DISCUSSION PAPER SERIES

DP12829
(v. 3)

FIRMS AND ECONOMIC PERFORMANCE: A VIEW FROM TRADE

Gino Gancia, Alessandra Bonfiglioli and Rosario Crinò

INTERNATIONAL MACROECONOMICS AND FINANCE

INTERNATIONAL TRADE AND REGIONAL ECONOMICS

MACROECONOMICS AND GROWTH

ISSN 0265-8003

FIRMS AND ECONOMIC PERFORMANCE: A VIEW FROM TRADE

Gino Gancia, Alessandra Bonfiglioli and Rosario Crinò

Discussion Paper DP12829
First Published 29 March 2018
This Revision 03 June 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

FIRMS AND ECONOMIC PERFORMANCE: A VIEW FROM TRADE

Abstract

This paper uses US import data to compare firms from multiple origins in a single destination. Guided by structural equations, we decompose countries' market shares into the contribution of the number of firm-products, their average attributes (quality and unit values) and reallocations among heterogeneous firms. We show that heterogeneity in attributes is important for explaining average exports and why these are higher from richer and larger countries. While we find evidence consistent with selection and misallocation, they do not seem the sole drivers of the results. In conclusion, reallocations among heterogeneous firms are important for trade, economic development and welfare.

JEL Classification: F12, F14

Keywords: US Imports, Firm Heterogeneity, international trade, prices, Quality, Variety, granularity

Gino Gancia - g.gancia@qmul.ac.uk
Queen Mary University of London and CEPR

Alessandra Bonfiglioli - a.bonfiglioli@qmul.ac.uk
Queen Mary University of London and CEPR

Rosario Crinò - rosario.crino@unicatt.it
Università Cattolica del Sacro Cuore and CEPR

Firms and Economic Performance: A View from Trade*

Alessandra Bonfiglioli† Rosario Crinò‡ Gino Gancia§

This draft: June 2020

Abstract

This paper uses US import data to compare firms from multiple origins in a single destination. Guided by structural equations, we decompose countries' market shares into the contribution of the number of firm-products, their average attributes (quality and unit values) and reallocations among heterogeneous firms. We show that heterogeneity in attributes is important for explaining average exports and why these are higher from richer and larger countries. While we find evidence consistent with selection and misallocation, they do not seem the sole drivers of the results. In conclusion, reallocations among heterogeneous firms are important for trade, economic development and welfare.

JEL Classification: F12, F14.

Keywords: US Imports, Firm Heterogeneity, Superstar Firms, International Trade, Prices, Quality, Misallocation.

---

*We thank 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 Barcelona GSE Summer Forum (2018), CAGE International Research Day (Warwick), CEMFI, DICE Workshop (Tuscany), ECB, ESADE, ESSEC Business School (Cergy-Pontoise), ESSIM (Oslo), IMT (Lucca), Italian Trade Study Group (Milan), LSE, OECD (Paris), the 1st QMUL-SUFE Workshop (London), 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, Università della Svizzera Italiana and WIEN 2019 (Vienna) for valuable comments. We acknowledge financial support from the Spanish Ministry of Economy and Competitiveness through the National Plan (ECO2014-59805-P), and the Leverhulme Trust through a 2019 Research Fellowship.

†Queen Mary University of London and CEPR, Mile End Road, London E1 4NS, UK. E-mail: a.bonfiglioli@qmul.ac.uk.

‡Università Cattolica del Sacro Cuore, Dept. of Economics and Finance, CEPR and CESifo. Via Necchi 5, 20123, Milan, ITALY. E-mail: rosario.crino@unicatt.it.

§Queen Mary University of London and CEPR. Mile End Road, London E1 4NS, UK. E-mail: g.gancia@qmul.ac.uk.

1 Introduction

Understanding differences in economic performance across countries has always been one of the great challenges in economics. Until recently, efforts to address this question were aimed at measuring aggregate productivity from national accounts.[1] The availability of firm-level data revolutionized the field by showing that productivity varies enormously even across firms within countries.[2] One of the most astonishing facts emerging from this literature is the increasingly dominant role of top firms (e.g., Autor et al., 2020). For instance, according to The Economist (17 September 2016), 10% of the world's public companies generate 80% of all profits. Large firms also dominate exports. For instance, 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). Yet, due to the lack of comprehensive and comparable data, to date there is still little systematic evidence on the role of firms in explaining aggregate performance.[3]

In this paper, we use detailed import data to compare firms from virtually all countries in the world selling in the US market. In doing so, we provide a comprehensive account of how the distribution of firm-level characteristics explain aggregate sales and trade flows. This allows us to document new patterns and tackle a number of questions that have remained largely unanswered: How important quantitatively are reallocations towards top firms for explaining sales across countries? Why do average exports per firm vary with country size and the level of development? And, how does the distribution of firm characteristics vary across countries?

Following recent methodological advances in trade theory, we show that data on unit values and volumes of exports to a single destination market, together with few and commonly made assumptions on demand, are sufficient to map the market share captured by each country into the attributes of its firms.[4] We apply this methodology to a unique transaction-level data on US seaborne imports in 2002 and 2012 containing information on unit values, volumes and the identity of exporting firms for 6-digit products from over 100 countries. As a preliminary step, we decompose the variation in countries' market shares of US imports 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 in a given country. This decomposition shows that each margin accounts on average for half of the overall variation in countries' market shares.

We then decompose average sales per firm-product into two parts: the average "appeal" of the

[1] See, for instance, Hall and Jones (1999), Caselli (2005), Gancia, Mueller and Zilibotti (2013).

[2] See, for instance, Hsieh and Klenow (2009), Restuccia and Rogerson (2008), Syverson (2011) and, more recently, Baqaee and Farhi (2020).

[3] Existing international comparisons of firm-level characteristics are confined to a handful of countries only, such as 20 in Gennaioli et al. (2013), 24 in Bartelsman, Haltiwanger and Scarpetta (2013), 34 in Bloom, Sadun and Van Reenen (2016) and 50 in Poschke (2018). All these studies are based on national data which cannot be used to perform the structural decompositions in this paper.

[4] In particular, see Hottman, Redding and Weinstein (2016) and Redding and Weinstein (2020).

firm-product and a "heterogeneity" term, capturing the effect of reallocations towards top firm-products. Intuitively, countries with more appealing firms can sell more and capture larger market shares. However, dispersion in appeal also affects the total value of sales because consumers can substitute low-appeal firm-products for high-appeal firm-products. In particular, 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 implies larger average sales. We find that reallocations among heterogeneous firms explains roughly half of the cross-country variation in average sales per firm-product in our data.

We extensively explore the robustness of this decomposition and study whether it is driven by specific groups of firms, industries or countries. One possibility could be that the importance of reallocations is disproportionately driven by "exceptional" firms, which are few in numbers but nevertheless account for a large fraction of total sales. This is not the case: we find that the relative contribution of dispersion is the same after removing exceptional firms, suggesting that these firms are more a manifestation of export performance than a cause of it. Similarly, the results could be driven by country-industry-year triplets with few firm-products, where all observations are influential. Yet, the main results are confirmed even when restricting the analysis to less granular triplets. The decomposition is also robust to various changes in the composition of the sample, such as excluding small or large market shares or specific groups of countries or industries.

Next, we use our novel decomposition of the intensive margin of countries' exports to explain why larger and richer countries sell more per firm. This is an empirical regularity that has been documented before (Fernandes, Freund and Pierola, 2016) but is still not well understood. Focusing on within-industry variation, we find that countries with a higher GDP per capita have higher average sales per firm because they have both a higher average appeal and a higher dispersion in appeal. On the other hand, the correlation between population size and average exports works mostly through a higher dispersion. In other words, more populous countries do not have better firms on average, but they have more top firms.

The rest of the paper turns to the interpretation of these results. Since our sample only includes firms that export to the US, one interesting question is whether differences in dispersion are likely to be driven by selection or rather reflect more general country characteristics. In particular, progressively less appealing firms may export from origins selling more to the US, and this may explain the correlation between sales and dispersion. We can test this hypothesis by imposing more restrictions on the distribution of firm characteristics. If sales are log-normal, Quantile-Quantile (QQ) regressions provide estimates of the shape parameter of the distribution that are independent of any truncation. We then look for evidence of selection by comparing the variances of the observed log sales with the QQ estimates. Interestingly, we find that log-normal distributions with parameters that vary across triplets provide a remarkably good fit of the data, and we do find evidence consistent with selection. Yet, we also derive an alternative decomposition that shows the contribution of dispersion to be only marginally reduced when controlling for truncation. We also find evidence consistent with a lower cutoff for exporting to the US in larger and richer countries. Yet, even after controlling for truncation, larger and richer countries still display a significantly higher heterogeneity in sales among exporters.

Many of these results confirm the predictions of theories of trade with heterogeneous firms that followed Melitz (2003). However, the effects of country size on average exports per firm and selection are not natural implications of these models. They might be consistent with Ricardian models, where larger countries are less specialized and their exporters more heterogeneous (i.e., Eaton and Kortum, 2002), or with models of innovation in which ideas, reflected in firms' attributes, are more heterogeneous in larger and richer countries (see, among others, Bonfiglioli, Crinò and Gancia 2018, 2019b, Benhabib, Perla and Tonetti, 2017, and König, Lorenz and Zilibotti, 2016). Yet, another plausible explanation is the presence of distortions such that firms with better attributes have inefficiently high prices and hence remain too small (e.g., Bento and Restuccia, 2017, and Edmond, Midrigan and Xu, 2015, 2019). While we cannot estimate productivity or markups with our data, we can use the information on unit values to decompose appeal into quality and quality-adjusted prices. We can then assess if high-quality firms from less developed and smaller countries have disproportionately higher quality-adjusted prices. This exercise suggests that this covariance can be quantitatively important, but it does not seem to be the main driver of the results.

Finally, we argue that firm heterogeneity is important not just to explain observed trade flows, but also for welfare. From an empirical standpoint, we implement an accounting exercise showing that the price index of the basket of imported varieties is significantly lower from origins with higher dispersion. From a theoretical perspective, we show how both welfare and the trade elasticity depend on the extent of heterogeneity in the workhorse model of trade with free entry and selection. This suggests that cross-country differences in the dispersion of firms' attributes can have an important effect on the level and distribution of the gains form trade and should therefore be taken into account in quantitative models (e.g., Costinot and Rodriguez-Clare, 2014).

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., 2019, Bernard et al. 2018, Fernandes, Freund and Pierola, 2016, Bernard et al. 2009, Chaney, 2008, Hummels and Klenow, 2005). This strand of the literature has shown that larger and richer countries have both more and bigger exporters. The reason why country size affects the intensive margin is still unclear. Our data and decomposition allow us to move from market shares to firm *attributes*, and to further study the role of reallocations in explaining the intensive margin. The finding that larger and richer countries export more per firm partly because their firms are more unequal is to our knowledge novel. Other papers have studied export patterns at the firm level by destination (e.g., Eaton, Kortum and Kramarz, 2011, Mayer, Melitz and Ottaviano, 2014). This strand of the literature has found evidence of selection in the form of a "pecking order" of destinations for exporters from a given country. We instead exploit variation across exporters serving the same destination market from virtually all possible origins. Our results on selection are consistent with a "pecking order" across origin countries, whereby progressively less appealing firms export from countries selling more in a given market.[5]

A closely related paper is Redding and Weinstein (2018), who use a similar framework for aggregating transaction-level US import data. We differ in several important ways. First, we ask a different question. Redding and Weinstein (2018) are interested in quantifying the contribution of prices, quality and variety for comparative advantage and price indexes. Instead, we decompose total sales by country so as to identify the firm-level determinants of economic performance. Second, we propose a different decomposition that fully separates the effect of average and dispersion in the *level* of attributes. Compared to our results, the log-linear decomposition in Redding and Weinstein (2018) overstates the contribution of heterogeneity. Third, one of our main contributions is to move beyond an accounting exercise by studying the role of selection, misallocation and country characteristics. Finally, we use different data. In Bonfiglioli, Crinò and Gancia (2019a), we instead use the same data to study how concentration has changed among firms exporting to the US.

The remainder of the paper is organized as follows. Section 2 introduces the theoretical framework that guides us through the decomposition of countries' market shares. Section 3 describes the firm-level data on US imports that we use in the empirical analysis. Section 4 reports the results from the main decomposition. In Section 5, we use our decomposition to study why richer and larger countries sell more per firm. Section 6 studies the role of selection. Section 7 discusses the interpretation of the results in light of models of trade with heterogeneous firms and theories of misallocation. Section 8 concludes.

2 Theoretical Framework

We now show how to map countries' market shares in a given industry into firm-level characteristics. The only restriction that we impose is a Constant Elasticity of Substitution (CES) demand system.[6] We assume that firms produce differentiated varieties and we identify each of 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".[7]

[5] Other papers have documented the importance of quality for explaining trade flows, but do not provide a full decomposition of all the firm-level margins (e.g., Crinò and Ogliari, 2017, Feenstra and Romalis, 2014, Hallak and Schott, 2011, and Khandelwal, 2010).

[6] CES preferences are a dominant paradigm in the literature. See Matsuyama and Ushchev (2017), Mrázová and Neary (2017) and Mrázová, Neary and Parenti (2017) for interesting discussions of more general demand systems.

[7] 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 any case, we will show in a robustness check that using firms, rather than firm-products, as the

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 $i \in \{1, ..., I\}$ 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 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}:

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

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 Market Shares

Using (4), the expenditure share of a single variety ω can be written as:

$$s_d(\omega) \equiv \frac{p_d(\omega) c_d(\omega)}{\sum\limits_{\omega \in \Omega_{di}} p_d(\omega) c_d(\omega)} = \frac{\left[\gamma_d(\omega)/p_d(\omega)\right]^{\sigma_i - 1}}{\sum\limits_{\omega \in \Omega_{di}} \left[\gamma_d(\omega)/p_d(\omega)\right]^{\sigma_i - 1}}. \tag{5}$$

basic unit of analysis has no bearing on the main results.

Market shares are increasing in the quality-to-price ratio, $\gamma_d(\omega)/p_d(\omega)$. More importantly, this equation illustrates that the distribution of quality-to-price ratios and the demand elasticity are sufficient statistics to compute any market share. In particular, defining for convenience $\tilde{\gamma}_d(\omega) \equiv \gamma_d(\omega)/p_d(\omega)$, the market share captured by all varieties sold from a single country of origin o in industry i, denoted by S_{doi}, is:

$$S_{doi} = \frac{\sum\limits_{\omega \in \Omega_{doi}} [\tilde{\gamma}_d(\omega)]^{\sigma_i - 1}}{\sum\limits_{\omega \in \Omega_{di}} [\tilde{\gamma}_d(\omega)]^{\sigma_i - 1}}, \tag{6}$$

where Ω_{doi} is the set of varieties sold in market d from origin o in industry i. Starting from this equation, we are interested in understanding what makes a country capture larger market shares.

To this end, note first that the scale of $\tilde{\gamma}$ is irrelevant, because what matters is its ratio to the sector aggregate. For this reason, and without loss of generality, we can normalize $\sum\limits_{\omega \in \Omega_{di}} [\tilde{\gamma}_d(\omega)]^{\sigma_i - 1} = 1$. Then, adding and subtracting the arithmetic mean of $\tilde{\gamma}_d(\omega)$ from a single origin o, $\mathbb{E}[\tilde{\gamma}_{doi}]$, to the power of $\sigma_i - 1$ yields:

$$S_{doi} = \sum_{\omega \in \Omega_{doi}} \left\{ \tilde{\gamma}_d(\omega)^{\sigma_i - 1} + \mathbb{E}[\tilde{\gamma}_{doi}]^{\sigma_i - 1} - \mathbb{E}[\tilde{\gamma}_{doi}]^{\sigma_i - 1} \right\}. \tag{7}$$

This expression allows us to decompose countries' market shares as follows. Consider an industry i. First, defining N_{doi} (N_{di}) as the number of varieties sold from o (from all origins) and $\bar{s}_{doi}$ ($\bar{s}_{di}$) as their average market share, we can write:

$$S_{doi} = \frac{N_{doi} \cdot \bar{s}_{doi}}{N_{di} \cdot \bar{s}_{di}}. \tag{8}$$

Second, the average market share per variety from country o is:

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

while $\bar{s}_{di} = 1/N_{di}$.

Equation (8) decomposes countries' market shares into the contribution of the number of varieties (extensive margin) versus average sales of each variety (intensive margin). More interestingly, equation (9) shows that average sales per variety can be further decomposed into two terms. The first term captures the average quality-to-price ratio of varieties from a given country. The second term captures the importance of heterogeneity in quality-to-price ratios from that origin. Clearly, equation (9) shows that the heterogeneity term is zero if all the quality-to-price ratios from a given country are identical. But what is the sign of this term if quality-to-price ratios do vary across varieties? The answer to this question depends on the value of σ_i, because the latter captures the strength of reallocations towards better firms.

To see how, note from equation (5) that sales are a convex function of the quality-to-price ratio 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 the quality-to-price ratio, so that only its average, and not its distribution, matters. Finally, when $\sigma_i < 2$, sales are a concave function of the quality-to-price ratio, so that more heterogeneity has a negative contribution to the overall market share. 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 attractive varieties increases average sales when holding constant the mean quality-to-price ratio.[8]

Interestingly, we can also rewrite (9) in terms of the observed individual market shares and σ_i. In particular, substituting $\tilde{\gamma}_d(\omega) = s_d(\omega)^{1/(\sigma_i - 1)}$ and $\mathbb{E}[\tilde{\gamma}_{doi}] = \mathbb{E}\left[s_{doi}^{1/(\sigma_i-1)}\right]$ into (9) yields:

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

where again $\mathbb{E}\left[s_{doi}^{1/(\sigma_i-1)}\right]^{\sigma_i-1}$ is the market share captured by a firm with an attribute equal to the simple average of all firms from o. Equations (8) and (10) can be used to decompose the market share of any country o relative to any other country j or any other group of countries, such as the set of all exporters to destination d. To that end, all is needed is estimates of σ_i and the observed firm-level market shares from any group of countries. Note also that while documenting the properties of the distribution of market shares, such as the variance or concentration at the top, is an interesting exercise, it is not sufficient to assess the importance of heterogeneity in attributes, 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. Finally, although the second term on the right-hand side of (10) captures the effect of reallocations between heterogeneous firms, in what follows we will often refer to it simply as a measure of "heterogeneity".

3 The Data

To perform the empirical analysis, we need data on the sales of individual products in a single destination market by firms of different origin countries. We obtain this information using transaction-level data on US imports from Piers, a database administered by IHS Markit. Piers contains the complete detail of the bill of lading of any shipment imported into the US by sea. IHS markit collects the bills of lading filed with 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

[8] Note that while $\sigma_i > 1$ is sufficient for demand to increase more than proportionally with $\tilde{\gamma}$, due to the negative effect of quantity on prices, the elasticity of revenues is $\sigma_i - 1$. This is why an elasticity greater than two is needed for revenues to increase more than proportionally with $\tilde{\gamma}$.

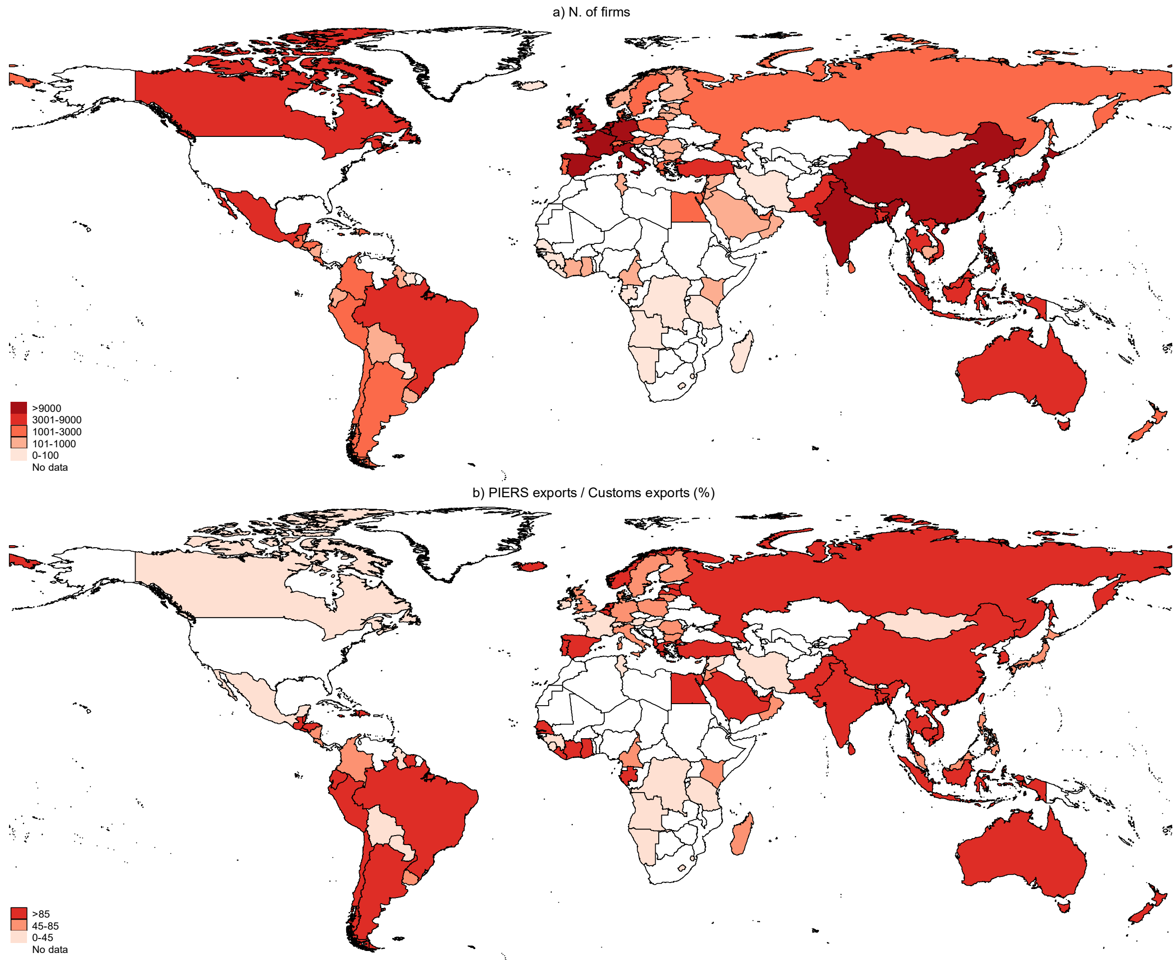

Source: Piers (IHS Markit), US seaborne import data for 2002 and 2012. Darker colors indicate a higher number of manufacturing firms exporting to the US (map a) or a higher ratio between the value of total manufacturing exports to the US obtained from Piers and the same value obtained from US Customs data (map b). All figures are averages between 2002 and 2012.

Figure 1: Data Coverage

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 (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.[9] Maritime trade represents the bulk of US trade in manufacturing, with seaborne imports accounting for 84% of total US imports in 2012.

Unlike other transaction-level data sets, Piers is not restricted and can thus be accessed by anyone at a fee. The fact that all firms in Piers use the same export mode (by sea) favors comparability. Importantly, Piers contains the full name of each firm. This unique feature allows researchers

[9] In the case of firms with multiple shipments (bills of lading) of the same product in a 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.

using Piers to precisely identify firms, reducing the risk of over-counting them. We assign varieties to industries by mapping each HS6 product exported by a firm to a 4-digit SIC industry, using a correspondence table developed by the World Integrated Trade Solutions. The 4-digit level of industry aggregation strikes a balance between number and comparability of products. On the one hand, we need each industry to encompass sufficiently comparable products, which would call for the use of highly disaggregated industries. On the other hand, the approach we use for estimating the elasticity of substitution in each industry (details below) rests on the set of common varieties imported in both years; identification therefore requires industries to be broad enough to encompass a sufficiently large number of varieties. In Section 6.2, we introduce an alternative decomposition that does not require an estimate of the elasticity of substitution. We thus perform an extensive sensitivity analysis using alternative levels of industry aggregation. As shown in Appendix Table A1, our results are robust to using industries that are either more or less aggregated than our baseline definition. Our 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.[10]

Figure 1 provides a visual representation of the geographic coverage of our sample. The first map focuses on the number of firms exporting to the US. The map shows that this number is particularly high in neighboring countries (Canada and Mexico), in large Latin American economies such as Brazil, in Europe and in South-East Asia (especially China). The second map describes the coverage of Piers in terms of export value. Darker colors indicate a better coverage, as measured by the ratio between the value of total exports computed from Piers and the same value computed from US Customs data at the product level (Feenstra, Romalis and Schott, 2002). The map shows that the coverage of Piers is remarkably good also across individual countries, with Piers covering more than 80% of total exports to the US for the average economy.[11] The coverage of Piers is

[10] We use a string matching algorithm to match and consolidate firms appearing in the data more than once with similar names. The algorithm first homogenizes standard expressions (e.g., it converts the extensions "Lim." and "LTD" in "Limited") and then computes the Levenshtein edit distance between all pairwise combinations of firm names sharing the same first character. The distance is normalized by the length of the longest string and a match is formed if the normalized edit distance is below a 5% threshold. To mitigate the risk of including transactions contaminated by reporting mistakes, we exclude firms whose total exports across all products to the US in a given industry and year are below $1,000 (4% of all firms) and firms with extreme unit values for their products, defined as unit values falling above the top or below the bottom 0.01% of the distribution in a given year. We focus on country-industry-year triplets with at least two varieties exported to the US, as variances are not defined for triplets with a single variety.

[11] We have also compared the number of firms exporting to the US in our sample with the corresponding number in the World Bank Exporter Dynamics Database (EDD), which uses information for the universe of export transactions obtained from each country's government custom agency. In 2012, out of the 48 countries covered by the version of EDD dedicated to firms exporting more than $1,000, 34 countries were also part of our sample. We found that our sample accounts for 63% of the total number of firms selling in the US for the average or median country, confirming that the majority of firms in the typical country export to the US by sea. Knowing the full name of firms allows us to significantly reduce the risk of over-counting them. For instance, Kamal, Krizan and Monarch (2015) perform the same comparison for the restricted-access US Customs and Border Protection (CBP) database, which covers the universe of US trade transactions (including non-seaborne trade) but does not contain the complete name of firms.

Table 1: Descriptive Statistics on Sample Coverage and Composition

	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)	60347	2360	536000
Average exports per variety ($1000)	1273	230	11058

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 this variable across all countries and years. The variables in panel b) are computed for each country-industry-year triplet. Reported statistics are the mean, median and standard deviation of these variables across all triplets.

unsurprisingly less extensive for Canada and Mexico, two countries for which maritime trade is not the main mode of export to the US. Nevertheless, these countries have a large number of firms exporting to the US, as shown in the first map. Because our decompositions are valid for any subset of firms and sales, we keep Canada and Mexico in our main baseline sample. In Section 4.2, however, we find that excluding all countries for which the coverage of Piers is less extensive (i.e., the first group of countries in the bottom map) has no bearing on our results.

Table 1 provides further details on sample coverage and composition. Panel a) confirms the high coverage of Piers, showing that for the average (median) country in our sample, Piers accounts for 83% (77%) of total exports to the US. These numbers are similar to the figures reported by Feenstra and Weinstein (2017) for an earlier and more limited vintage of the Piers database. Panel b) provides details on sample composition. All variables in this panel are computed separately for each country-industry-year triplet, and reported statistics are calculated across all triplets in the data set. 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.

4 Decomposing US Imports

4.1 Preliminaries: Sales, Quality and Prices

To implement the decompositions illustrated in Section 2.2, we need to estimate the elasticity of substitution between varieties in each industry, σ_i. We can then use these estimates to back out the quality of each variety, $\gamma_d(\omega)$. These tasks can be accomplished using data on prices and market

The authors show that,for most countries, the CBP overshoots the number of foreign firms exporting to the US in the EDD, with an average over-counting rate of 25%. Some of the firms in Piers could in principle be trading companies, but we find no such company among the top-10 exporting firms in any 2-digit industry, suggesting that the majority of firms in our sample are actual exporters. Finally, we have compared the information on unit values contained in Piers with the unit values for maritime trade obtained from US Customs data at the product level (Feenstra, Romalis and Schott, 2002). Regressing the unit values in the custom data on the unit values in Piers, across exporting countries and 6-digit products in 2002 and 2012, yields a coefficient of 0.836 (s.e. 0.003) and an R^2 of 0.58.

shares, together with the structural equations of the model. We now discuss how.

The first step consists of estimating the elasticities of substitution. To obtain these parameters, we use the reverse-weighting (RW) estimator introduced by Redding and Weinstein (2017). As detailed in Appendix A, this estimator looks for the value of σ_i that minimizes the sum of squared deviations of the forward and backward differences of the price index—which measure changes in the cost of imported varieties using initial period (2002) and final period (2012) expenditure shares, respectively, as weights—from a money-metric price index—which depends solely on prices and expenditure shares and is independent of demand parameters. The identifying assumption underlying the RW estimator is that changes in relative demand cancel out across varieties, consistent with the theoretical framework outlined in Section 2. Under this assumption, σ_i is identified out of time variation in prices and market shares for varieties that are imported in both years. In our sample, the number of common varieties allows identifying σ_i in most industries (259); all estimates satisfy the theoretical restriction that $\sigma_i > 1$.[12]

For the median industry in our sample, σ_i equals 3.54, a value that is in between the estimates of the elasticity of substitution across firms and across firms' products obtained by Redding and Weinstein (2018) using restricted-access transaction-level data on US imports from the US Customs. Regarding the distribution of σ_i across industries, the 10th, 25th, 75th and 90th percentiles of our estimates are equal to 1.92, 2.50, 5.06 and 10.34, respectively. Importantly, our estimates of σ_i are thus larger than 2 for the vast majority of industries. As explained in the previous section, this implies that greater heterogeneity in attributes should translate into higher average market shares.

With the estimates of σ_i in hand, we can use information on prices and market shares to back out the quality of each variety relative to the industry average. In particular, it holds that:

$$\gamma_{d,t}(\omega) = p_{d,t}(\omega) \cdot s_{d,t}(\omega)^{1/(\sigma_i - 1)}, \tag{11}$$

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. Thus, similar to Khandelwal, Schott, and Wei (2013), the quality of each variety is inferred from variation in market shares conditional on prices: higher-quality varieties exhibit higher market shares in an industry-year for given price.

Having estimated elasticities of substitution at the industry level and computed firm attributes that rationalize observed sales, we now study the role of firms in shaping trade flows. We start by presenting some new stylized facts about how sales and firm attributes vary across industries and countries. In Table 2, we report summary statistics on a number of important moments. The first two columns show the mean and standard deviation of each variable across all country-industry

[12] As a robustness check, we will also perform our decomposition using an alternative set of estimates of σ_i, obtained by exploiting cross-industry variation in sales dispersion, as explained in Appendix A. With two time observations per variety, σ_i cannot be estimated using the approach of Feenstra (1994), which requires at least three years of data.

Table 2: Descriptive Statistics on Key Moments

	Mean (2012)	Std. Dev. (2012)	Change (02-12)
Var. log sales	3.69	3.11	0.10
Var. log prices	0.41	0.68	-0.01
Var. log quality	2.92	9.87	0.11
Var. log quality-to-price ratio	2.24	9.42	0.18
Cov. log quality-log prices	0.52	1.00	-0.07

Notes. All variables are computed separately for each country-industry-year triplet. The first two columns report the mean and the standard deviation of each variable across all countries and industries in the year 2012. The third column reports the percentage change in the average value of each variable between 2002 and 2012.

pairs in 2012; the third column shows the change in the average value of each variable between 2002 and 2012. Sales dispersion is high, varies markedly across countries and industries, and has increased by 10% over the sample period.[13] Given that we know the identity of firms, with our data we can also compute the change in sales dispersion driven by reallocations among firms active in both years. We find that, in the subsample of continuing varieties, sales dispersion has increased by 29%.[14] In the rest of the sample, sales dispersion has increased by approximately 8%. Note also that our measure of sales dispersion coincides with the dispersion of market shares, given that the latter are defined relative to the total sales in a given industry.

Quality dispersion shows similar patterns, and the variance of log quality is close to the variance of log sales on average. Conversely, price dispersion is relatively small, exhibits a low cross-sectional variation, and has remained stable over time. Consistent with these patterns, the table also documents a substantial dispersion in quality-to-price ratios, as well as a tendency for it to increase over time. Finally, the covariance between log quality and log prices is positive, suggesting that higher-quality products are more costly to produce and more expensive, but it has become weaker over time.

To have a first sense of how firm heterogeneity varies across countries, Figure 2 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.[15] 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

[13] 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, Crinò and Gancia, 2018, 2019b).

[14] Continuing varieties account for 28% of total exports to the US in the average country-industry pair in 2012. To save space, statistics on the subsample of continuing firms are not tabulated in Table 2.

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

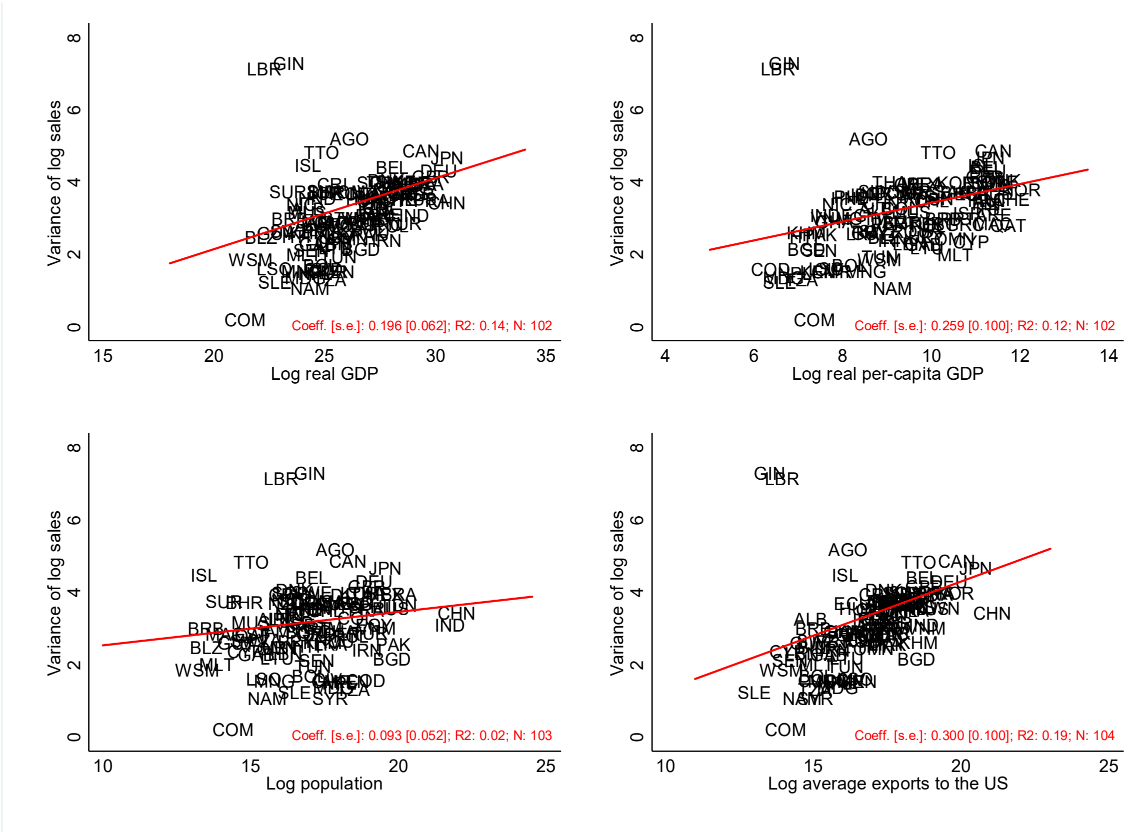

Notes. Variance of log sales and average exports to the US are computed separately for each country-industry-year triplet and then averaged across all industries and years for each country. The other country characteristics are simple averages between the years 2002 and 2012.

Figure 2: Dispersion of Market Shares and Country Characteristics

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 a strong positive correlation between sales dispersion and both country characteristics: sales are significantly more dispersed both in richer and in 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.

Having described the main features of the data, we next proceed with an exact decomposition of firm sales in the US market, which allows us to quantify the importance of firm attributes, and especially their dispersion, for economic success. We will then explore more in depth the origin and consequences of firm heterogeneity.

4.2 Decomposing Sales: Firms, Average Attributes and Heterogeneity

We now implement the decompositions presented in Section 2.2. We start by decomposing countries' market shares into the contributions of the extensive and intensive margin. To this purpose, we take logs of (8) and run separate regressions of $(\ln N_{doi,t} - \ln N_{di,t})$ and $(\ln \bar{s}_{doi,t} - \ln \bar{s}_{di,t})$ on $\ln S_{doi,t}$ across all available triplets. In all regressions, we control for industry-year fixed effects, so that our decomposition focuses on variation in market shares across countries within each industry and year.[16] The properties of OLS imply that the coefficients on $\ln S_{doi,t}$ obtained from these regressions add up to one, and thus provide the percentage contribution of each margin to explaining variation in countries' market shares. We similarly decompose the intensive margin into the contribution of average attributes and heterogeneity in attributes, by regressing each term on the right-hand side of (10) on $\bar{s}_{doi,t}$ and industry-year fixed effects.

The results of these decompositions are reported in Table 3. Panel a) shows the decompositions performed on the full sample. The results in columns (1) and (2) refer to the contributions of the extensive and intensive margin. Each margin explains roughly half of the variation in countries' market shares. Hence, countries that capture larger market shares in 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 attributes and heterogeneity in attributes. The results show that reallocations between heterogeneous firms contribute at least as much as average attributes to explaining variation in average market share per variety. Hence, firm heterogeneity is an important factor for understanding countries' export performance.

These results are consistent across a large number of robustness checks, which are presented in the remaining panels of Table 3. In panel b), we exclude countries for which Piers covers less than 45% of total exports to the US (i.e., the first group of countries in Figure 1b). The coefficients are essentially unchanged, suggesting that the decompositions are not influenced by countries for which the coverage of Piers is less extensive. 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 each industry-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%.[17] Also in this case, the results are virtually unchanged, suggesting that the decompositions are not driven by industries in which US imports predominantly consist of intermediate inputs

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

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

Table 3: Decomposition of Countries' Market Shares

	First Step		Second Step	
	Extensive Margin	Intensive Margin	Average	Dispersion
	(1)	(2)	(3)	(4)
a) Baseline	0.471***	0.529***	0.492***	0.508***
	[0.003]	[0.003]	[0.068]	[0.068]
b) No small countries	0.481***	0.519***	0.448***	0.552***
	[0.003]	[0.003]	[0.067]	[0.067]
c) No countries with small market shares	0.516***	0.484***	0.489***	0.511***
	[0.007]	[0.007]	[0.069]	[0.069]
d) No countries with large market shares	0.466***	0.534***	0.436***	0.564***
	[0.003]	[0.003]	[0.052]	[0.052]
e) No industries with high shares of imported inputs	0.466***	0.534***	0.497***	0.503***
	[0.004]	[0.004]	[0.081]	[0.081]
f) No top-1 firm in each triplet	0.510***	0.490***	0.509***	0.491***
	[0.003]	[0.003]	[0.122]	[0.122]
g) No firms with sales 2SD above triplet average	0.503***	0.497***	0.570***	0.430***
	[0.004]	[0.004]	[0.063]	[0.063]
h) Triplets with 9+ varieties	0.490***	0.510***	0.329***	0.671***
	[0.006]	[0.006]	[0.026]	[0.026]
i) Triplets with 27+ varieties	0.499***	0.501***	0.506***	0.494***
	[0.010]	[0.010]	[0.061]	[0.061]
j) Alternative elasticity of substitution	0.466***	0.534***	0.533***	0.467***
	[0.003]	[0.003]	[0.044]	[0.044]

Notes. Columns (1) and (2) perform the decomposition in eq. (8). 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 countries' market shares, 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 market share per variety, controlling for industry-year fixed effects. Panel a) uses the whole sample of triplets (18800 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 1 (16313 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 (11606 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 (18449 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% (15231 observations). Panel f) uses the sample that excludes the top-1 firm in each triplet (15872 observations). Panel g) 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 (18800 observations). Panel h) uses the sample of triplets with at least 9 firm-products exported to the US (9451 observations). Panel i) uses the sample of triplets with at least 27 firm-products exported to the US (4727 observations). Panel j) uses estimates of the elasticity of substitution obtained as in Appendix A.2 (24754 observations). The standard errors are corrected for heteroskedasticity. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

reflecting firms' participation in global value chains.

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 attributes 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.[18] But are these firms really "exceptional", in the sense that they significantly affect the quantitative role of dispersion in explaining market shares? 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 f), we exclude

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

the top-1 firm from each triplet, defined as the firm with the highest sales across products. In panel g), we instead remove all firms whose total exports (across all products) to the US are at least two standard deviations above the average exports for their triplet. The decomposition of the extensive and intensive margin is essentially unchanged, and heterogeneity in attributes now explains a slightly smaller share of the variation in average market shares per variety compared to panel a). While this was expected, the reduction in the size of the coefficients is quite small (the contribution of heterogeneity drops at most from 50 to 43%), which suggests that the importance of heterogeneity in attributes for explaining the intensive margin of countries' exports does not merely reflect the presence of exceptional firms.

Aside from superstar firms, any observation is likely to be influential in small samples. Accordingly, one concern could be that the results of our decompositions are driven by triplets with very few varieties. To study this possibility, we now consider 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). The results, reported in panel h) for the sample of triplets with 9+ varieties and in panel i) for the sample of triplets with 27+ varieties, show no change in the first decomposition. At the same time, the results suggest that the contribution of dispersion may exceed 50% once small triplets are excluded.

The importance of heterogeneity in attributes for the decomposition of average market shares per variety also depend on the elasticity of substitution, σ_i, as suggested by (10). Thus, for robustness, we now repeat the decompositions using an alternative set of estimates of σ_i to compute the two terms on the right-hand side of (10). As explained in Appendix A, these estimates are obtained by exploiting cross-industry variation in sales dispersion, and are available for all industries. While they cannot be given a structural interpretation, they absorb industry averages and thus isolate the cross-country variation in appeal that we are interested in. The results are reported in panel j). The coefficients in columns (1) and (2) are slightly different from their counterparts in panel a) due to the use of a larger sample; 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 remarkably similar to 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 attributes and dispersion of attributes explain half of the variation in average market shares per variety.

We now pause to briefly discuss the relationship between these results and the existing literature. That the extensive margin explains about half of the variation in trade flows is consistent with the findings in Hummels and Klenow (2005) using product-level data, Fernandes, Freund and Pierola (2016) and Fernandes et al. (2019) using firm-level data for up to 50 countries, and Redding and Weinstein (2018) using US import data from the US Customs. The contribution of firm heterogeneity in affecting the volume of trade has received less attention. Redding and Weinstein (2018), 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 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%.

5 Average Exports across Countries

Larger and richer countries have been shown to have higher exports per firm (see, for instance, Fernandes, Freund and Pierola, 2016). Our decomposition allows us to probe deeper into the mechanisms underlying these relationships. In particular, we now 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 decompose the correlation of population and income per capita with average market shares per variety into their two components: average attributes and heterogeneity in attributes across a country's firms. 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. Given that OLS is a linear operator, the coefficients obtained from the regressions for average attributes and heterogeneity in attributes add up to the coefficients obtained from the regression for average market shares. Hence, the two coefficients provide an additive decomposition of the margins along which population and income per capita correlate with the intensive margin of countries' exports. To account for the role of distance in explaining trade flows, we control for the distance of each country to the US in all specifications.[19]

The results are reported in Table 4. Column (1) shows that population and income per capita are both strongly correlated with average market shares. The coefficient on the log of real per-capita GDP implies that a 1 s.d. increase in this variable is associated with a rise by 52% of a s.d. in average market shares. The effect of a commensurate increase in log population is a rise in average market shares equal to 53% of a s.d.. For comparison, average market shares would rise by only 24% of a s.d. following a 1 s.d. decrease in log distance. Hence, our data confirm that population and income per capita are important correlates of the intensive margin of countries' exports.

Columns (2) and (3) decompose the relationship between average market shares and the two country characteristics into the contributions of average attributes and heterogeneity in attributes. In particular, column (2) reports the coefficients obtained from the regression for average attributes, while column (3) shows the estimated coefficients from the regression for heterogeneity in attributes.

[19] Distance is the population-weighted number of kilometers between each exporting country and the US constructed by CEPII. Population and per-capita GDP are averaged between the years 2002 and 2012.

Table 4: Average Exports and Country Characteristics

	Country-Level Averages			Country Fixed Effects		
	Market share	Average	Dispersion	Market share	Average	Dispersion
	(1)	(2)	(3)	(4)	(5)	(6)
Real per-capita GDP	0.030***	0.019***	0.011***	0.011**	0.006*	0.005***
	[0.004]	[0.003]	[0.002]	[0.004]	[0.003]	[0.002]
Population	0.027***	0.014***	0.012***	0.007	0.002	0.005**
	[0.004]	[0.002]	[0.002]	[0.004]	[0.003]	[0.002]
Distance	-0.044**	-0.023*	-0.021***	-0.037*	-0.016	-0.021***
	[0.020]	[0.014]	[0.007]	[0.022]	[0.015]	[0.008]
Obs.	101	101	101	101	101	101
R2	0.52	0.45	0.52	0.16	0.08	0.23

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. Distance is the population-weighted number of kilometers between each country and the US. All explanatory variables are in logs and all dependent variables are multiplied by 100. The standard errors are corrected for heteroskedasticity. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

Population and income per capita are positively correlated with each margin. For both country characteristics, coefficients are similar across the two regressions, implying that average attributes and heterogeneity in attributes have a quantitatively similar importance for explaining the correlations of the intensive margin of exports with population and income per capita.

The above results are consistent with two interpretations. On the one hand, larger and richer countries may host better and more heterogeneous firms within each industry. On the other hand, larger and richer countries may be specialized in industries in which firms are better and more heterogeneous. In this paper, we are especially interested in the first explanation: in order to know whether firms really differ across countries, we want to hold constant the industry composition of exports. Yet, the second explanation is not completely ruled out by taking country-level averages of the three terms in (10), unless all countries exported to the US in all sectors.

To 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 properties of OLS ensure that the country effects obtained from the regressions for average attributes and heterogeneity in attributes add up to those obtained from the regression for average market shares. This allows us to additively decompose the contributions of the two margins while cleaning the results from compositional effects.[20]

The results are reported in columns (4)-(6) of Table 4. The estimated coefficients confirm that larger and richer countries have higher average market shares, the coefficient on log population now being only marginally insignificant (*p-value* 0.11). The correlation between income per capita and the intensive margin of exports continues to be driven by both margins with similar contributions. Conversely, the correlation between population and the intensive margin of exports now mostly works through a higher dispersion. Comparing this result with the estimates in columns (1)-(3) suggests that more populous countries tend to specialize in industries where firms have better attributes on average. Once compositional effects are neutralized, however, larger countries no longer appear to have better firms. On the contrary, these countries systematically host firms with more heterogeneous attributes within industries. The higher incidence of top firms within sectors in larger countries almost entirely explains why population is correlated with the intensive margin of exports.

Why larger countries have higher exports per firm is still an open question. Our decomposition helps us shed some more light on it. On the one hand, the results imply that more populous countries tend to specialize in sectors dominated by large firms, possibly because they can take advantage of scale economies. On the other hand, within sectors, larger countries have more heterogeneous, rather than better, exporters, suggesting that larger markets may be more conducive to reallocations in favor of top firms.

6 Firm Heterogeneity and Selection

One reason for the importance of firm heterogeneity in explaining sales documented in Section 4.2 may be selection into exporting. In countries where the costs of exporting to the US is lower, firms with relatively worse attributes will find it profitable to export. This will induce a positive correlation between the overall volume of sales to the US and the difference between the top and the worst exporting firm from a given origin. In other words, observed sales reflect the primitive distribution of attributes above some cutoff level, and the truncation affects both the volume of exports and the dispersion of observed sales.

[20] Donald and Lang (2007) show that in settings like ours, in which the number of groups (countries) is not too small and the number of observations (industry-years) per group is large, this two-step approach delivers coefficient estimates that are both consistent and fully efficient (see also Wooldrige, 2003, for a discussion). Consistent with this, we obtain similar and slightly more noisy estimates of the coefficients on the country characteristics using a one-step approach, which consists of regressing each term in (10) on the country characteristics, controlling for industry-year fixed effects and adjusting the standard errors for clustering within countries. We focus here on the two-step estimates, which are computationally less demanding and directly comparable with the regressions using country-level averages reported in columns (1)-(3).

We now study this possibility. The main challenge is that the effect of truncation depends on the shape of the unconditional distribution of attributes. Hence, to make progress, we now make some additional structural assumptions. Several papers argue that sales are well approximated by a log-normal distribution (see, for instance, Cabral and Mata, 2003, Head, Mayer and Thoenig, 2014, Bas, Mayer and Thoenig, 2017). We will confirm this feature of the sales data in our sample. Under the assumption of log-normality, 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 for our decomposition.

6.1 Identifying Selection

Assume that, in a given country-industry-year triplet, the logarithm of sales follows a normal distribution with mean $\mu_{doi,t}$ and variance $\varsigma^2_{doi,t}$, truncated from below at $\ln r^{\min}_{doi,t}$. As we will see, this will be the case if attributes follows a log-normal distribution. To estimate the variance of the unconditional distribution, $\varsigma^2_{doi,t}$, we follow Head, Mayer and Thoenig (2014) and exploit a linear relationship linking the theoretical and empirical quantiles of log sales. This approach is known as the Quantile-Quantile (QQ) estimator, and its asympthotic properties have been studied by 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 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 the 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 $\hat{F}_{doi,t}(\omega) = (\omega - 0.3) / (N_{doi,t} + 0.4)$. The theoretical quantiles are defined as:

$$\mathbb{Q}^T_{doi,t}(\omega) = \mathbb{E}(\ln r_{doi,t}) + \varsigma_{doi,t}\Phi^{-1}(\hat{F}_{doi,t}(\omega)), \tag{12}$$

where Φ is the CDF of the standard normal distribution. 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 the 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 offer us enough degrees of freedom for estimating the parameters of the QQ regression in each triplet. We find that the QQ regressions fit the data remarkably well: in the sample of triplets with 9+ varieties, the R^2 of the QQ regression equals 0.95 for the average triplet, and the standard deviation of R^2 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 is even tighter (s.d. 0.03). These findings suggest that the log-normal distribution provides a very good approximation of the distribution of exports to the US in our sample.

To have visual feeling of the distribution of our data, Appendix Figure A1 plots kernel density distributions of log exports to the US for rich and poor countries and for small and large countries separately.[21] Each kernel density distribution is drawn by pooling together all varieties exported to the US by a given group of countries over the two years, and is centered around zero by deviating the log exports of each variety from the average log exports of the corresponding exporting country. The distributions shown in the figure do indeed resemble normal distributions. Moreover, in line with our previous results, the distributions are relatively more spread out for rich and large countries, which exhibit lower mass around the mean and more mass at the tails compared to poor and small countries.

We now turn to discussing the relationship between the unconditional variance, ς^2_{doi}, and the observed variance, $\mathbb{V}(\ln s_{doi})$.[22] 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. Moreover, the ratio $\varsigma^2_{doi}/\mathbb{V}(\ln s_{doi})$ is proportional to and increasing in the cutoff, $\ln r^{\min}_{doi}$: as smaller sales are removed, the variance of the remaining distribution falls. More precisely:

$$\frac{\varsigma^2_{doi}}{\mathbb{V}(\ln s_{doi})} = \left[1 + \frac{\theta_{doi}\phi(\theta_{doi})}{1-\Phi(\theta_{doi})} - \left(\frac{\phi(\theta_{doi})}{1-\Phi(\theta_{doi})}\right)^2\right]^{-1}, \tag{13}$$

where $\theta_{doi} \equiv (\ln r^{\min}_{doi} - \mu_{doi})/\varsigma_{doi}$, and ϕ is the density of the standard normal.

Since the volume of exports is a negative function of the truncation point, we can test if the ratio $\varsigma^2_{doi}/\mathbb{V}(\ln s_{doi})$ is indeed decreasing in S_{doi}. This would be evidence of selection, i.e., it would suggest that the cutoff for exporting to the US is lower for countries that capture larger US market shares in a given industry-year. In Table 5, we report the results obtained by regressing the log ratio between the unconditional and the observed variance of log sales in each triplet, $\ln\left(\varsigma^2_{doi,t}/\mathbb{V}(\ln s_{doi,t})\right)$, on the log of countries' market shares, $\ln S_{doi,t}$, using the sample of triplets with 9+ varieties in column (1) and with 27+ varieties in column (3). We also report the results obtained by regressing $\ln\left(\varsigma^2_{doi,t}/\mathbb{V}(\ln s_{doi,t})\right)$ on the two components of market shares according to (8), namely, the extensive margin of exports, $(\ln N_{doi,t} - \ln N_{di,t})$, and the intensive margin of exports, $(\ln \bar{s}_{doi,t} - \ln \bar{s}_{di,t})$ (columns 2 and 4). We control for industry-year fixed effects in all specifications, so as to focus on cross-country variation within industry-years. Across all specifications and for both margins of trade, we find that larger exports relative to other countries are associated with a lower ratio $\varsigma^2_{doi,t}/\mathbb{V}(\ln s_{doi,t})$, which is consistent with a lower cutoff. We interpret these results as evidence of selection. In particular, they are consistent with a "pecking order" across origin countries, whereby

[21] Rich (poor) countries are those with real per-capita GDP above (below) the 75th (25th) percentile. Large (small) countries are those with population above (below) the 75th (25th) percentile.

[22] Since $\mathbb{V}(\ln s_{doi}) = \mathbb{V}(\ln r_{doi})$, from now on we use $\mathbb{V}(\ln s_{doi})$ for consistency with our previous notation.

Table 5: Export Margins and Selection

	(1)	(2)	(3)	(4)
Country market share	-0.014***		-0.008***	
	[0.000]		[0.000]	
Extensive margin		-0.026***		-0.013***
		[0.000]		[0.000]
Intensive margin		-0.003***		-0.003***
		[0.000]		[0.001]
Obs.	12437	12437	6314	6314
R2	0.31	0.38	0.31	0.34

Notes. The dependent variable is the ratio between the estimated and the actual variance of log sales in each country-industry-year triplet; estimated variances are obtained by applying the QQ estimator on each triplet, using theoretical quantiles of sales implied by the log-normal distribution. The explanatory variables are the three terms in eq. (8). All variables are in logs. All regressions are run across triplets and include industry-year fixed effects. The sample consists of triplets with at least 9 firm-products exported to the US in columns (1)-(2) and of triplets with at least 27 firm-products exported to the US in columns (3)-(4). The standard errors are corrected for heteroskedasticity. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

progressively less appealing firms export from countries selling more in a given market.

6.2 Decomposing US Imports with Selection

We now build on the previous results to implement a new decomposition that allows us to assess how selection affects the contribution of heterogeneity at explaining variation in market shares across countries. Assume that, in each country-industry pair, $\tilde{\gamma}$ is drawn from a log-normal distribution with possibly different parameters, and disregard truncation for the moment. Since sales are a power function of $\tilde{\gamma}$, they inherit the log-normal distribution, albeit with different parameters.[23] We can then use the properties of this distribution to obtain a formula that decomposes market shares across country pairs within any industry. We start from (6) and take the log of the market share captured by country o *relative to* another country x in industry i:

$$\ln \frac{S_{doi}}{S_{dxi}} = \ln \sum_{\omega \in \Omega_{doi}} \tilde{\gamma}_d(\omega)^{\sigma_i - 1} - \ln \sum_{\omega \in \Omega_{dxi}} \tilde{\gamma}_d\,(\omega)^{\sigma_i - 1}\,. \tag{14}$$

[23] The overall distribution of imports from all origins will no longer be log-normal, although it may still be approximately so.

Next, since $\sum_{\omega\in\Omega_{doi}}\tilde{\gamma}_d(\omega)^{\sigma_i-1}=\mathbb{E}\left[\tilde{\gamma}_{doi}^{\sigma_i-1}\right]\cdot N_{doi}$ and using the properties of the log-normal distribution, this equation can be rewritten as:[24]

$$\ln\frac{S_{doi}}{S_{dxi}}=(\sigma_i-1)\left\{\mathbb{E}\left[\ln\tilde{\gamma}_{doi}\right]-\mathbb{E}\left[\ln\tilde{\gamma}_{dxi}\right]\right\}+\frac{(\sigma_i-1)^2}{2}\left[\mathbb{V}\left(\ln\tilde{\gamma}_{doi}\right)-\mathbb{V}\left(\ln\tilde{\gamma}_{dxi}\right)\right]+\ln\frac{N_{doi}}{N_{dxi}}. \tag{15}$$

Equation (15) provides a decomposition that depends on readily observable variables only, and not on estimates of σ_i. To see this, use $\tilde{\gamma}_d(\omega)=s_d(\omega)^{1/(\sigma_i-1)}$ and $\mathbb{V}(\ln\tilde{\gamma}_{doi})=\mathbb{V}(\ln s_{doi})/(\sigma_i-1)^2$ to obtain:

$$\ln\frac{S_{doi,t}}{S_{dxi,t}}=\left\{\mathbb{E}\left[\ln s_{doi,t}\right]-\mathbb{E}\left[\ln s_{dxi,t}\right]\right\}+\frac{\left[\mathbb{V}\left(\ln s_{doi,t}\right)-\mathbb{V}\left(\ln s_{dxi,t}\right)\right]}{2}+\ln\frac{N_{doi,t}}{N_{dxi,t}}, \tag{16}$$

where we have added the time subscript t in view of the following empirical analysis. The reason is that, as anticipated, with CES demand a log-normal distribution of $\tilde{\gamma}$ generates log-normally distributed sales. While convenient, this new decomposition does not fully separate the effect of heterogeneity in attributes across firms. Since the logarithm is a concave function, by Jensen's inequality dispersion in sales also affects the mean of log sales, i.e., the first term of (16). Moreover, without using σ_i, (16) does not identify the role of reallocations.[25] Nevertheless, this new decomposition does quantify the importance of heterogeneity when aggregating log sales.

Equation (16) does not control for selection. If observed sales are truncated at $r_{doi,t}^{\min}$, we can rewrite the decomposition in terms of the parameters of the unconditional distribution and the truncation point:

$$\ln\frac{S_{doi,t}}{S_{dxi,t}}=(\mu_{doi,t}-\mu_{dxi,t})+\frac{\varsigma_{doi,t}^2-\varsigma_{dxi,t}^2}{2}+\ln\frac{N_{doi,t}}{N_{dxi,t}}+\ln\frac{T_{doi,t}}{T_{dxi,t}}, \tag{17}$$

where

$$T_{doi,t}=\frac{1-\Phi(\theta_{doi,t}-\varsigma_{doi,t})}{1-\Phi(\theta_{doi,t})} \tag{18}$$

is a function of $r_{doi,t}^{\min}$ and $\theta_{doi,t}=(\ln r_{doi,t}^{\min}-\mu_{doi,t})/\varsigma_{doi,t}$. Comparing (16) and (17) immediately reveals that, to control for the effect of truncation on dispersion, one simply needs to use $\varsigma_{doi,t}^2$ instead of $\mathbb{V}\left(\ln s_{doi,t}\right)$.

We can now use the new decompositions to perform various exercises. First, to assess what fraction of the variation in $\ln(S_{doi,t}/S_{dxi,t})$ is explained by each component of (16), we regress each

[24] Recall that, if $x\sim\log Normal$, then $\ln\mathbb{E}\left[x^n\right]=n\mathbb{E}\left[\ln x\right]+\frac{n^2\text{var}(\ln x)}{2}$.

[25] To see instead the effect of reallocations, we can use the properties of the log-normal distribution to substitute $\mathbb{E}\left[\ln\tilde{\gamma}\right]=\mathbb{E}\left[\tilde{\gamma}\right]-\mathbb{V}\left(\ln\tilde{\gamma}\right)/2$ into (15):

$$\ln\left(\frac{S_{doi}}{S_{dxi}}\right)^{1/(\sigma_i-1)}=\left\{\mathbb{E}\left[\tilde{\gamma}_{doi}\right]-\mathbb{E}\left[\tilde{\gamma}_{dxi}\right]\right\}+\frac{(\sigma_i-2)}{2}\left[\mathbb{V}\left(\ln\tilde{\gamma}_{doi}\right)-\mathbb{V}\left(\ln\tilde{\gamma}_{dxi}\right)\right]+\frac{\left[\ln N_{doi}-\ln N_{dxi}\right]}{\sigma_i-1}.$$

This formula shows once more that market shares are increasing in the variance of firm attributes when $\sigma_i>2$.

Table 6: Decomposition of Countries' Market Shares under Log Normality

	Difference in Log Number of Varieties	Difference in Average Log Sales	Difference in Variance of Log Sales	Difference in Variance of Log Sales (QQ)
	(1)	(2)	(3)	(4)
a) Baseline	0.487***	0.236***	0.228***	-
	[0.000]	[0.000]	[0.001]	-
b) Triplets with 9+ varieties	0.505***	0.194***	0.203***	0.184***
	[0.001]	[0.001]	[0.001]	[0.001]
c) Triplets with 27+ varieties	0.512***	0.176***	0.175***	0.160***
	[0.002]	[0.001]	[0.001]	[0.001]

Notes. The table performs the decomposition in eq. (16). Each coefficient is obtained from a separate regression of the variable indicated in the column's heading on the log relative market share between country o and country x in each industry and year. In particular, the dependent variables are: the difference in the log number of varieties exported to the US between country o and country x in each industry and year (column 1); the difference in average log sales between country o and country x in each industry and year (column 2); the difference in the actual variance of log sales, times one half, between country o and country x in each industry and year (column 3); and the difference in the QQ estimate of the variance of log sales, times one half, between country o and country x in each industry and year (column 4). Panel a) uses the whole sample of triplets (1081182 observations). Panel b) uses the sample of triplets with at least 9 firm-products exported to the US (347516 observations). Panel c) uses the sample of triplets with at least 27 firm-products exported to the US (113690 observations). The standard errors are corrected for heteroskedasticity. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

term on the right-hand side on $\ln(S_{doi,t}/S_{dxi,t})$.[26] The results for the whole sample are presented in panel a) of Table 6 and broadly confirm the findings discussed in Section 4.2: the number of firms accounts for about half of the total variation in market shares, while average and variance play a comparable role. Given that (16) is no longer an exact decomposition, there is now a residual, which however explains less than 5% of the overall variation.[27]

Next, we redo the exercise for the samples of triplets with 9+ and 27+ varieties. For these samples, we also have estimates of the unconditional variances, $\varsigma^2_{doi,t}$, obtained with the QQ estimator. The results are reported in panels b) and c). In particular, column (3) shows the contribution of differences in the observed variance, $\mathbb{V}(\ln s_{doi,t})$, while column (4) shows the contribution of differences in $\varsigma^2_{doi,t}$. Comparing these numbers shows that using an estimate of dispersion that is independent of truncation lowers the contribution of the variance by less than 2 percentage points, which corresponds to less than 10% of the overall contribution. This suggests that selection effects do not play

[26] Note that these regressions only exploit cross-country variation within industry-years, so industry-year fixed effects are redundant. To see this, note that a regression with industry-year fixed effects is equivalent to a regression with all variables in deviation from industry-year means. Because (16) decomposes relative market shares across all country pairs within each industry-year, the industry-year averages of both the dependent and the explanatory variable in these regressions are equal to zero by construction.

[27] Because the decomposition in (16) does not require estimates of σ_i, it represents an appropriate framework for assessing the sensitivity of our results to the use of alternative levels of industry aggregation. In Appendix Table A1, we reestimate the regressions reported in panel a) of Table 6 using various alternative definitions of industry and alternative industrial classifications. In particular: (i) we consider more aggregate industries according to the SIC classification, namely, 2-digit and 3-digit SIC industries (20 and 129 industries, respectively); (ii) we let each industry coincide with a 6-digit code in the HS classification (2982 industries); and (iii) we define industries as 2-digit or 4-digit HS codes (93 and 925 industries, respectively). The results are always remarkably similar to the baseline estimates. We also repeat the decomposition by aggregating sales across all products within each firm. Appendix Table A1 shows that using firms, rather than firm-products, as the basic unit of analysis has no bearing on the results.

Table 7: Sales Dispersion, Selection and Country Characteristics

	Actual Variance of Log Sales	Estimated (QQ) Variance of Log Sales	QQ/Actual Variance of Log Sales	Actual Variance of Log Sales	Estimated (QQ) Variance of Log Sales	QQ/Actual Variance of Log Sales
	(1)	(2)	(3)	(4)	(5)	(6)
Real per-capita GDP	0.130***	0.126***	-0.005*	0.091***	0.090***	-0.001
	[0.024]	[0.024]	[0.003]	[0.019]	[0.019]	[0.001]
Population	0.065***	0.056***	-0.010***	0.041**	0.034*	-0.007***
	[0.018]	[0.018]	[0.002]	[0.019]	[0.019]	[0.001]
Distance	-0.225***	-0.230***	-0.005	-0.210***	-0.201***	0.009*
	[0.058]	[0.058]	[0.006]	[0.063]	[0.064]	[0.005]
Obs.	96	96	96	79	79	79
R2	0.42	0.39	0.20	0.31	0.30	0.32

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 firm-products exported to the US (columns 1-3) and the sample of triplets with at least 27 firm-products exported to the US (columns 4-6). Real per-capita GDP and population are simple averages of these variables between the years 2002 and 2012. Distance is the population-weighted number of kilometers between each country and the US. All explanatory variables are in logs. The standard errors are corrected for heteroskedasticity. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

a major role in the decomposition.

6.3 Heterogeneity and Selection across Countries

In this section, we analyze the role of selection for explaining the correlations between firm heterogeneity and country characteristics documented in Section 5. In particular, we are interested in assessing whether larger and richer countries have a higher dispersion of firms' attributes, and hence sales, because of differences in the underlying distributions, because of a lower export cutoff, or both. To answer this question, we first regress $\ln \mathbb{V}(\ln s_{doi,t})$, $\ln \varsigma^2_{doi,t}$ and $\ln \left(\varsigma^2_{doi,t}/\mathbb{V}(\ln s_{doi,t})\right)$ 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, controlling for the log of distance to the US.

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. Two main results stand out. First, the negative and statistically significant coefficient on population in the regression for the variance ratio (columns 3 and 6) is consistent with larger countries having a lower export cutoff. We find similar evidence for real per-capita GDP, and also obtain that larger distance is associated with tougher selection, although these estimates are less precise. Second, dispersion is significantly higher in larger and richer countries even when controlling for possible truncation in the data. 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), but the difference between the two sets of estimates is tiny. Overall, all the tests presented in this section are consistent with the hypothesis of selection, but they also suggest that selection alone is unlikely to explain the observed differences in heterogeneity of sales across countries.

7 Implications

We now discuss some of the main implications of the results obtained so far. We start by showing that heterogeneity is important for welfare. 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. Next, we show how both welfare and the trade elasticity depend on heterogeneity in the workhorse model of trade with free entry and selection. Finally, we discuss our decomposition of the intensive margin in light of existing theories of trade and misallocation.

7.1 Price Indexes

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, which is in turn informative about welfare. We start by rewriting the price index in destination d and industry i as:

$$P_{di} = \left(\sum_{o} P_{doi}^{1-\sigma_i} \right)^{\frac{1}{1-\sigma_i}}, \quad \text{with} \quad P_{doi} = \left[\sum_{\omega \in \Omega_{doi}} \left(\tilde{\gamma}_{doi} \right)^{\sigma_i - 1} \right]^{\frac{1}{1-\sigma_i}}. \tag{19}$$

Note that P_{doi} is the price index of the basket of goods imported from o. Using the properties of the log-normal distribution yields:

$$\ln(1/P_{doi}) = \mathbb{E}(\tilde{\gamma}_{doi}) + \frac{\sigma_i - 2}{2} \mathbb{V}(\ln \tilde{\gamma}_{doi}) + \frac{\ln N_{doi}}{\sigma_i - 1}. \tag{20}$$

This expression shows that heterogeneity lowers price indexes when $\sigma_i > 2$.[28]

These formulas allow us to evaluate the effect of heterogeneity, as captured by the variance of log attributes, holding constant the average. Consider for example two countries, o and x, differing only in the variance of firms' attributes. Then, their relative price index is:

$$\frac{P_{doi}}{P_{dxi}} = \exp \left\{ \frac{\sigma_i - 2}{2 \left(\sigma_i - 1 \right)^2} \left[\mathbb{V}(\ln s_{dxi}) - \mathbb{V}(\ln s_{doi}) \right] \right\}, \tag{21}$$

where we substituted $\mathbb{V}(\ln \tilde{\gamma}_{doi}) = \mathbb{V}(\ln s_{doi}) / \left(\sigma_i - 1 \right)^2$. Hence, the effect of heterogeneity on price indexes can be computed from the elasticity of substitution and the difference in the variance of log

[28] To obtain (20), first follow a decomposition similar to that in (15) to obtain:

$$\ln(1/P_{doi}) = \mathbb{E}(\ln \tilde{\gamma}_{doi}) + \frac{\sigma_i - 1}{2} \mathbb{V}(\ln \tilde{\gamma}_{doi}) + \frac{\ln N_{doi}}{\sigma_i - 1}.$$

Using again the properties of the log-normal distribution to substitute $\mathbb{E}(\ln \tilde{\gamma}_{doi})$ yields the equation in the text. See Epifani and Gancia (2011) for a related result.

sales. Our estimates of σ_i have a mean of 5.6 and a median of 3.5. In this interval, the factor $\frac{(\sigma_i-2)}{2(\sigma_i-1)^2}$ ranges from 0.8 to 0.12.[29] We therefore choose the intermediate value of 0.1. Regarding $\mathbb{V}(\ln s_{doi})$, the average value in our data is 3.69, with a standard deviation of 3.11. With these numbers in mind, according to (21), a difference of 3 in $\mathbb{V}(\ln s_{doi})$ implies a price index about 35% higher in the low-variance case. Moreover, an increase in $\mathbb{V}(\ln s_{doi})$ by 0.3, comparable to the observed change over the decade covered by our data, implies a fall in the price index of 3%. Although it is important to stress that these numbers do not come from a counterfactual exercise, they nonetheless indicate that reallocations between heterogenoeus firms are an important determinant of the price index.

7.2 Welfare with Free Entry and Selection

So far, we have considered the effect of heterogeneity for price indexes conditioning on the number of firms selling in a market. When entry is endogenous, however, heterogeneity also operates through selection. To understand why, we now consider a model of trade with monopolistic competition and heterogeneous firms as in Melitz (2003) and Melitz and Redding (2014). The model is standard and hence we relegate all the derivations to Appendix B. Firms pay an entry cost F_{oi} to draw attributes ϕ (a composite of quality and efficiency) from a known distribution, $G_{oi}(\phi)$, which can vary across industries and countries.[30] To sell in each destination market d, firms must incur a fixed cost f_{doi} and an iceberg cost τ_{doi}.

In these models, the real wage, expressed in terms of the price index of any industry i, is proportional to the exit cutoff, ϕ^*_{ooi}. In turn, the exit cutoff is determined by the free-entry condition:

$$\mathbb{E}\left[\pi_{oi}\right] = f_{doi} \sum_d \int_{\phi^*_{doi}}^{\infty} \left[\left(\frac{\phi}{\phi^*_{doi}}\right)^{\sigma_i - 1} - 1\right] \mathrm{d}G_{oi}(\phi) = F_{oi}, \tag{22}$$

where

$$\phi^*_{doi} = A_{di}^{-1} \tau_{doi} \left(w_o^{\sigma_i} f_{doi}\right)^{1/(\sigma_i - 1)} \tag{23}$$

is the minimum ϕ required to serve destination d, w_o is the wage in country o and A_{di} is parameter proportional to demand. In words, entry pushes up the exit cutoff, ϕ^*_{doi}, up to the point where expected profits from selling to all possible destination, $\mathbb{E}[\pi_{oi}]$, are equal to the entry cost, F_{oi}.

What is the effect of dispersion in ϕ on welfare with endogenous entry? Whenever $\pi_{oi}(\phi)$ is a convex function of ϕ, any mean-preserving spread of the distribution $G_{oi}(\phi)$ increases expected profits from entry. Then, the free-entry condition implies that the exit cutoff and hence real wages

[29] The term $\frac{\sigma_i-2}{2(\sigma_i-1)^2}$ reaches a maximum of 0.125 for $\sigma_i = 3$. 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.

[30] In the Appendix, we allow firms to differ along two dimensions, quality, γ, and efficiency, φ. Yet, we show that the equilibrium only depends on the composite variable $\phi \equiv \varphi\gamma$, which can be taken as a synthetic measure of firm heterogeneity.

rise. Since profits are proportional to revenues, equation (22) shows the familiar result that profits per market are convex in ϕ if $\sigma_i > 2$. Hence, $\sigma_i > 2$ is a sufficient condition for welfare to increase with dispersion in ϕ. Additionally, with endogenous entry, $\pi_{oi}(\phi)$ tends to be convex because profits are bounded by zero and because more productive firms enter more markets. As a result, welfare can increase with more dispersion even if $\sigma_i < 2$.[31]

Dispersion also affects the trade elasticity, which can be used to compute the welfare gains from an observed volume of trade. As shown in Bas, Mayer and Thoenig (2017), the bilateral aggregate trade elasticity can be expressed as:

$$\frac{d\ln(\bar{r}_{doi}N_{doi})}{d\ln\tau_{doi}} = (1-\sigma_i) + \frac{r^{\min}_{doi}}{\bar{r}_{doi}}\frac{d\ln N_{doi}}{d\ln\phi^*_{doi}}, \tag{24}$$

where $\bar{r}_{doi}$ and $r^{\min}_{doi}$ are average and minimum exports per firm from origin o. The first term is the fall in trade by existing exporters, while the second is the fall in trade due to firms that stop exporting, taking into account that exiting firms have below-average exports. The expression immediately shows that dispersion, by making marginal firms smaller than the average firm, makes the trade elasticity smaller in absolute value, which in turn implies larger gains from trade.

7.3 The Intensive Margin and Selection: Theory vs Evidence

Next, we ask what our decomposition of the intensive margin of exports can teach us about models of trade with selection. Average sales per exporter from two origins, o and x, are:

$$\frac{\bar{r}_{doi}}{\bar{r}_{dxi}} = \left(\frac{\tau_{dxi}w_x}{\tau_{doi}w_o}\right)^{\sigma_i-1}\frac{\int_{\phi^*_{doi}}^{\infty}\phi^{\sigma_i-1}\frac{\mathrm{d}G_{oi}(\phi)}{1-G_{oi}\left(\phi^*_{doi}\right)}}{\int_{\phi^*_{dxi}}^{\infty}\phi^{\sigma_i-1}\frac{\mathrm{d}G_{xi}(\phi)}{1-G_{xi}\left(\phi^*_{dxi}\right)}}, \tag{25}$$

where ϕ^*_{doi} is given by (23). Equation (25) shows that dispersion of attributes may vary across countries because of differences in ϕ^*_{doi} (selection) or in the distribution $G_{oi}(\phi)$. We did find evidence of selection. In particular, the result that a higher volume of exports is associated with a lower truncation point is entirely consistent with the model, which predicts ϕ^*_{doi} to be lower if export costs (τ_{doi} and f_{doi}) are smaller. On the other hand, other things equal, a lower cutoff ϕ^*_{doi}, by adding less productive firms, lowers average exports per firm:

$$\frac{d\ln\bar{r}_{doi}}{d\ln\phi^*_{doi}} = -\left(1-\frac{r^{\min}_{doi}}{\bar{r}_{doi}}\right)\frac{d\ln N_{doi}}{d\ln\phi^*_{doi}} > 0. \tag{26}$$

[31] These results are related to Bonfiglioli, Crinò and Gancia (2018, 2019b) who show how firm heterogeneity affects the value of entry through similar channels and develop a model in which the extent of heterogeneity depends on innovation decisions.

On the contrary, what we observe in the data is a strong positive correlation between dispersion and average exports per firm. Hence, selection alone, for instance due to differences in the fixed costs f_{doi}, is not sufficient to explain the evidence.

Can differences in τ_{doi} replicate the correlation found in the data? Variable trade costs have both a direct effect on $\bar{r}_{doi}$ and an effect through selection. In particular, a lower τ_{doi} increases exports of existing firms, but it also induces firms with below-average sales to enter the market:

$$\frac{d \ln \bar{r}_{doi}}{d \ln \tau_{doi}} = (1 - \sigma_i) - \left(1 - \frac{r_{doi}^{\min}}{\bar{r}_{doi}}\right) \frac{d \ln N_{doi}}{d \ln \tau_{doi}}. \tag{27}$$

Which effect dominates depends on the distribution $G_{oi}(\phi)$. For instance, the two effects exactly cancel out if $G_{oi}(\phi)$ is Pareto, leading to the counterfactual prediction of a constant intensive margin. As shown in Fernandes et al. (2019), instead, the direct effect dominates if $G_{oi}(\phi)$ is log-normal.

Finally, can the model explain the positive correlation between total income, dispersion and average exports observed in the data? Country size should have no effect either on ϕ_{doi}^* or on average exports. On the contrary, our evidence suggests that larger countries have lower export cutoffs *and* more heterogeneous exporters, and that these are the main reasons why bigger countries have larger average exports. Holding constant τ_{doi} and f_{doi}, a higher income per capita should increase wages and hence the export cutoff, thereby *lowering* dispersion. Again, this is not what we observe in the data. In sum, in the context of this model, systematic differences in the moments of $G_{oi}(\phi)$ seems needed to fit the data.

7.4 Prices, Quality and Misallocation

While our evidence is consistent with the existence of differences in the distribution of attributes across countries, another possible interpretation is misallocation. For instance, credit constraints may prevent firms from realizing their full growth potential (e.g., Bento and Restuccia, 2017) or top firms may charge systematically higher markups in less competitive markets (e.g., Edmond, Midrigan and Xu, 2015, 2019).[32] In both cases, in the presence of misallocation, firms with better attributes would have inefficiently high prices and hence remain too small. At least to some extent, this could explain why firms are more heterogeneous in richer and larger countries. We now use the information on prices to test if the evidence is consistent with this view.

With no data on inputs, we cannot estimate productivity or markups. However, we can assess to what extent the covariance between quality and quality-adjusted prices explains why firms in less developed and smaller countries are less heterogeneous. To do so, we need to impose more structural assumptions. Following a large literature on endogenous quality choice, we posit that producing

[32] Competition in the domestic market is likely to matter for exports because it may affect the cost of producing quality.

higher-quality products requires a higher marginal cost. Hence, let prices be:

$$p_d(\omega) = \gamma_d^{\alpha}(\omega) \cdot \tilde{p}_d(\omega), \tag{28}$$

where $\alpha > 0$ is the elasticity of prices with respect to quality and $\tilde{p}_d(\omega)$ is the quality-adjusted price. Using this into the definition of $\tilde{\gamma}_d(\omega)$, we can write appeal as follows:

$$\ln \tilde{\gamma}_d(\omega) = (1-\alpha)\ln \gamma_d(\omega) - \ln \tilde{p}_d(\omega). \tag{29}$$

Given α, since we have extracted both $\tilde{\gamma}_d(\omega)$ and $\gamma_d(\omega)$ from the data, we can retrieve $\tilde{p}_d(\omega)$ as well. Equation (29) then implies that

$$\mathbb{V}(\ln \tilde{\gamma}_{doi,t}) = (1-\alpha)^2\, \mathbb{V}(\ln \gamma_{doi,t}) + \mathbb{V}(\ln \tilde{p}_{doi,t}) - 2(1-\alpha)\, \texttt{Cov}(\ln \gamma_{doi,t}, \ln \tilde{p}_{doi,t}). \tag{30}$$

Hence, dispersion of attributes, $\mathbb{V}(\ln \tilde{\gamma}_{doi,t})$, depends positively on the variance of quality and quality-adjusted prices, and negatively on their covariance. We want to test to what extent less-developed and smaller countries have a lower dispersion due to a higher covariance between $\gamma_d(\omega)$ and $\tilde{p}_d(\omega)$.

In order to implement the decomposition implied by (30), a value of α is needed. Lacking a credible source of exogenous variation in quality to estimate (28) using our data, we assess the importance of the covariance for different values of α. In particular, we consider values within the $(0, 1)$ interval, which is consistent with revenues being increasing in quality. For the sake of space, here we focus on four equally distant values within this range, i.e., $\alpha \in \{0.2, 0.4, 0.6, 0.8\}$, as this suffices to highlight the key patterns in the data. For each value of α, we compute the three components of $\mathbb{V}(\ln \tilde{\gamma}_{doi,t})$ according to (30). Across triplets, the average value of the covariance ranges from -0.016 for $\alpha = 0.2$ to -1.69 for $\alpha = 0.8$, consistent with the view that high appeal firms are also more efficient. Then, we regress $\mathbb{V}(\ln \tilde{\gamma}_{doi,t})$ and each of its components on country and industry-year fixed effects, and relate the country effects obtained from these regressions to the log of countries' population and real per-capita GDP, controlling for the log of distance to the US as in previous tables.

The results are reported in Table 8. Column (1) confirms that larger and richer countries have higher dispersion of attributes within industries. In each of the four subsequent panels, we study the relationship of population and income per capita with the three components of dispersion in (30), computed using a different value of α as indicated in the panels' headings. Independent of the chosen α, we find that larger and richer countries have higher dispersion of both quality and quality-adjusted prices. On the contrary, the covariance is unrelated to country characteristics when α is low. In fact, for $\alpha = 0.2$, the coefficients on log population and log real per-capita GDP in the covariance regression are actually negative, albeit imprecisely estimated. If anything, this suggests that, for low levels of α, smaller and less-developed countries tend to have lower covariance between

Table 8: Dispersion of Attributes and Country Characteristics

	Variance of Log Attributes	Quality Component	Price Component	Covariance Component	Quality Component	Price Component	Covariance Component
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
			a) alpha = 0.2			**b) alpha = 0.4**	
Real per-capita GDP	0.188***	0.174***	0.025***	-0.011	0.098***	0.034***	0.057*
	[0.063]	[0.045]	[0.005]	[0.022]	[0.025]	[0.010]	[0.030]
Population	0.168**	0.153***	0.021***	-0.006	0.086***	0.029**	0.053
	[0.071]	[0.050]	[0.005]	[0.024]	[0.028]	[0.011]	[0.033]
Distance	-0.397***	-0.318***	-0.071***	-0.009	-0.179***	-0.093***	-0.126**
	[0.142]	[0.110]	[0.015]	[0.046]	[0.062]	[0.024]	[0.062]
Obs.	101	101	101	101	101	101	101
R2	0.11	0.17	0.37	0.00	0.17	0.14	0.05
			c) alpha = 0.6			**d) alpha = 0.8**	
Real per-capita GDP	0.188***	0.043***	0.064***	0.081***	0.011***	0.115***	0.062***
	[0.063]	[0.011]	[0.022]	[0.030]	[0.003]	[0.040]	[0.021]
Population	0.168**	0.038***	0.056**	0.073**	0.010***	0.103**	0.056**
	[0.071]	[0.013]	[0.025]	[0.034]	[0.003]	[0.045]	[0.023]
Distance	-0.397***	-0.079***	-0.155***	-0.163**	-0.020***	-0.256***	-0.121**
	[0.142]	[0.027]	[0.050]	[0.067]	[0.007]	[0.089]	[0.047]
Obs.	101	101	101	101	101	101	101
R2	0.11	0.17	0.10	0.09	0.17	0.10	0.11

Notes. The dependent variables are the country fixed effects obtained by regressing the variables indicated in the columns' headings on country and industry-year fixed effects across country-industry-year triplets. The variance of log attributes is additively decomposed into a quality, price and covariance component as in eq. (30). The quality component is equal to the variance of log quality times the square of one minus the quality elasticity of prices (alpha). The price component is equal to the variance of log quality-adjusted prices. The covariance component is equal to minus twice the covariance between log quality and log quality-adjusted prices, times one minus the quality elasticity of prices. The value of alpha used to construct the three components is equal to 0.2 in panel a), 0.4 in panel b), 0.6 in panel c) and 0.8 in panel d). Real per-capita GDP and population are simple averages of these variables between the years 2002 and 2012. Distance is the population-weighted number of kilometers between each country and the US. All explanatory variables are in logs. The standard errors are corrected for heteroskedasticity. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

quality and quality-adjusted prices.[33] As α increases, however, the coefficients on population and income per capita initially rise. For $\alpha = 0.4$, both coefficients are positive, and for $\alpha = 0.6$ they are also highly statistically significant. As α increases further, the two coefficients remain positive and precisely estimated but decrease in size, reflecting the fact that the covariance term approaches zero as $\alpha \to 1$.

Overall, these results imply that, except for small and possibly implausible values of the quality elasticity of prices, the covariance between quality and quality-adjusted prices is indeed higher in smaller and less-developed countries. This could be because of differences in the cost of producing quality or because top-quality firms have inefficiently high prices, consistent with the view that misallocation plays an important role in compressing sales. Be as it may, misallocation alone is unlikely to be the main explanation. Indeed, even the largest coefficients obtained for $\alpha = 0.6$ imply

[33] Note that the covariance term in (30) is multiplied by -1, so a negative (positive) coefficient on a given country characteristic implies that the covariance is increasing (decreasing) in that characteristic.

that misallocation could explain at most 43% of the correlation of heterogeneity in attributes with countries' population and income per capita.

8 Conclusions

In this paper, we used highly-disaggregated, transaction-level, US import data to compare firms from virtually all countries in the world competing in a single destination market. With the help of commonly-made assumptions on demand, we decomposed the economic performance of countries into the contribution of the number of firm-products, their average attributes and reallocations around the mean. The most important and novel lessons from our analysis are that reallocations between heterogeneous firms are important for explaining countries' exports, and that firm-level heterogeneity correlates systematically with country characteristics. In particular, we found that an important reason why larger and richer countries have higher exports per firm is the fact that they have more heterogeneous exporters. While we found significant evidence consistent with selection and misallocation, we also argued that these factors do not seem to be the exclusive drivers of our results. We also showed that differences in the dispersion of firm characteristics matter not just for exports, but also for 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 of the main driving forces. For instance, richer and larger markets may be more conducive to drastic innovation with more dispersed outcomes (e.g., Bonfiglioli, Crinò and Gancia 2018, 2019b) than imitation (Benhabib, Perla and Tonetti, 2017, König, Lorenz and Zilibotti, 2016). Alternatively, random ideas, reflected in firms' attributes, may be more heterogeneous in a larger population. Another possibility is that agglomeration economies, or more in general increasing returns, may explain the effect of market size. 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). Identifying the exact mechanism through which the distribution of attributes across firms is generated and evolves seems an important direction for future research.

Appendix A Estimates of the Elasticity of Substitution

In this Appendix, we explain how we estimate the elasticity of substitution. We first present the Reverse-Weighting (RW) estimator introduced by Redding and Weinstein (2017) and then move to an alternative approach that exploits differences in sales dispersion across industries.

A.1 The Reverse-Weighting Estimator

We start by considering three equivalent expressions for the change in the price index of the basket of imported varieties in a given industry between 2002 ($t-1$) and 2012 (t). Dropping industry, origin and destination subscripts to save on notation, these expressions read as follow:

$$\frac{P_t}{P_{t-1}} = \left(\frac{\lambda_{t,t-1}}{\lambda_{t-1,t}}\right)^{\frac{1}{\sigma-1}} \left\{\sum_{\omega\in\Omega_{t,t-1}} s^*_{t-1}(\omega)\left[\frac{p_t(\omega)/\gamma_t(\omega)}{p_{t-1}(\omega)/\gamma_{t-1}(\omega)}\right]^{1-\sigma}\right\}^{\frac{1}{1-\sigma}}, \tag{A1}$$

$$\frac{P_t}{P_{t-1}} = \left(\frac{\lambda_{t,t-1}}{\lambda_{t-1,t}}\right)^{\frac{1}{\sigma-1}} \left\{\sum_{\omega\in\Omega_{t,t-1}} s^*_{t}(\omega)\left[\frac{p_t(\omega)/\gamma_t(\omega)}{p_{t-1}(\omega)/\gamma_{t-1}(\omega)}\right]^{-(1-\sigma)}\right\}^{-\frac{1}{1-\sigma}}, \tag{A2}$$

$$\frac{P_t}{P_{t-1}} = \left(\frac{\lambda_{t,t-1}}{\lambda_{t-1,t}}\right)^{\frac{1}{\sigma-1}} \frac{\widetilde{P}^*_t}{\widetilde{P}^*_{t-1}}\left(\frac{\widetilde{S}^*_t}{\widetilde{S}^*_{t-1}}\right)^{\frac{1}{\sigma-1}}, \tag{A3}$$

where $\Omega_{t,t-1}$ denotes the set of common varieties in both years; $s^*(\omega)$ denotes the share of common variety ω in expenditure on all common varieties; $\widetilde{S}^*$ and $\widetilde{P}^*$ denote the geometric averages of $s^*(\omega)$ and $p(\omega)$, respectively, computed on common varieties; and $(\lambda_{t,t-1}/\lambda_{t-1,t})^{1/(\sigma-1)}$ is the variety-adjustment term, which adjusts the common varieties price index for entering and exiting varieties.

While the three ways of expressing the change in the price index are equivalent, the formulation in (A3) is the only one that exclusively depends on prices and expenditure shares, and not also on the demand shifter γ, i.e., this formulation is money-metric. Note also that the three expressions depend on the elasticity of substitution, σ. Hence, the idea behind the RW estimator is to look for the value of σ that renders the three expressions for the change in the price index consistent with the same money-metric utility function.

Combining (A1)-(A3) and rearranging terms yields:

$$\Theta^F_{t-1,t}\left\{\sum_{\omega\in\Omega_{t,t-1}} s^*_{t-1}(\omega)\left[\frac{p_t(\omega)}{p_{t-1}(\omega)}\right]^{1-\sigma}\right\}^{\frac{1}{1-\sigma}} = \frac{\widetilde{P}^*_t}{\widetilde{P}^*_{t-1}}\left(\frac{\widetilde{S}^*_t}{\widetilde{S}^*_{t-1}}\right)^{\frac{1}{\sigma-1}}, \tag{A4}$$

$$\left(\Theta^B_{t,t-1}\right)^{-1}\left\{\sum_{\omega\in\Omega_{t,t-1}} s^*_{t}(\omega)\left[\frac{p_t(\omega)}{p_{t-1}(\omega)}\right]^{-(1-\sigma)}\right\}^{-\frac{1}{1-\sigma}} = \frac{\widetilde{P}^*_t}{\widetilde{P}^*_{t-1}}\left(\frac{\widetilde{S}^*_t}{\widetilde{S}^*_{t-1}}\right)^{\frac{1}{\sigma-1}}, \tag{A5}$$

where

$$\Theta^F_{t-1,t} \equiv \left\{ \frac{\sum_{\omega\in\Omega_{t,t-1}} s^*_{t-1}(\omega)\left[\frac{p_t(\omega)}{p_{t-1}(\omega)}\right]^{1-\sigma}\left[\frac{\gamma_t(\omega)}{\gamma_{t-1}(\omega)}\right]^{\sigma-1}}{\sum_{\omega\in\Omega_{t,t-1}} s^*_{t-1}(\omega)\left[\frac{p_t(\omega)}{p_{t-1}(\omega)}\right]^{1-\sigma}} \right\}^{\frac{1}{1-\sigma}}, \tag{A6}$$

$$\Theta^B_{t,t-1} \equiv \left\{ \frac{\sum_{\omega\in\Omega_{t,t-1}} s^*_{t}(\omega)\left[\frac{p_t(\omega)}{p_{t-1}(\omega)}\right]^{-(1-\sigma)}\left[\frac{\gamma_t(\omega)}{\gamma_{t-1}(\omega)}\right]^{-(\sigma-1)}}{\sum_{\omega\in\Omega_{t,t-1}} s^*_{t}(\omega)\left[\frac{p_t(\omega)}{p_{t-1}(\omega)}\right]^{-(1-\sigma)}} \right\}^{\frac{1}{1-\sigma}} \tag{A7}$$

are forward and backward aggregate demand shifters, respectively. These demand shifters summarize the impact of changes in the relative demand for individual varieties on the overall price index.

Identification of σ requires the following identifying assumption:

$$\Theta^F_{t-1,t} = \left(\Theta^B_{t,t-1}\right)^{-1} = 1, \tag{A8}$$

which means that changes in relative demand cancel out across varieties, so that the aggregate demand shifters are both equal to 1. Using (A8) together with (A4)-(A5), one can construct a generalized method of moment estimator for σ. In particular, the following moment functions obtain:

$$M(\sigma) \equiv \begin{pmatrix} \frac{1}{1-\sigma}\ln\left\{\sum_{\omega\in\Omega_{t,t-1}} s^*_{t-1}(\omega)\left[\frac{p_t(\omega)}{p_{t-1}(\omega)}\right]^{1-\sigma}\right\} - \ln\left[\frac{\widetilde{P}^*_t}{\widetilde{P}^*_{t-1}}\left(\frac{\widetilde{S}^*_t}{\widetilde{S}^*_{t-1}}\right)^{\frac{1}{\sigma-1}}\right] \\ -\frac{1}{1-\sigma}\ln\left\{\sum_{\omega\in\Omega_{t,t-1}} s^*_{t}(\omega)\left[\frac{p_t(\omega)}{p_{t-1}(\omega)}\right]^{-(1-\sigma)}\right\} - \ln\left[\frac{\widetilde{P}^*_t}{\widetilde{P}^*_{t-1}}\left(\frac{\widetilde{S}^*_t}{\widetilde{S}^*_{t-1}}\right)^{\frac{1}{\sigma-1}}\right] \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}. \tag{A9}$$

The RW estimator $\hat{\sigma}$ solves:

$$\hat{\sigma} = \arg\min\left\{M(\hat{\sigma})' \times \mathbb{I} \times M(\hat{\sigma})\right\}, \tag{A10}$$

where $\mathbb{I}$ is the identity matrix. Weighting the two moment conditions by the identity matrix implies that the RW estimator minimizes the sum of squared deviations of the aggregate demand shifters from zero. Hence, the RW estimator selects the value of σ that minimizes the squared deviations of the forward and backward differences of the price index from a money-metric utility function.

A.2 Exploiting Variation in Sales Dispersion across Industries

As a robustness check on our main decomposition, in panel j) of Table 3 we use an alternative set of values of σ_i. To estimate them, we exploit the observed variation in sales dispersion across industries, building on the model's insight that a higher substitutability between varieties should generate more dispersion of sales for a given distribution of attributes.

To illustrate the approach, we start by using (5) to write:

$$\ln \mathbb{V}\left(\ln s_{doi,t}\right) = 2\ln\left(\sigma_i - 1\right) + \ln \mathbb{V}\left(\ln \tilde{\gamma}_{doi,t}\right). \tag{A11}$$

This equation illustrates that a given dispersion of attributes, $\mathbb{V}\left(\ln \tilde{\gamma}_{doi,t}\right)$, translates into a larger dispersion of sales, $\mathbb{V}\left(\ln s_{doi,t}\right)$, in industries where varieties are more substitutable. If $\mathbb{V}\left(\ln \tilde{\gamma}_{doi,t}\right)$ was observed, one could estimate the structural parameter measuring the elasticity of substitution by first regressing $\ln \mathbb{V}\left(\ln s_{doi,t}\right)$ on $\ln \mathbb{V}\left(\ln \tilde{\gamma}_{doi,t}\right)$ and industry fixed effects, and then backing out the elasticities from the estimates of the fixed effects.

Unfortunately, $\mathbb{V}\left(\ln \tilde{\gamma}_{doi,t}\right)$ cannot be computed without an estimate of σ_i. Hence, we proxy for this term using observable variables that are known to influence the dispersion of attributes. The first variable is the variance of log prices, $\mathbb{V}\left(\ln p_{doi,t}\right)$. While prices are just one component of $\tilde{\gamma}$, controlling for their variance would be sufficient to proxy for $\mathbb{V}\left(\ln \tilde{\gamma}_{doi,t}\right)$ if there was a one-to-one mapping between quality and prices, as in several models of endogenous quality. Indeed, an ample empirical evidence exists that prices are a good proxy for quality (see Hottman, Redding and Weinstein, 2016, and Johnson, 2012). The second variable is the number of varieties in each triplet, $N_{doi,t}$. Indeed, previous evidence shows that dispersion may vary systematically with the number of observations over which it is computed (Bonfiglioli, Crinò and Gancia, 2018, 2019b). Finally, we control for country-time fixed effects, $\nu_{o,t}$. The latter remove time-varying country characteristics that affect sales dispersion uniformly across industries, e.g., by systematically inducing some countries to specialize in high- or low-dispersion industries. Hence, we estimate the following specification:

$$\ln \mathbb{V}\left(\ln s_{doi,t}\right) = \alpha_i + \beta_1 \ln \mathbb{V}\left(\ln p_{doi,t}\right) + \beta_2 \ln N_{doi,t} + \nu_{o,t} + \varepsilon_{doi,t}, \tag{A12}$$

where α_i are industry fixed effects and $\varepsilon_{doi,t}$ is a error term. Using the estimates of α_i, we then solve for σ_i as $\sigma_i = \left[\exp\left(\alpha_i/2\right) + 1\right]$ from (A11). The resulting estimates of σ_i are in the same ballpark as those obtained with the RW estimator, with a median value of 4.22.

It is important to note that this approach does not identify the structural parameter measuring the elasticity of substitution, for two main reasons. First, the control variables included in (A12) are not perfect proxies for $\mathbb{V}\left(\ln \tilde{\gamma}_{doi,t}\right)$. Hence, part of the dispersion of attributes remains unobserved and ends up in the error term. If the unobserved component of $\mathbb{V}\left(\ln \tilde{\gamma}_{doi,t}\right)$ systematically varied across industries, the value of σ_i backed out from the industry fixed effects would not coincide with its structural counterpart. Second, even if the control variables were perfect proxies for $\mathbb{V}\left(\ln \tilde{\gamma}_{doi,t}\right)$, some industry characteristic excluded from the model could exist that affects sales dispersion across industries. In this case, the industry fixed effect would identify not just the elasticity of substitution but also this other industry-specific component of sales dispersion.

These are important *caveats*. Nevertheless, we believe it is useful to check how the results of our main decomposition change when using this alternative approach. On the one hand, by absorbing the industry average, this approach allows us to isolate the cross-country variation in attributes that we are interested in. On the other hand, unlike the RW estimator, this alternative approach does not rest on the use of continuing varieties for estimating σ_i. It thus delivers estimates for all the 366 industries included in our sample, allowing us to check the results of our main decomposition using the full sample size.

Appendix B Trade with Heterogeneous Firms and Selection

The decompositions in Section 2.2 hold irrespective of any supply-side assumptions, that is, for any production function, any distribution of product characteristics and any market structure. However, imposing more structure allows us to gain further insights. In particular, we now embed our accounting framework into a standard model of trade with heterogeneous firms and endogenous entry as in Melitz and Redding (2014). The model features a continuum of firms, monopolistic competition and fixed costs of selling into different markets.

In each industry, every variety ω is produced by firms that are heterogeneous in their labor productivity, φ, and quality, γ. In this section, we interpret γ as capturing an intrinsic product characteristic. Since all firms with the same attributes (φ, γ) behave similarly, we can index firms by (φ, γ) and identify firms with products. The equilibrium price of a firm with attributes (φ, γ) serving market d from country o is:

$$p_{doi}\left(\varphi, \gamma\right) = \frac{\sigma_i}{\sigma_i - 1} \frac{\tau_{doi} w_o}{\varphi}, \tag{A13}$$

where w_o is the wage in country o, $\tau_{doi} \geq 1$ is the iceberg cost of shipping from o to d in industry i and $\frac{\sigma_i}{\sigma_i - 1}$ is the markup over the marginal cost charged by the firm. Note that, as long as σ, τ and w do not vary across products sold in d from a given origin in a given industry, dispersion of prices at the destination are entirely driven by dispersion of efficiency: $\mathbb{V}(\ln \varphi_{doi}) = \mathbb{V}(-\ln p_{doi})$. In more general models with endogenous markups, prices would vary also because of differences in market power across firms.

Revenue earned from selling to market d is a power function of $\varphi\gamma$, which captures the overall appeal of a firm.:

$$r_{doi}(\varphi\gamma) = P_{di}^{\sigma_i} C_{di} \left(\frac{\sigma_i - 1}{\sigma_i} \frac{\gamma\varphi}{\tau_{doi} w_o} \right)^{\sigma_i - 1}. \tag{A14}$$

Profits earned in market d are a fraction σ_i of revenue minus the fixed cost of serving the market, $w_o f_{doi}$:

$$\pi_{doi}\left(\varphi\gamma\right) = \frac{r_{doi}(\varphi\gamma)}{\sigma_i} - w_o f_{doi}. \tag{A15}$$

A firm finds it profitable to serve market d if and only if $\varphi\gamma$ is sufficiently high. Define $(\varphi\gamma)^*_{doi}$ as the minimum level of $\varphi\gamma$ such that a firm breaks even in market d: $\pi_{doi}\left((\varphi\gamma)^*_{doi}\right) = 0$. Then, revenue from market d of a firm located in country o and operating in industry i can be expressed as:

$$r_{doi}(\varphi\gamma) = r^*_{doi}\left[\frac{\varphi\gamma}{(\varphi\gamma)^*_{doi}}\right]^{\sigma_i - 1}, \tag{A16}$$

where $r^*_{doi} = \sigma_i w_o f_{doi}$. Note that export participation, quantities and the price index all depend on the composite variable $\varphi\gamma$, which can be taken as a synthetic measure of firm heterogeneity. Selection into exporting of the most productive firms implies that the distribution of sales in a foreign destination depends on the distribution of characteristics of domestic firms in any country of origin truncated at the cutoff $(\varphi\gamma)^*_{doi}$.

Next, we study how firm attributes are determined. We assume that, upon paying an entry cost $w_o F_{oi}$, firms can draw their attributes from a known distribution. Although attributes are two-dimensional (φ, γ), tractability is preserved by the fact that, for the purpose of determining the equilibrium allocation, quality and efficiency can be collapsed into a one-dimensional object, the product $\varphi\gamma$, which can be taken as a single measure of performance. We simplify the notation by defining this variable $\phi \equiv \varphi\gamma$ and denote with $G_{oi}(\phi)$ the distribution from which ϕ is drawn. This distribution is allowed to vary across industries and countries.

Firms enter until expected profits are equal to the entry cost. Using (A15) and (A16), expected profits from selling to market d can be expressed as:

$$\mathbb{E}\left[\pi_{doi}\right] = \int_0^\infty \pi_{doi}\left(\varphi\gamma\right) \mathrm{d}G_{oi}\left(\varphi\gamma\right) = w_o f_{doi} \int_{\phi^*_{doi}}^\infty \left[\left(\frac{\phi}{\phi^*_{doi}}\right)^{\sigma - 1} - 1\right] \mathrm{d}G_{oi}\left(\phi\right), \tag{A17}$$

where ϕ^*_{doi} is defined by $\pi_{doi}\left(\phi^*_{doi}\right) = 0$. Using (A14) we can solve:

$$\phi^*_{doi} = A^{-1}_{di} \tau_{doi} \left(w_o^{\sigma_i} f_{doi}\right)^{1/(\sigma_i - 1)}, \tag{A18}$$

where $A_{di} \equiv (\sigma_i - 1)\left[C_{di}(P_{di}/\sigma_i)^{\sigma_i}\right]^{1/(\sigma_i - 1)}$.

Expected profits from selling in all potential markets are

$$\mathbb{E}\left[\pi_{oi}\right] = \sum_d \mathbb{E}\left[\pi_{doi}\right] = w_o f_{doi} \sum_d \int_{\phi^*_{doi}}^\infty \left[\left(\frac{\phi}{\phi^*_{doi}}\right)^{\sigma_i - 1} - 1\right] \mathrm{d}G_{oi}\left(\phi\right). \tag{A19}$$

Next, we express the cutoff for serving any market as a function of the domestic cutoff (or exit cutoff):

$$\frac{\varphi_{ooi}^*}{\varphi_{doi}^*} = \tau_{doi}^{-1} \left(\frac{w_d L_d P_{di}^{\sigma_i - 1}}{w_o L_o P_{oi}^{\sigma_i - 1}} \frac{f_{ooi}}{f_{doi}} \right)^{1/(\sigma_i - 1)} \equiv \rho_{doi}. \tag{A20}$$

Setting $\mathbb{E}\left[\pi_{oi}\right]$ equal to the entry cost $w_o F_{oi}$ and using $\varphi_{doi}^* = \varphi_{ooi}^* / \rho_{doi}$ we obtain an equation that defines implicitly the domestic cutoff ϕ_{ooi}^*:

$$F_{oi} = f_{doi} \sum_d \int_{\phi_{ooi}^*/\rho_{doi}}^{\infty} \left[\left(\frac{\phi \rho_{doi}}{\phi_{ooi}^*} \right)^{\sigma_i - 1} - 1 \right] \mathrm{d}G_{oi}\left(\phi\right). \tag{A21}$$

Finally, we can show that real wages, expressed in terms of the price index of any industry i, are also a function of the domestic cutoffs, ϕ_{ooi}^*. Combining the break-even condition for the marginal firm, $r\left(\phi_{ooi}^*\right) = \sigma_i w_o f_{ooi}$, with (A14) and substituting $C_{oi} = \beta_i w_o L_o / P_{oi}$ yields:

$$\frac{w_o}{P_{oi}} = \phi_{ooi}^* \frac{\sigma_i - 1}{\sigma_i} \left(\frac{\beta_i L_o}{\sigma_i f_{ooi}} \right)^{\frac{1}{\sigma_i - 1}}. \tag{A22}$$

This expression shows that the effect of the distribution of firms' attributes on real wages is entirely summarized by its effect on the domestic cutoffs. A higher cutoff, i.e., more selection, means lowers prices and hence higher real wages.

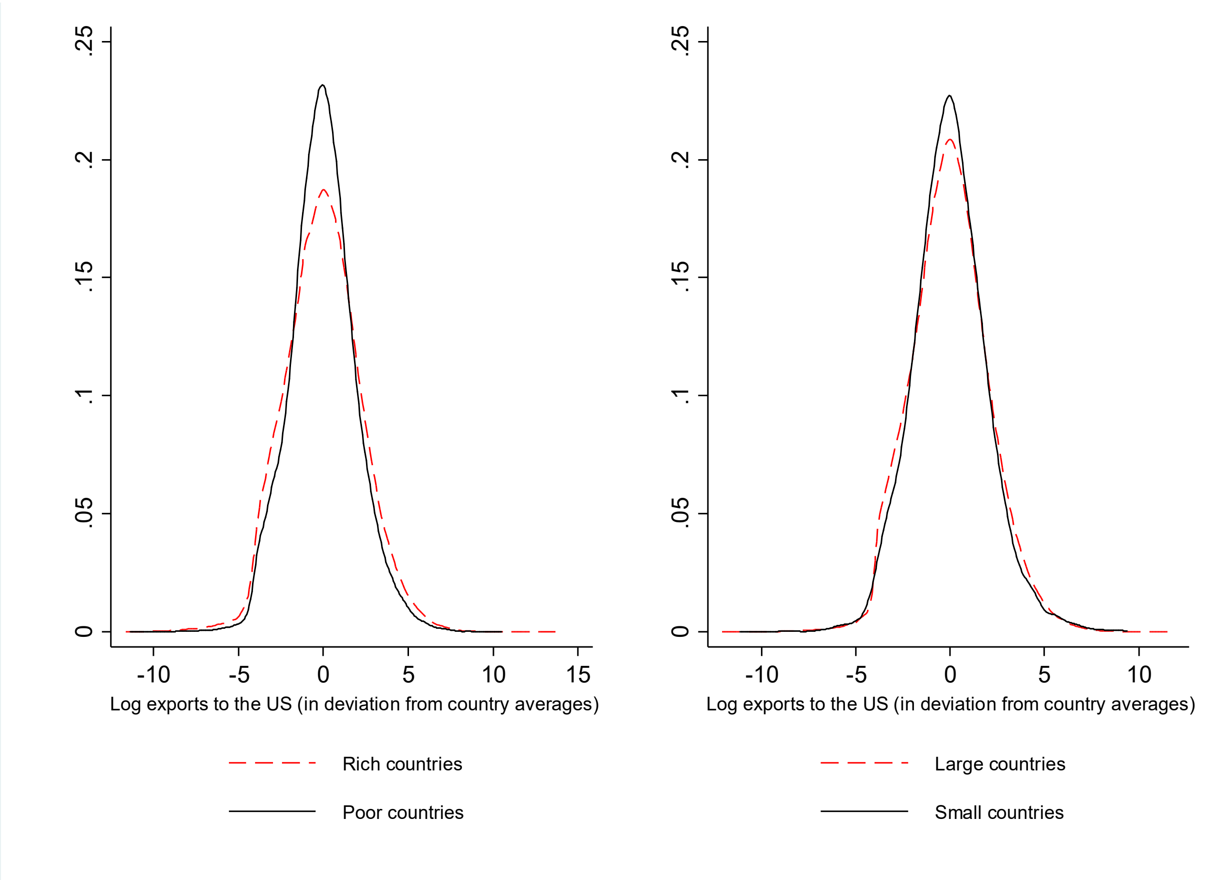

Notes. Each curve corresponds to the kernel density distribution of log exports to the US for a different group of exporting countries. Rich (poor) countries are those whose real per-capita GDP (averaged between the years 2002 and 2012) is above (below) the 75th (25th) percentile. Large (small) countries are those whose population (averaged between the years 2002 and 2012) is above (below) the 75th (25th) percentile. Each distribution is drawn by pooling together all the varieties exported to the US by a given group of countries over the two years, and is centered around zero by deviating the log exports of each variety from the average log exports of the corresponding exporting country.

Figure A1: Distribution of Log Exports to the US by Group of Exporting Countries

Table A1: Decomposition of Countries' Market Shares under Log Normality (Alternative Definitions of Industry and Variety)

	Difference in Log Number of Varieties	Difference in Average Log Sales	Difference in Variance of Log Sales
	(1)	(2)	(3)
a) Industry: 3-digit SIC	0.532***	0.190***	0.203***
	[0.000]	[0.000]	[0.001]
b) Industry: 2-digit SIC	0.619***	0.115***	0.162***
	[0.001]	[0.001]	[0.001]
c) Industry: 6-digit HS	0.417***	0.317***	0.235***
	[0.000]	[0.000]	[0.001]
d) Industry: 4-digit HS	0.453***	0.276***	0.229***
	[0.000]	[0.000]	[0.001]
e) Industry: 2-digit HS	0.541***	0.189***	0.198***
	[0.001]	[0.001]	[0.001]
f) Variety: Firm	0.464***	0.248***	0.233***
	[0.000]	[0.000]	[0.001]

Notes. The table performs the decomposition in eq. (16). Each coefficient is obtained from a separate regression of the variable indicated in the column's heading on the log relative market share between country o and country x in each industry and year. In particular, the dependent variables are: the difference in the log number of varieties exported to the US between country o and country x in each industry and year (column 1); the difference in average log sales between country o and country x in each industry and year (column 2); and the difference in the actual variance of log sales, times one half, between country o and country x in each industry and year (column 3). Industries are defined according to: the 3-digit level of the SIC classification in panel a) (688008 observations); the 2-digit level of the SIC classification in panel b) (220024 observations); the 6-digit level of the HS classification in panel c) (1531266 observations); the 4-digit level of the HS classification in panel d) (1300250 observations); and the 2-digit level of the HS classification in panel e) (502242 observations). In panel f), industries are defined according to the 4-digit level of the SIC classification and firm-level sales (obtained by summing sales across all products exported by a firm) are used instead of firm-product-level sales (1061352 observations). The standard errors are corrected for heteroskedasticity. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

References

[1] Autor, David, David Dorn, Lawrence Katz, Christina Patterson and John Van Reenen (2020). "The Fall of the Labor Share and the Rise of Superstar Firms," *Quarterly Journal of Economics*, 135, 645–709.

[2] Baqaee, David Rezza and Emmanuel Farhi (2020). "Productivity and Misallocation in General Equilibrium," *Quarterly Journal of Economics*, 135, 105–163.

[3] Bartelsman, Eric, John Haltiwanger and Stefano Scarpetta (2013). "Cross-Country Differences in Productivity: The Role of Allocation and Selection," *American Economic Review*, 103, 305–334.

[4] Bas, Maria, Thierry Mayer and Mathias Thoenig (2017). "From Micro to Macro: Demand, Supply and Heterogeneity in the Trade Elasticity," *Journal of International Economics*, 108, 1–19.

[5] Benhabib, Jess, Jesse Perla and Christopher Tonetti (2017). "Reconciling Models of Diffusion and Innovation: A Theory of the Productivity Distribution and Technology Frontier," *NBER Working Paper No. 23095.*

[6] Bento, Pedro and Diego Restuccia (2017). "Misallocation, Establishment Size, and Productivity," *American Economic Journal: Macroeconomics*, 9, 267–303.

[7] Bernard Andrew, Bradford Jensen, Stephen Redding and Peter Schott (2009). "The Margins of US Trade," *American Economic Review Papers and Proceedings*, 99, 487–493.

[8] Bernard, Andrew, Bradford Jensen, Stephen Redding and Peter Schott (2018). "Global Firms," *Journal of Economic Literature*, 56, 565–619.

[9] Bloom, Nicholas, Raffaella Sadun and John Van Reenen (2016). "Management as a Technology?" *NBER Working Paper No. 22327.*

[10] Bonfiglioli, Alessandra, Rosario Crinò and Gino Gancia (2018). "Betting on Exports: Trade and Endogenous Heterogeneity," *Economic Journal*, 128, 612–651.

[11] Bonfiglioli, Alessandra, Rosario Crinò and Gino Gancia (2019a). "Concentration in International Markets: Evidence from US Imports," *CEPR Discussion Paper No. 13566.*

[12] Bonfiglioli, Alessandra, Rosario Crinò and Gino Gancia (2019b). "Trade, Finance and Endogenous Heterogeneity," *Journal of the European Economic Association*, 17, 79–130.

[13] Bonfiglioli, Alessandra and Gino Gancia (2019). "Heterogeneity, Selection and Labor Market Disparities," *Review of Economic Dynamics*, 31, 305–325.

[14] Cabral, Luís and José Mata (2003). "On the Evolution of the Firm Size Distribution: Facts and Theory," *American Economic Review*, 93, 1075–1090.

[15] Caselli, Francesco (2005). "Accounting for Cross-Country Income Differences," in *Handbook of Economic Growth* (Aghion and Durlauf, eds.), Elsevier, Volume 1, Part A, 679–741.

[16] Chaney Thomas (2008). "Distorted Gravity: The Intensive and Extensive Margins of International Trade," *American Economic Review*, 98, 1707–1721.

[17] Costinot, Arnaud and Andrés Rodríguez-Clare (2014). "Trade Theory with Numbers: Quantifying the Consequences of Globalization," in *Handbook of International Economics* (Gopinath, Helpman and Rogoff, eds.), Elsevier, Volume 4, 197–261.

[18] Crinò, Rosario and Laura Ogliari (2017). "Financial Imperfections, Product Quality, and International Trade," *Journal of International Economics*, 104, 63–84.

[19] Donald, Stephen and Kevin Lang (2007). "Inference with Difference-in-Differences and Other Panel Data," *Review of Economics and Statistics*, 89, 221–233.

[20] Eaton, Jonathan and Samuel Kortum (2002). "Technology, Geography and Trade," *Econometrica*, 70, 1741–1779.

[21] Eaton Jonathan, Samuel Kortum and Francis Kramarz (2011). "An Anatomy of International Trade: Evidence From French Firms," *Econometrica*, 79, 1453–1498.

[22] Edmond, Chris, Virgiliu Midrigan and Daniel Xu (2015). "Competition, Markups, and the Gains from International Trade," *American Economic Review*, 105, 3183–3221.

[23] Edmond, Chris, Virgiliu Midrigan and Daniel Xu (2019). "How Costly Are Markups?" *NBER Working Paper No. 24800*.

[24] Epifani, Paolo and Gino Gancia (2011). "Trade, Markup Heterogeneity and Misallocations," *Journal of International Economics*, 83, 1–13.

[25] Feenstra, Robert (1994). "New Product Varieties and the Measurement of International Prices," *American Economic Review*, 84, 157–177.

[26] Feenstra, Robert and John Romalis (2014). "International Prices and Endogenous Quality," *Quarterly Journal of Economics*, 129, 477–527.

[27] Feenstra, Robert, John Romalis and Peter Schott (2002). "U.S. Imports, Exports and Tariff Data, 1989–2001." *NBER Working Paper No. 9387*.

[28] Feenstra, Robert and David Weinstein (2017). "Globalization, Competition, and U.S. Welfare," *Journal of Political Economy*, 125, 1041–1074.

[29] Fernandes, Ana, Caroline Freund and Martha Denisse Pierola (2016). "Exporter Behavior, Country Size and Stage of Development: Evidence from the Exporter Dynamics Database," *Journal of Development Economics*, 119, 121–137.

[30] Fernandes, Ana, Peter Klenow, Sergii Meleshchuk, Martha Denisse Pierola and Andres Rodriguez-Clare (2019). "The Intensive Margin in Trade," *Mimeo*, Stanford University.

[31] Freund, Caroline and Martha Denisse Pierola (2015). "Export Superstars," *Review of Economics and Statistics*, 97, 1023–1032.

[32] Gancia, Gino, Andreas Mueller and Fabrizio Zilibotti (2013). "Structural Development Accounting," in *Advances in Economics and Econometrics: Theory and Applications* (Acemoglu, Arellano and Dekel, eds.), Tenth World Congress of the Econometric Society, NY/Cambridge: Cambridge University Press, Volume 2, 373–418.

[33] Gennaioli, Nicola, Rafael La Porta, Florencio López-de-Silanes and Andrei Shleifer (2013). "Human Capital and Regional Development," *Quarterly Journal of Economics*, 128, 105-164.

[34] Hall, Robert and Charles Jones (1999). "Why Do Some Countries Produce So Much More Output Per Worker Than Others?," *Quarterly Journal of Economics*, 114, 83–116.

[35] Hallak, Juan Carlos and Peter Schott (2011). "Estimating Cross-Country Differences in Product Quality," *Quarterly Journal of Economics*, 126, 417–474.

[36] Head, Keith, Thierry Mayer and Mathias Thoenig (2014). "Welfare and Trade Without Pareto," *American Economic Review Papers and Proceedings*, 104, 310–316.

[37] Hottman, Colin, Stephen Redding and David Weinstein (2016). "Quantifying the Sources of Firm Heterogeneity," *Quarterly Journal of Economics*, 131, 1291–1364.

[38] Hsieh, Chang-Tai and Peter Klenow (2009). "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124, 1403–1448.

[39] Hummels, David and Peter Klenow (2005). "The Variety and Quality of a Nation's Exports," *American Economic Review*, 95, 704–723.

[40] Johnson, Robert (2012). "Trade and Prices with Heterogeneous Firms," *Journal of International Economics*, 86, 43–56.

[41] Kamal, Fariha, C.J. Krizan and Ryan Monarch (2015). "Identifying Foreign Suppliers in U.S. Merchandise Import Transactions," *Board of Governors of the Federal Reserve System - International Finance Discussion Paper No. 1142.*

[42] Khandelwal, Amit (2010). "The Long and Short (of) Quality Ladders," *Review of Economic Studies*, 77, 1450–1476.

[43] Khandelwal, Amit, Peter Schott, and Shang-Jin Wei (2013). "Trade Liberalization and Embedded Institutional Reform: Evidence from Chinese Exporters," *American Economic Review*, 103, 2169–2195.

[44] König, Michael, Jan Lorenz and Fabrizio Zilibotti (2016). "Innovation vs. Imitation and the Evolution of Productivity Distributions," *Theoretical Economics*, 11, 1053–1102.

[45] Kratz, Marie and Sidney Resnick (1996). "The QQ-Estimator and Heavy Tails," *Stochastic Models*, 12, 699–724.

[46] Matsuyama, Kiminori and Philip Ushchev (2017). "Beyond CES: Three Alternative Classes of Flexible Homothetic Demand Systems," *CEPR Discussion Paper No. 12210.*

[47] Mayer Thierry, Marc Melitz and Gianmarco Ottaviano (2014). "Market Size, Competition, and the Product Mix of Exporters," *American Economic Review*, 104, 495–536.

[48] Melitz, Marc (2003). "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71, 1695–1725.

[49] Melitz, Marc and Stephen Redding (2014). "Heterogeneous Firms and Trade," in *Handbook of International Economics* (Gopinath, Helpman and Rogoff, eds.), Elsevier, Volume 4, 1–54.

[50] Mrázová, Monika and Peter Neary (2017). "Not So Demanding: Demand Structure and Firm Behavior," *American Economic Review*, 107, 3835–3874.

[51] Mrázová, Monika, Peter Neary and Mathieu Parenti (2017). "Sales and Markup Dispersion: Theory and Empirics," *Mimeo*, University of Oxford.

[52] Poschke, Markus (2018). "The Firm Size Distribution across Countries and Skill-Biased Change in Entrepreneurial Technology," *American Economic Journal: Macroeconomics*, 10, 1–41.

[53] Redding, Stephen and David Weinstein (2017). "A Unified Approach to Aggregate Price and Welfare Measurement," *NBER Working Paper No. 22479*.

[54] Redding, Stephen and David Weinstein (2018). "Accounting for Trade Patterns," *Mimeo*, Princeton University.

[55] Redding, Stephen and David Weinstein (2020). "Measuring Aggregate Price Indexes with Taste Shocks: Theory and Evidence for CES Preferences," *Quarterly Journal of Economics*, 135, 503–560.

[56] Restuccia, Diego and Richard Rogerson (2008). "Policy Distortions and Aggregate Productivity with Heterogeneous Establishments," *Review of Economic Dynamics* 11, 707–720.

[57] Sampson, Thomas (2014). "Selection into Trade and Wage Inequality," *American Economic Journal: Microeconomics*, 6, 157–202.

[58] Schott, Peter (2004). "Across-Product versus Within-Product Specialization in International Trade," *Quarterly Journal of Economics,* 119, 647–678.

[59] Syverson, Chad (2011). "What Determines Productivity," *Journal of Economic Literature*, 49, 326–365.

[60] Wooldridge, Jeffrey (2003). "Cluster-Sample Methods in Applied Econometrics," *American Economic Review Papers and Proceedings*, 93, 133–138.