown in the first four columns of Table 4. Column (1) reports the number of triplets on which the tests are run. Column (2)-(4) indicate the fraction of triplets for which the null hypothesis of the tests, i.e., that log exports to the US are drawn from a normal distribution, is not rejected at the 5% significance level. We henceforth refer to these cases as “LN triplets” for brevity.

We start with the Kolmogorov–Smirnov (KS) test, which non-parametrically compares the empirical distribution of log exports to the US with the theoretical normal distribution.[26] As shown in column (2), panel (a), the null hypothesis is not rejected in 98% of the 24754 triplets in our sample. In columns (3) and (4), we complement the KS test with two tests proposed by D’Agostino et al. (1990), which are designed to detect deviations from the normal distribution associated with skewness (asymmetric tails) or non-normal kurtosis (thick or light tails). The fraction of LN triplets remains well above 80% in both cases, ranging from 82% for the skewness test to 87% for the kurtosis test.[27]

[25] See Appendix Figure B1 for preliminary visual evidence.

[26] Given a random variable X, with empirical distribution function $\hat{F}_n(x)$ over n i.i.d. ordered observations of X, the KS statistics is $D_n = \sup_x \left|\hat{F}_n(x) - F(x)\right|$, where $F(x)$ is a theoretical (continuous) distribution function. Intuitively, D_n provides the largest difference, in absolute value, between the empirical distribution function and the theoretical distribution function. Under the null hypothesis that the sample of observations of X is drawn from the theoretical distribution,

Table 4
Tests for log normality of exports.

	Statistical Tests				R2 from QQ Regressions		Coefficient on Theoretical Quantiles (R2) from QQ Regressions			
	Obs.	% LN (KS)	% LN (Sk)	% LN (Ku)	Mean	Std. Dev.	All	Top 75%	Top 50%	Top 25%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
(a) All triplets	24 754	0.98	0.82	0.87	–	–	–	–	–	–
(b) Triplets with 9+ varieties	12 437	0.96	0.78	0.83	0.95	0.05	1.93 (0.96)	1.93 (0.96)	1.98 (0.96)	2.05 (0.95)
(c) Triplets with 27+ varieties	6314	0.92	0.67	0.75	0.96	0.03	1.97 (0.97)	1.98 (0.98)	2.05 (0.97)	2.14 (0.96)

Notes. Column (1) reports the number of country–industry–year triplets on which the tests for the normality of log exports to the US are performed. Column (2) reports the fraction of triplets for which the null hypothesis of normality is not rejected at the 5% level of significance according to the Kolmogorov–Smirnov test. Columns (3) and (4) report analogous results for the D'Agostino et al. (1990) test based on the skewness and kurtosis of the distribution. The sample consists of all available triplets (panel a) and of triplets with at least 9 (panel b) or 27 (panel c) firm–products exported to the US; in panel (a), the D'Agostino et al. (1990) tests are performed on the sample of 16893 triplets with at least 5 firm–products exported to the US, as the test statistics are not defined for smaller triplets. Columns (5)–(6) report descriptive statistics on the R2 obtained by estimating the QQ regression on each triplet, using theoretical quantiles of sales implied by the log-normal distribution; reported statistics are the mean and standard deviation of the R2 across all triplets. Columns (7)–(10) report the coefficient on the theoretical quantiles and the R2 (in parenthesis) obtained by estimating the QQ regression on each triplet, using either the whole sample of firm–products in the triplet (column 7) or the sub-sample of firm–products in the top 75th, 50th and 25th percentile of the sales distribution in the triplet (columns 8–10); reported statistics are the median coefficient and the median R2 across all triplets.

One reason for the large fraction of LN triplets in panel (a) could be the limited power of the KS and D'Agostino et al. (1990) tests in small samples. Accordingly, we repeat the tests on two less granular samples, which consist of triplets containing at least 9 varieties (i.e., the sample median, as shown in Table 1) or 27 varieties (the top quartile), respectively. The results are reported in panel (b) of Table 4 for the sample of triplets with 9+ varieties and in panel (c) for the sample of triplets with 27+ varieties.[28] As expected, the fraction of LN triplets is slightly reduced as we progressively focus on larger samples. Yet, it still exceeds 90% for the KS test, and oscillates between two-thirds and four-fifths of the sample for the D'Agostino et al. (1990) tests.

The above results speak in favor of the log-normal distribution. This is important because, under the assumption of log-normality, we can estimate the parameters of the unconditional distribution of exports. In particular, to estimate the variance, ς^2_{doi}, we follow Head et al. (2014) and exploit a linear relationship linking theoretical and empirical quantiles of log sales. This approach is known as the Quantile-Quantile (QQ) estimator, and its asymptotic properties are studied in Kratz and Resnick (1996).

We start by ranking varieties within each triplet in ascending order of their sales. Let $\ln r_{doi,t}(\omega)$ denote the log sales in the US of variety ω, with $\omega = 1$ indicating the variety with the minimum sales and $\omega = N_{doi,t}$ the variety with the maximum sales in the triplet. The empirical quantiles of sorted log sales within the triplet are $\mathbb{Q}^E_{doi,t}(\omega) = \ln r_{doi,t}(\omega)$, while the empirical CDF of log sales is given by $\hat{F}_{doi,t}(\omega) = (\omega - 0.3) / \left(N_{doi,t} + 0.4\right)$. The theoretical quantiles are defined as:

$$\mathbb{Q}^T_{doi,t}(\omega) = \mathbb{E}\left(\ln r_{doi,t}\right) + \varsigma_{doi,t}\Phi^{-1}(\hat{F}_{doi,t}(\omega)). \tag{15}$$

The QQ estimator regresses the empirical quantiles on the theoretical quantiles, so the variance of the unconditional distribution can be recovered from the coefficient on $\Phi^{-1}(\hat{F})$. We run a separate regression of $\mathbb{Q}^E_{doi,t}(\omega)$ on $\mathbb{Q}^T_{doi,t}(\omega)$ for each triplet, and recover the variance of the unconditional distribution in the triplet, $\varsigma^2_{doi,t}$, as the square of the coefficient on $\Phi^{-1}(\hat{F}_{doi,t}(\omega))$ from the corresponding regression. Since the relationship between theoretical and empirical quantiles is linear under log-normality, QQ regressions provide an estimate of $\varsigma^2_{doi,t}$ independently of truncation in the observed data.

We implement the QQ estimator on the two samples of triplets with 9+ and 27+ varieties, which provide us with enough degrees of freedom to estimate the parameters of the QQ regression in each triplet. Columns (5)-(6) of Table 4 report descriptive statistics on the R^2 of the QQ regressions. We find the latter to fit the data remarkably well: in the sample of triplets with 9+ varieties, the R^2 equals 0.95 for the average triplet and its standard deviation across triplets is a tiny 0.05. In the sample of triplets with 27+ varieties, the average R^2 is even higher (0.96) and the distribution of R^2 across triplets even tighter (s.d. 0.03).

The above findings corroborate the results of the KS and D'Agostino et al. (1990) tests. However, as is well-known, the fit of the log-normal distribution may be relatively worse on the right tail. We can use the results of the QQ regressions to also shed light on this issue. To this purpose, we study how the fit of these regressions changes as we progressively trim larger parts of the left tail of the sales data in each triplet. In addition, we analyze how the coefficient on the theoretical quantiles changes on the truncated samples, as this coefficient should not be sensitive to truncation under log-normality.

$\sqrt{n}D_n$ converges to the Kolmogorov distribution K, which is independent of F. The null hypothesis is rejected at significance level α when $\sqrt{n}D_n > K_\alpha$, where K_α is the critical value of the Kolmogorov distribution, i.e., the value such that $\Pr\left(K \le K_\alpha\right) = 1 - \alpha$.

[27] These results are based on the sample of 16893 triplets with at least 5 firm–products exported to the US, as the skewness and kurtosis statistics of the D'Agostino et al. (1990) test (see Eq. (13) and 19 in their paper) are not defined for smaller triplets. See Sager and Timoshenko (2019) for a previous application of the D'Agostino et al. (1990) tests to micro-level export data for Brazil.

[28] D'Agostino et al. (1990) tests are recommended for sample sizes $n \ge 9$.

Table 5
Descriptive statistics on heterogeneity and selection.

	Estimated (QQ) Variance of Log Sales				QQ/Observed Variance of Log Sales			
	Mean	Median	Std. Dev.	Change	Mean	Median	Std. Dev.	Change
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(a) Triplets with 9+ varieties	4.02	3.71	2.00	0.08	1.04	1.03	0.06	0.01
(b) Triplets with 27+ varieties	4.09	3.88	1.58	0.05	1.01	1.01	0.04	0.01

Notes. Columns (1)–(4) report descriptive statistics on the estimated variance of log sales, obtained by applying the QQ estimator on each country–industry–year triplet, using theoretical quantiles of sales implied by the log-normal distribution. Columns (5)–(8) report descriptive statistics on the ratio between the estimated and the observed variance of log sales. Reported statistics are the mean, median and standard deviation of each variable across all triplets, as well as the percentage change in the average value of each variable between 2002 and 2012. The sample consists of triplets with at least 9 (panel a) or 27 (panel b) firm–products exported to the US.

The results of these exercises are reported in the last four columns of Table 4, which show the median coefficient on $\Phi^{-1}(\hat{F}_{doi,t}(\omega))$ and the median R^2 obtained by running the QQ regressions on different samples. In particular, in column (7), we use the whole sample of varieties in each triplet, while in columns (8)-(10) we restrict to varieties in the top 75th, 50th and 25th percentile of the sales distribution in each triplet. The median R^2 ranges between 0.95 and 0.98, and the median coefficient always remains in a tight neighborhood of 2, suggesting that the log-normal distribution provides a good fit of the export data also on the right tail.

In Table 5, we present some statistics on the estimates of $\varsigma^2_{doi,t}$, the variance of the unconditional distribution of log sales in each triplet, both for the sample of triplets with 9+ varieties (panel a) and for the sample of triplets with 27+ varieties (panel b). Three facts stand out. First, heterogeneity is typically large, with $\varsigma^2_{doi,t}$ exceeding 4 in the average triplet and 3.7 in the median triplet. Second, heterogeneity has increased over time, with the average value of $\varsigma^2_{doi,t}$ across countries and industries rising by 5–8% between 2002 and 2012. Third, and more importantly, heterogeneity varies markedly across triplets. The standard deviation of $\varsigma^2_{doi,t}$ is roughly half its mean, ranging between 1.6 and 2 depending on the sample. More generally, the distribution of $\varsigma^2_{doi,t}$ across triplets spans a wide interval, ranging between 0.83 (1st percentile) and 10.83 (99th percentile) in the sample of triplets with 9+ varieties, and between 1.29 and 8.83 in the sample of triplets with 27+ varieties.[29]

Finally, we discuss the relationship between the unconditional variance, ς^2_{doi}, and the observed variance of log sales, $\mathbb{V}\left(\ln r_{doi}\right)$. Recall from (14) that the variance of the truncated distribution should be less than the variance of the unconditional distribution. Moreover, this ratio should be increasing in the cutoff, $\ln r^{\min}_{doi}$, and hence in selection. Table 5 reports descriptive statistics on $\varsigma^2_{doi,t}/\mathbb{V}\left(\ln r_{doi,t}\right)$ across triplets. The ratio exceeds 1 both in the average and in the median triplet, with a tiny standard deviation of 0.04–0.06 across triplets depending on the sample. Over time, $\varsigma^2_{doi,t}/\mathbb{V}\left(\ln r_{doi,t}\right)$ has remained largely stable, with the change in its average value not exceeding 1% on either sample.

5.3. Firm heterogeneity, trade and economic performance

We now discuss the implications of firm heterogeneity for trade. To this purpose, we run regressions of the log of exports to the US, $\ln R_{doi,t}$, on the log of the unconditional variance, $\ln \varsigma^2_{doi,t}$, across country–industry–year triplets. We control for two sets of fixed effects. First, we include industry×year fixed effects, so as to exploit cross-country variation within industry-years and neutralize differences in σ_i across industries. Second, we include country×year fixed effects, which absorb the influence of gravity variables such as distance, common language and common border, as long as these bilateral determinants of trade affect exports uniformly across industries.

The results are reported in Table 6, using the sample of triplets with 9+ varieties in column (1) and with 27+ varieties in column (5). The results show that, within a given industry and year, countries characterized by larger firm heterogeneity export significantly more to the US. The relationship between heterogeneity and trade is also remarkably tight.[30] In columns (2) and (6), we augment the specification by adding the log ratio between the unconditional variance and the observed variance in each triplet, $\ln\left(\varsigma^2_{doi,t}/\mathbb{V}\left(\ln r_{doi,t}\right)\right)$. Since the volume of exports is a negative function of the truncation point, a negative coefficient on $\ln\left(\varsigma^2_{doi,t}/\mathbb{V}\left(\ln r_{doi,t}\right)\right)$ would be consistent with selection, i.e., it would suggest that the cutoff for exporting to the US is lower for countries that sell more in the US in a given industry and year. In both samples, we find that larger exports to the US are indeed significantly associated with a lower value of $\ln\left(\varsigma^2_{doi,t}/\mathbb{V}\left(\ln r_{doi,t}\right)\right)$. We interpret this result as consistent with a "pecking order" across origin countries, whereby progressively less appealing firms export from countries selling more in a given market. Yet, the correlation between heterogeneity and trade is largely unaffected.

In columns (3) and (7), we include additional control variables. To account for differences in the average appeal of varieties, we add the log of the average quality-to-price ratio across all varieties in each triplet. We construct this variable using Eqs. (6) and (8), along with sales data from Piers and estimates of σ_i from Broda and Weinstein (2006); the industry-specific term A_{di} is

[29] The distribution of $\varsigma^2_{doi,t}$ across triplets is shown in Appendix Figure B2 for both samples.

[30] The scatterplots corresponding to these regressions are reported in Appendix Figure B3.

Table 6
Firm heterogeneity and exports.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Estimated (QQ) variance of log sales	1.655***	1.640***	1.042***	0.671***	2.072***	2.059***	1.309***	0.757***
	(0.032)	(0.031)	(0.036)	(0.046)	(0.058)	(0.055)	(0.062)	(0.080)
QQ/observed variance of log sales		−9.663***	−8.301***	−3.497***		−10.694***	−9.795***	−4.069***
		(0.349)	(0.320)	(0.370)		(0.667)	(0.659)	(0.595)
Average appeal			0.751***	0.690***			0.775***	0.719***
			(0.037)	(0.052)			(0.043)	(0.073)
Skill endowment × skill intensity			1.091***	−0.723			0.744***	−0.521
			(0.126)	(0.762)			(0.160)	(0.974)
Capital endowment × capital intensity			0.122***	−0.024			0.176***	0.015
			(0.016)	(0.071)			(0.021)	(0.076)
Distance × bulk weight			−0.170***				−0.151***	
			(0.032)				(0.044)	
Common language × bulk weight			0.002				0.048	
			(0.027)				(0.035)	
Common border × bulk weight			0.194**				0.263**	
			(0.078)				(0.117)	
US import tariff			0.632	−0.149			−0.043	−1.006
			(0.614)	(0.925)			(0.847)	(1.054)
Industry × year FE	yes	yes	yes	yes	yes	yes	yes	yes
Country × year FE	yes	yes	yes	yes	yes	yes	yes	yes
Country × industry FE	no	no	no	yes	no	no	no	yes
Obs.	12 437	12 437	11 646	11 646	6314	6314	6037	6037
R2	0.57	0.63	0.73	0.96	0.60	0.64	0.73	0.97

Notes. The dependent variable is the log of total exports to the US in each country–industry–year triplet. *Estimated (QQ) variance of log sales* is the log of the estimated variance of log sales, obtained by applying the QQ estimator on each triplet, using theoretical quantiles of sales implied by the log-normal distribution. *QQ/observed variance of log sales* is the log ratio between the estimated and the observed variance of log sales in each triplet. *Average appeal* is the log of the average quality-to-price ratio across all firm–products in each triplet. See Appendix A.2 for a detailed description of the other variables used in columns (3)–(4) and (7)–(8). All regressions are run across triplets. The sample consists of triplets with at least 9 (columns 1–4) or 27 (columns 5–8) firm–products exported to the US. The standard errors are corrected for heteroskedasticity. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

absorbed by the industry×year fixed effects. We also control for other determinants of trade. Because the specification includes industry× year and country×year fixed effects, all factors influencing exports to the US either at the industry level or at the country level are already accounted for. Hence, we add variables capturing determinants of trade that could operate both across countries and industries. In particular, to account for endowment-based comparative advantage, we include interactions between proxies for the skill and capital endowment of each country in a given year and proxies for the skill and capital intensity of each industry. As proxies for country–industry transportation costs, we include the log bulk weight of each industry interacted with three different variables: the log distance of each country from the US; and two dummies equal to 1 if a country shares the same language or a common border, respectively, with the US. To control for trade barriers, we also add the log of one plus the US import tariff in each triplet.[31]

While the coefficients on tariffs and the interaction between bulk weight and the common language dummy are imprecisely estimated, all other variables enter with the expected sign and their coefficients are statistically significant. Nevertheless, the inclusion of these controls implies just a moderate reduction in the coefficient on $\ln \varsigma^2_{doi,t}$, which remains highly significant. Hence, firm heterogeneity remains a strong correlate of a country's exports even after accounting for average appeal and other determinants of trade. Finally, in columns (4) and (8), we add country×industry fixed effects. These fixed effects subsume the proxies for country–industry transportation costs used in columns (3) and (7). Moreover, they imply that the remaining coefficients are identified only from variation over time within country–industry pairs. Heterogeneity remains a strong correlate of trade also in this very demanding specification.

Next, we revisit the correlations between firm heterogeneity and country characteristics. In particular, we are interested in assessing whether larger and richer countries have a higher dispersion of sales because of differences in the underlying distributions, because of a lower export cutoff, or both. To answer this question, we first regress $\ln \mathbb{V}\left(\ln r_{doi,t}\right)$, $\ln \varsigma^2_{doi,t}$ and $\ln\left(\varsigma^2_{doi,t}/\mathbb{V}\left(\ln r_{doi,t}\right)\right)$ on country and industry×year fixed effects, and then relate the country effects from these regressions to the log of countries' real per-capita GDP and population.

The results are reported in Table 7. In columns (1)-(3), we focus on the sample of triplets with 9+ varieties, whereas in columns (4)-(6) we use the sample of triplets with 27+ varieties. The estimates imply that heterogeneity is higher in larger and richer countries even when controlling for possible truncation in the data. Moreover, the coefficients from the regression for the unconditional variance (columns 2 and 5) are smaller than those from the regression for the observed variance (columns 1 and 4), consistent with

[31] See Appendix A.2 for a detailed description of these variables and their sources.

Table 7
Firm heterogeneity, selection and country characteristics.

	Observed Variance of Log Sales (1)	Estimated (QQ) Variance of Log Sales (2)	QQ/Observed Variance of Log Sales (3)	Observed Variance of Log Sales (4)	Estimated (QQ) Variance of Log Sales (5)	QQ/Observed Variance of Log Sales (6)
Real per-capita GDP	0.135***	0.130***	−0.005*	0.095***	0.093***	−0.001
	(0.026)	(0.027)	(0.003)	(0.023)	(0.022)	(0.001)
Population	0.054***	0.045**	−0.009***	0.034*	0.028	−0.006***
	(0.017)	(0.018)	(0.002)	(0.018)	(0.018)	(0.001)
Obs.	97	97	97	80	80	80
R2	0.32	0.29	0.18	0.20	0.19	0.25

Notes. The dependent variables are the country fixed effects obtained by regressing the variables indicated in the columns' headings (in logs) on country and industry–year fixed effects across triplets, using the sample of triplets with at least 9 (columns 1–3) or 27 (columns 4–6) firm–products exported to the US. Real per-capita GDP and population are simple averages of these variables between the years 2002 and 2012. All explanatory variables are in logs. The standard errors are corrected for heteroskedasticity. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

the hypothesis of selection. However, the difference between the two sets of estimates is small, suggesting that selection is unlikely to be the main determinant of the observed differences in sales dispersion across countries.

6. Implications for price indexes

We now show that the results obtained so far are important also for welfare. In particular, using again our accounting framework, we show that the price index of the basket of imported varieties is significantly lower from origins with higher dispersion. We instead discuss in Appendix D the effect of heterogeneity on welfare in the workhorse trade model with free entry and selection.[32]

Since log-normal distributions provide a good approximation of the data, we can derive a simple formula mapping few and easy-to-compute statistics about firms into prices. These formulas allow us to study the effect of heterogeneity on the price index. We start by rewriting the price index in destination d and industry i as:

$$P_{di} = \left(\sum_o P_{doi}^{1-\sigma_i} \right)^{\frac{1}{1-\sigma_i}}, \quad \text{with} \quad P_{doi} = \left[\sum_{\omega \in \Omega_{doi}} (\tilde{\gamma}_{doi})^{\sigma_i - 1} \right]^{\frac{1}{1-\sigma_i}}. \tag{16}$$

Note that P_{doi} is the price index of the basket of goods imported from country o.

Taking logs, we can rewrite:

$$\ln(1/P_{doi}) = \frac{1}{\sigma_i - 1} \ln \mathbb{E}\left[(\tilde{\gamma}_{doi})^{\sigma_i - 1} \right] + \frac{\ln N_{doi}}{\sigma_i - 1}, \tag{17}$$

where $\mathbb{E}\left[(\tilde{\gamma}_{doi})^{\sigma_i-1}\right] = \left[\sum_{\omega\in\Omega_{doi}} (\tilde{\gamma}_{doi})^{\sigma_i-1}\right]/N_{doi}$. Since $\tilde{\gamma}_d(\omega)^{\sigma_i-1} = r_d(\omega)/A_{di}$, we can express the inverse of the price index as a function of the parameters of the unconditional distribution of sales, which is log-normal $\left(\mu_{doi}, \varsigma_{doi}^2\right)$. Due to truncation, the conditional mean is:

$$\ln \mathbb{E}\left[(\tilde{\gamma}_{doi})^{\sigma_i - 1} \right] = \mu_{doi} + \frac{1}{2}\varsigma_{doi}^2 + \ln \frac{\Phi\left(\varsigma_{doi} - \theta_{doi}\right)}{1 - \Phi\left(\theta_{doi}\right)} - \ln A_{di}.$$

Substituting this into (17) yields:

$$\ln(1/P_{doi}) = \frac{1}{\sigma_i - 1} \left[\mu_{doi} + \frac{1}{2}\varsigma_{doi}^2 + \ln \frac{\Phi\left(\varsigma_{doi} - \theta_{doi}\right)}{1 - \Phi\left(\theta_{doi}\right)} - \ln A_{di} \right] + \frac{\ln N_{doi}}{\sigma_i - 1}. \tag{18}$$

This equation shows how the price index depends on the mean and variance of the unconditional distribution of log sales. However, we want to study the effect of heterogeneity holding constant the mean of the unconditional distribution of appeal: $\ln \mathbb{E}^u(\tilde{\gamma}_{doi}) = \widetilde{\mu}_{doi} + \widetilde{\varsigma}_{doi}^2/2$. Using this together with (11)–(12) to solve for μ_{doi} and substituting it into (18), we obtain:

$$\ln(1/P_{doi}) = \ln \mathbb{E}^u(\tilde{\gamma}_{doi}) + \frac{\sigma_i - 2}{2\left(\sigma_i - 1\right)^2}\varsigma_{doi}^2 + \frac{1}{\sigma_i - 1} \ln \frac{\Phi\left(\varsigma_{doi} - \theta_{doi}\right)}{1 - \Phi\left(\theta_{doi}\right)} + \frac{\ln N_{doi}}{\sigma_i - 1}. \tag{19}$$

Holding constant $\mathbb{E}^u(\tilde{\gamma}_{doi})$ and in the absence of truncation, the second term in this expression shows that heterogeneity, as measured by ς_{doi}^2, lowers price indexes when $\sigma_i > 2$, i.e., when sales are convex in appeal.[33] Moreover, the third term in (19) shows that heterogeneity has an additional effect due to the truncation point. Since truncation also contributes to making sales convex, it

[32] In this section, we quantify the effect of firm heterogeneity for the welfare-relevant price index derived from CES preferences. Our results do not refer to the bias in measuring price indexes due to taste shocks discussed in Redding and Weinstein (2020).

[33] See Epifani and Gancia (2011) for a related result.

reinforces the impact of heterogeneity on the price index. Intuitively, shifting probability mass to the tails of the unconditional distribution lowers the observed average price when actual prices are cut above a certain level.[34]

These formulas allow us to evaluate the role of heterogeneity, as captured by the variance of the log of appeal. We now use them to perform some simple quantifications. It is important to stress that, by keeping constant both mean appeal and the number of varieties, these exercises allow us to isolate the relevance of heterogeneity in an accounting sense. They do not correspond to counterfactual quantifications in which other variables are allowed to adjust. They should therefore be interpreted as partial equilibrium results or as outcomes of a decomposition that remains agnostic on how the various components of the price index are related to each other. Compared to the decompositions in Section 4, this exercise has the advantage of isolating the role of a structural parameter, rather than the observed dispersions.

We take as a benchmark the average industry in the average country. We therefore set $\sigma_i = 4$, which corresponds to the simple mean across industries in our sample and is well within the conventional values for the elasticity of substitution between varieties. As we will see later on, the results are not very sensitive to this parameter. For the remaining parameters, we focus on the sample of triplets with 27+ varieties, although results are very similar using triplets with 9+ varieties. We choose $\varsigma^2_{doi,t} = 4$, which corresponds to the mean of the QQ estimates of the variance of the unconditional sales distribution (see Table 5). For the truncation point, we set $\ln r^{\min}_{doi,t} = 8.7$, which corresponds to the simple average across country–industry–year triplets of the lowest decile of the sales distribution in each triplet. Finally, we obtain the unconditional mean, $\mu_{doi,t}$, so as to match the simple mean of the observed average sales, $\ln \mathbb{E}\left[r_{doi,t}\right] = 11.16$, according to (13), given $\ln r^{\min}_{doi,t} = 8.7$ and $\varsigma_{doi,t} = 2$. This yields $\mu_{doi,t} = 10.51$.

We now compare this benchmark industry i with another industry x that differs only in the dispersion of sales, $\varsigma^2_{dox,t}$, and compute the implied difference in their price indexes, holding everything else constant. In particular, we consider a second hypothetical industry with an unconditional variance that is one standard deviation higher than the mean. This corresponds to $\varsigma^2_{dox,t} = 5.22$ in the sample of triplets with 27+ varieties (the difference would be even larger, with $\varsigma^2_{dox,t} = 5.72$, in the sample of triplets with 9+ varieties). As a concrete example, these differences roughly correspond to the average industry in Germany (higher dispersion) versus China (lower dispersion). Then, their relative price index is:

$$\ln \frac{P_{doi}}{P_{dox}} = \frac{\sigma_i - 2}{2\left(\sigma_i - 1\right)^2}\left(\varsigma^2_{dox} - \varsigma^2_{doi}\right) + \frac{1}{\sigma_i - 1}\left[\ln \frac{\Phi\left(\varsigma_{dox} - \theta_{dox}\right)}{1 - \Phi\left(\theta_{dox}\right)} - \ln \frac{\Phi\left(\varsigma_{doi} - \theta_{doi}\right)}{1 - \Phi\left(\theta_{doi}\right)}\right]. \tag{20}$$

Using the above-mentioned values, the first term is equal to 0.135 and the second to 0.013. Together, they imply a price index 16% higher in the industry with lower dispersion. These are sizeable differences. In terms of sales, given $\sigma_i = 4$, they would imply a 56% higher value from the industry with higher dispersion. Once again, these results confirm that differences in heterogeneity are quantitatively important for explaining economic performance.

How robust are these results to alternative parametrizations? As (20) shows, the direct effect of differences in ς^2_{doi} (disregarding truncation) only depends on σ_i. Estimates of this parameter are often in the $[2.5, 5]$ range. It turns out that, within this interval, the coefficient multiplying ς^2_{doi} is not very sensitive to the exact value of σ_i. In particular, it is equal to 0.11 for $\sigma_i = 2.5$, reaches a maximum of 0.125 for $\sigma_i = 3$ and falls to 0.094 for $\sigma_i = 5$.[35] Regarding the estimated differences in $\varsigma^2_{doi,t}$, as already mentioned, variation is even higher in the sample of triplets with 9+ varieties. In this case, the direct effect from the first term of (20) would increase to 0.19, which is 40% higher than in the more conservative case considered above. The effect through truncation depends instead on more parameters. However, its impact on the price index is much smaller than the direct effect, about a tenth, and hence it is unlikely to significantly alter the overall quantitative implications under alternative parametrizations.

7. Conclusions

In this paper, we used transaction-level US import data to compare firms from virtually all countries in the world competing in a single destination market. By making distributional assumptions consistent with the data, we have identified new structural parameters that are useful in understanding trade and economic performance. In particular, we found that the shape parameter of the distribution of appeal varies systematically across countries and that this variation is quantitatively important to explain exports. We also found that richer countries export more per firm largely because they have more heterogeneous, and hence top, exporters. Finally, we showed that differences in the dispersion of appeal matter not just for exports, but also for price indexes and hence welfare. It is therefore important to take them into account, especially in quantitative models, and to understand their origins.

We conclude by discussing briefly some candidate explanations. First, it seems natural to conjecture that innovation be one driver of firm heterogeneity. For instance, richer markets may be more conducive to drastic innovation with more dispersed outcomes (e.g., Bonfiglioli et al., 2018, 2019) than imitation (Benhabib et al., 2021; König et al., 2016). It could also be that richer and thicker markets facilitate a stronger sorting between firms, suppliers and workers, which would amplify any pre-existing productivity differences (e.g., Bonfiglioli and Gancia, 2019; Sampson, 2014). Another hypothesis is that distortions in poorer countries may discourage large investments and prevent good firms from growing to their full potential. Since different models are likely to carry different policy prescriptions, identifying the mechanism through which firm heterogeneity is generated is an important direction for future research.

34 Note that $\ln(1/P_{doi})$ can also be expressed as a function of the dispersion in appeal, $\widetilde{\varsigma}^2_{doi}$. However, the advantage of (19) is that we have already obtained estimates of the parameter ς^2_{doi}.

35 The intuition for the non-monotonic effect is as follows. A high σ_i increases the effect of dispersion in attributes on the price index, but it also lowers the level of dispersion needed to match a given variance of sales.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.euroecorev.2024.104912.

Data availability

The authors do not have permission to share data.

EEREV-D-23-01191_Replication files.zip(Original data) (dropbox)

References

Autor, David, Dorn, David, Katz, Lawrence, Patterson, Christina, Van Reenen, John, 2020. The fall of the labor share and the rise of superstar firms. Q. J. Econ. 135, 645–709.

Bartelsman, Eric, Haltiwanger, John, Scarpetta, Stefano, 2013. Cross-country differences in productivity: The role of allocation and selection. Amer. Econ. Rev. 103, 305–334.

Bas, Maria, Mayer, Thierry, Thoenig, Mathias, 2017. From micro to macro: Demand, supply and heterogeneity in the trade elasticity. J. Int. Econ. 108, 1–19.

Benhabib, Jess, Perla, Jesse, Tonetti, Christopher, 2021. Reconciling models of diffusion and innovation: A theory of the productivity distribution and technology frontier. Econometrica 89, 2261–2301.

Bernard, Andrew, Jensen, Bradford, Redding, Stephen, Schott, Peter, 2009. The margins of US trade. Am. Econ. Rev. Pap. Proc. 99, 487–493.

Bernard, Andrew, Jensen, Bradford, Redding, Stephen, Schott, Peter, 2018. Global firms. J. Econ. Lit. 56, 565–619.

Bloom, Nicholas, Sadun, Raffaella, Van Reenen, John, 2016. Management as a technology? NBER Working Paper No. 22327.

Bonfiglioli, Alessandra, Crinò, Rosario, Gancia, Gino, 2018. Betting on exports: Trade and endogenous heterogeneity. Econom. J. 128, 612–651.

Bonfiglioli, Alessandra, Crinò, Rosario, Gancia, Gino, 2019. Trade, finance and endogenous heterogeneity. J. Eur. Econom. Assoc. 17, 79–130.

Bonfiglioli, Alessandra, Crinò, Rosario, Gancia, Gino, 2021. Concentration in international markets: Evidence from US imports. J. Monetary Econ. 121, 19–39.

Bonfiglioli, Alessandra, Gancia, Gino, 2019. Heterogeneity, selection and labor market disparities. Rev. Econ. Dyn. 31, 305–325.

Broda, Christian, Weinstein, David, 2006. Globalization and the gains from variety. Q. J. Econ. 121, 541–585.

Cabral, Luís, Mata, José, 2003. On the evolution of the firm size distribution: Facts and theory. Amer. Econ. Rev. 93, 1075–1090.

Chaney, Thomas, 2008. Distorted gravity: The intensive and extensive margins of international trade. Amer. Econ. Rev. 98, 1707–1721.

Costinot, Arnaud, Rodríguez-Clare, Andrés, 2014. Trade theory with numbers: Quantifying the consequences of globalization. In: Gopinath, Helpman, Rogoff (Eds.), Handbook of International Economics, Vol. 4. Elsevier, pp. 197–261.

Crinò, Rosario, Ogliari, Laura, 2017. Financial imperfections, product quality, and international trade. J. Int. Econ. 104, 63–84.

D'Agostino, Ralph, Belanger, Albert, D'Agostino, Jr., Ralph, 1990. A suggestion for using powerful and informative tests of normality. Amer. Statist. 44, 316–321.

di Giovanni, Julian, Levchenko, Andrei, Ranciere, Romain, 2011. Power laws in firm size and openness to trade: Measurement and implications. J. Int. Econ. 85, 42–52.

Eaton, Jonathan, Kortum, Samuel, Kramarz, Francis, 2011. An anatomy of international trade: Evidence from French firms. Econometrica 79, 1453–1498.

Epifani, Paolo, Gancia, Gino, 2011. Trade, markup heterogeneity and misallocations. J. Int. Econ. 83, 1–13.

Feenstra, Robert, 1994. New product varieties and the measurement of international prices. Amer. Econ. Rev. 84, 157–177.

Feenstra, Robert, Romalis, John, 2014. International prices and endogenous quality. Q. J. Econ. 129, 477–527.

Feenstra, Robert, Weinstein, David, 2017. Globalization, competition, and U.S. welfare. J. Polit. Econ. 125, 1041–1074.

Fernandes, Ana, Freund, Caroline, Pierola, Martha Denisse, 2016. Exporter behavior, country size and stage of development: Evidence from the exporter dynamics database. J. Dev. Econ. 119, 121–137.

Fernandes, Ana, Klenow, Peter, Meleshchuk, Sergii, Pierola, Martha Denisse, Rodriguez-Clare, Andres, 2023. The intensive margin in trade: How big and how important? Am. Econ. J.: Macroecon. 15, 320–354.

Flaaen, Aaron, Haberkorn, Flora, Lewis, Logan, Monken, Anderson, Pierce, Justin, Rhodes, Rosemary, Yi, Madeleine, 2023. Bill of lading data in international trade research with an application to the COVID-19 pandemic. Rev. Int. Econ. 31, 1146–1172.

Freund, Caroline, Pierola, Martha Denisse, 2015. Export superstars. Rev. Econ. Stat. 97, 1023–1032.

Gennaioli, Nicola, Porta, Rafael La, López-de Silanes, Florencio, Shleifer, Andrei, 2013. Human capital and regional development. Q. J. Econ. 128, 105–164.

Hallak, Juan Carlos, Schott, Peter, 2011. Estimating cross-country differences in product quality. Q. J. Econ. 126, 417–474.

Head, Keith, Mayer, Thierry, Thoenig, Mathias, 2014. Welfare and trade without Pareto. Am. Econ. Rev. Pap. Proc. 104, 310–316.

Hottman, Colin, Redding, Stephen, Weinstein, David, 2016. Quantifying the sources of firm heterogeneity. Q. J. Econ. 131, 1291–1364.

Hummels, David, 2001. Toward a geography of trade costs. Mimeo, Purdue University.

Hummels, David, Klenow, Peter, 2005. The variety and quality of a nation's exports. Amer. Econ. Rev. 95, 704–723.

Khandelwal, Amit, 2010. The long and short (of) quality ladders. Rev. Econ. Stud. 77, 1450–1476.

Kondo, Illenin O., Lewis, Logan T., Stella, Andrea, 2023. Heavy tailed but not Zipf: Firm and establishment size in the United States. J. Appl. Econometrics 38 (5), 767–785.

König, Michael, Lorenz, Jan, Zilibotti, Fabrizio, 2016. Innovation vs. Imitation and the evolution of productivity distributions. Theor. Econ. 11, 1053–1102.

Kratz, Marie, Resnick, Sidney, 1996. The QQ-estimator and heavy tails. Stoch. Models 12, 699–724.

Mayer, Thierry, Melitz, Marc, Ottaviano, Gianmarco, 2014. Market size, competition, and the product mix of exporters. Amer. Econ. Rev. 104, 495–536.

Mrázová, Monika, Neary, Peter, 2017. Not so demanding: Demand structure and firm behavior. Amer. Econ. Rev. 107, 3835–3874.

Mrázová, Monika, Neary, Peter, Parenti, Mathieu, 2021. Sales and markup dispersion: Theory and empirics. Econometrica 89, 1753–1788.

Poschke, Markus, 2018. The firm size distribution across countries and skill-biased change in entrepreneurial technology. Am. Econ. J.: Macroecon. 10, 1–41.

Redding, Stephen, Weinstein, David, 2020. Measuring aggregate price indexes with taste shocks: Theory and evidence for CES preferences. Q. J. Econ. 135, 503–560.

Redding, Stephen, Weinstein, David, 2024. Accounting for trade patterns. J. Int. Econ. (150).

Sager, Erick, Timoshenko, Olga, 2019. The double EMG distribution and trade elasticities. Can. J. Econ. 52, 1523–1557.

Sampson, Thomas, 2014. Selection into trade and wage inequality. Am. Econ. J. Microecon. 6, 157–202.

Schott, Peter, 2004. Across-product versus within-product specialization in international trade. Q. J. Econ. 119, 647–678.

UNCTAD, 2018. Review of Maritime Transport. United Nations Publications, New York.