

The Intensive Margin in Trade: How Big and How Important?

Ana M. Fernandes Peter J. Klenow Sergii Meleshchuk Martha Denisse Pierola

Andrés Rodríguez-Clare *

May 4, 2022

Abstract

In the benchmark trade models with monopolistic competition featuring a constant trade elasticity, variation in bilateral trade flows happens either entirely on the intensive margin of exports per exporting firm (Krugman) or entirely on the extensive margin of the number of exporting firms (Melitz-Pareto). Using the World Bank's Exporter Dynamics Database featuring firm-level exports from 50 countries, we find that around 50% of variation in exports is along each margin, implying that the trade elasticity may not be constant, and gains from trade may differ from those in benchmark models. We show that moving from a Pareto to a lognormal distribution gives a positive role for both margins, and we use likelihood methods to estimate a generalized Melitz model with a joint lognormal distribution for firm productivity, fixed costs, and demand shifters. Using “exact hat algebra” we quantify how trade costs affect trade flows and welfare in the estimated model. We find similar welfare effects to those in the Melitz-Pareto model but significant differences in the implied trade flows.

*Ana Fernandes, World Bank, afernandes@worldbank.org; Peter J. Klenow, Stanford and NBER, klenow@stanford.edu; Sergii Meleshchuk, IMF, smeleshchuk@imf.org; Martha Denisse Pierola, Inter-American Development Bank, mpierola@iadb.org; Andrés Rodríguez-Clare, UC Berkeley and NBER, arc@berkeley.edu. We are grateful to Arnaud Costinot, Caroline Freund, Cecile Gaubert, Keith Head, Sam Kortum, Thierry Mayer, Eduardo Morales and Jesse Perla for useful discussions, to seminar participants at various institutions, and to Matthias Hoelzlein and Nick Sander for outstanding research assistance. The World Bank provided access to the *Exporter Dynamics Database*. Research for this paper has in part been supported by the World Bank's Multidonor Trust Fund for Trade and Development and the Strategic Research Program on Economic Development, as well as the Stanford Institute for Economic Policy Research. The findings expressed in this paper are those of the authors and do not necessarily represent the views of the World Bank or its member countries. The views expressed herein are those of the authors and should not be attributed to the IMF, its Executive Board, or its management.

1. Introduction

The trade elasticity (i.e., the elasticity of trade with respect to trade costs) is a crucial statistic for the gains from trade (Arkolakis, Costinot and Rodríguez-Clare, 2012) (henceforth ACR). In workhorse trade models such as Krugman and Melitz-Pareto the trade elasticity is a constant pinned down by a single structural parameter. In the Krugman (1980) model firms sell in every destination and all variation in bilateral trade flows is on the intensive margin (i.e., average exports per firm), so the trade elasticity is the constant elasticity of substitution across products (minus one). Melitz (2003) brings the extensive margin to life with fixed costs of exporting, and emphasizes the importance of selection of firms into exporting. In a popular version of the Melitz model with a Pareto distribution of productivity introduced by Chaney (2008), average exports per firm is constant and all variation in bilateral trade flows is along the firm extensive margin — implying the trade elasticity is given by the Pareto shape parameter.

These stylized models have been criticized for being more tractable than realistic. The implication that all variation in bilateral trade flows happens along either the intensive or the extensive margins is clearly extreme and unlikely to be consistent with data. And, if both margins are operative, then the trade elasticity need not be constant. Head, Mayer and Thoenig (2014) and Melitz and Redding (2015) explore non-Pareto productivity distributions and show that they generate variation in the trade elasticity across countries and time, with potentially important implications for the gains from trade. Does their critique have empirical bite? Do deviations from the constant-elasticity polar models which better fit the data result in starkly different gains from trade?

We tackle these questions in three steps. First, we exploit firm-level export data for a large set of countries to investigate whether we are anywhere close to the all-intensive-margin or all-extensive-margin extremes implied by the Krugman and Melitz-Pareto models. We find that the two margins have a roughly equal role to play in accounting for the variation in bilateral trade flows. Second, we show that, when paired with a lognormal distribution of firm productivity, the Melitz model is entirely consistent with the empirical patterns we observe. Finally, we study the welfare effects of trade liberalization in our estimated Melitz-lognormal model and find them to be quite close to those in the standard Melitz-Pareto model. Despite the trade elasticity varying substantially across trade partners, the gains from trade do *not* differ much from the Melitz-Pareto benchmark. Thus, the ACR framework provides a surprisingly accurate approximation to the gains from trade even in a context in which the trade elasticity is variable and both the intensive and extensive margins of trade are active.

To elaborate, we use the World Bank's *Exporter Dynamics Database* (hereafter EDD) to systematically examine the importance of the firm extensive and intensive margins in driving bilateral trade flows. The EDD covers firm-level exports from 60 (mostly developing) countries to all destination countries in most years from 2003 to 2013. For 50 of the countries, every firm's exports to each destination in a given year

can be broken down into HS 6-digit products.¹ Having many origin and destination countries enables us to study firm margins while allowing for origin-year and destination-year fixed effects that control for differences in population, wages, and other country characteristics.² We find that between 40 and 60 percent of the variation in exports across origin-destination pairs is accounted for by the intensive margin, with the rest accounted for by the extensive margin. This breakdown is robust to using different country samples or sets of fixed effects, excluding country pairs with few exporters or tiny exporters, and looking within industries. If we place exporting firms into percentiles for each trading pair and look across pairs, the importance of the intensive margin in explaining overall exports rises steadily from around 20 percent for the smallest exporters to over 50 percent for the largest exporters.

We interpret the finding that up to 60 percent of the variation in bilateral trade flows are explained by the extensive margin as providing broad support for the Melitz (2003) model. But finding the intensive margin accounts for at least 40 percent of variation — even allowing for origin-year and destination-year fixed effects — contradicts the Melitz-Pareto model with fixed trade costs varying only because of separate origin and destination components. In this model all variation in bilateral exports should occur through the number of exporters (the extensive margin). Lower variable trade costs should stimulate exports of a given firm, but draw in marginal exporting firms to the point that average exports per firm (the intensive margin) is unchanged. This exact offset is a special property of the Pareto distribution.³

We explore several potential explanations for the prominent intensive margin in the EDD data while retaining a Melitz-Pareto core, namely fixed trade costs that vary across country pairs, multi-product firms, and firm granularity. We do this because Melitz-Pareto has become an important benchmark model in international trade. It is consistent with many firm-level facts (Eaton et al., 2011), generates a gravity equation (Chaney, 2008), and yields a simple summary statistic for the welfare gains from trade (Arkolakis et al., 2012). Unfortunately, none of the extensions of the Melitz-Pareto model that we consider fits the intensive margin stylized facts that we uncover with the EDD, so we drop the Pareto assumption of firm productivity and adopt instead a lognormal distribution. Head, Mayer and Thoenig (2014) analyze how the welfare gains from trade in the Melitz model differ with a lognormal instead of a Pareto distribution. Bas, Mayer and Thoenig (2017) show how the trade elasticity varies with a lognormal distribution. Both papers marshal evidence from firms in France and China in favor of the lognormal distribution.

¹ See Fernandes, Freund and Pierola (2016) for a detailed description of the dataset.

² Most firm-level empirical trade studies have one or at most a few exporting countries. Bernard, Jensen, Redding and Schott (2007) decompose exports from the U.S. to other countries. Eaton, Eslava, Kugler and Tybout (2008) analyze firm-level exports for Colombia, Eaton, Kortum and Kramarz (2011) do so for France, Eaton, Kortum and Sotelo (2012) for Denmark and France, Manova and Zhang (2012) for China, and Arkolakis, Ganapati and Muendler (2021) for Brazil, Chile, Denmark and Norway.

³ This property of the Melitz-Pareto model extends to environments with demand and fixed costs that are idiosyncratic to firm-destinations (Eaton et al., 2011); convex marketing costs (Arkolakis, 2010); non-CES preferences (Arkolakis, Costinot, Donaldson and Rodríguez-Clare, 2019); non-monopolistic competition (Bernard, Eaton, Jensen and Kortum, 2003), and multinational production (Arkolakis, Ramondo, Rodríguez-Clare and Yeaple, 2018).

We consider a Melitz model with demand and fixed trade cost shocks that are specific to each firm-destination, as in (Eaton et al., 2011), but with a firm productivity distribution that is lognormal rather than Pareto. In particular, we assume that each firm is characterized by a productivity parameter, an idiosyncratic demand shifter, and a fixed cost for each destination market, all drawn from a multivariate lognormal distribution. We allow for a non-zero covariance between the demand shifter and the fixed cost in each destination, but set all other covariances to zero. One appealing feature of this setup is that it is amenable to likelihood estimation methods. As the likelihood may not be a concave function of the parameters, and since we have a large number of parameters to estimate (means, variances, one covariance, and trade costs), we rely on the estimation methodology by Chernozhukov and Hong (2003).

Our estimation shows that a lognormal distribution for firm productivity can successfully generate a sizable intensive margin elasticity. When variable trade costs fall and fixed costs are constant, the productivity cutoff falls and the ratio of mean to minimum exports per firm increases under the lognormal distribution (while being constant under Pareto).⁴ As in the data, the intensive margin elasticity rises steadily with the size percentile of exporters under a lognormal productivity distribution.

We finish by studying the implications of our empirical findings for the impact of trade liberalization. We show how to extend the Dekle, Eaton and Kortum (2008) “exact hat algebra” to a Melitz model with a general distribution of firm-level productivity, fixed export costs, and destination-specific demand shifters. We then compute the effects of changes in trade costs on trade flows and welfare in our full Melitz-lognormal model. We compare these effects to those in the standard Melitz-Pareto model with the Pareto shape parameter estimated to fit the average trade elasticity implied by our estimated Melitz-lognormal model. The welfare effects of trade liberalization in this Melitz-Pareto approximation are very close to those in the Melitz-lognormal model, although the effects on trade flows do differ significantly.⁵

Our counterfactual analysis is related to Head et al. (2014) and Melitz and Redding (2015). Using a lognormal distribution and a bounded-Pareto distribution, respectively, they show that the trade elasticity is not constant across countries or time. They then draw implications for the welfare effects of trade in calibrated symmetric two-country models. Our conclusion — that the Melitz-Pareto model offers a good approximation to the welfare effects when the data generating process is our estimated full Melitz-lognormal model — is consistent with the finding in Head et al. (2014) that their “macro-data approach” to calibration leads to similar results across the lognormal and Pareto models. In contrast, Melitz and Redding (2015) show that the formula proposed in Arkolakis et al. (2012) to compute welfare

⁴The intensive margin comes alive under other thin-tailed productivity distributions, such as bounded Pareto as in Feenstra (2018). However, a bounded Pareto distribution loses the analytical convenience of the unbounded Pareto while lacking the estimation convenience of the lognormal distribution.

⁵Our approach and findings bear some resemblance to those in contemporaneous work by Head et al. (2021). They consider a model with rich patterns of substitutability across varieties and variable markups as the true data generating process and explore the extent to which counterfactuals differ if one wrongly estimates and applies a simple CES-monopolistic competition model on the generated data. They find that the CES model serves as a very good approximation.

changes given changes in trade shares (their formula for “ex-post welfare evaluation”) is no longer accurate if the Pareto productivity distribution is bounded from above. We find the Melitz-Pareto model to be a good approximation because variation in the trade elasticity is much smaller in our full Melitz-lognormal model estimated on the EDD data than in their symmetric Melitz model with a truncated Pareto distribution calibrated to match the relative size of exporting and non-exporting U.S. firms.⁶

To recap, this paper makes several contributions to the literature. First, we use the EDD to establish a new stylized fact, namely that between 40% and 60% of the variation in exports across country pairs takes place along the intensive margin, with this margin being important all along the firm-size distribution. Second, we show that the Melitz-Pareto model cannot match this fact, even allowing for a number of extensions. Third, we show that a lognormal firm productivity distribution generates a positive role for the intensive margin as required by the data. Fourth, we use likelihood methods to estimate a Melitz model with a lognormal distribution for productivity plus idiosyncratic demand shocks and idiosyncratic fixed costs. Finally, we extend the exact hat algebra approach to a generalized Melitz model and use it to explore counterfactual trade flow and welfare implications in Melitz-lognormal versus Melitz-Pareto.

The rest of the paper is organized as follows. Section 2 describes the EDD data and documents the importance of the extensive and intensive margins in export variation. Section 3 contrasts the EDD facts with the predictions of the Melitz-Pareto model (with a continuum of single product firms, multi-product firms, or a finite number of firms). Section 4 shows how the intensive margin in the Melitz model changes when we drop the Pareto assumption and instead assume that the firm productivity distribution is lognormal. Section 5 gauges the impact of trade cost shocks using “exact hat algebra.” Section 6 concludes.

2. The Intensive Margin in the Data

The Exporter Dynamics Database

We use the EDD described in [Fernandes et al. \(2016\)](#) to study the intensive and extensive margins of trade. The EDD is based on firm-level customs data covering the universe of export transactions provided by customs agencies from 59 countries (53 developing and 6 developed countries) that we complement with data for China.⁷ For each country, the raw firm-level customs data contains annual export flows (in current values) disaggregated by firm, destination and Harmonized System (HS) 6-digit product. Oil exports are excluded due to lack of accurate firm-level customs data for many of the oil-exporting countries. For most countries total non-oil exports in the EDD are close to total non-oil exports reported in COMTRADE/WITS. For the descriptive analysis in this section as well as for the regression and simu-

⁶Our paper is also related to [Adao, Costinot and Donaldson \(2017\)](#), who extend the exact-hat algebra approach to a setting with a variable trade elasticity due to variation in the elasticity of demand.

⁷China is not included in the publicly available EDD statistics due to confidentiality concerns.

lation work in the sections that follow, we focus on a *core sample* that consists of 50 countries (49 from the EDD and China) for which we have the firm-level data. However, for the motivating plots below we use an *extended sample* that includes 60 countries (59 from the EDD plus China). Both samples cover a subset of years between 2003 and 2013 — see Tables 1 and A1 in the Online Appendix.

We focus on EDD products in the manufacturing sector.⁸ We calculate variants of average exports per firm, number of exporting firms, and total exports at the origin-destination-year level or at the origin-product-destination-year level. The product disaggregations that we use are HS 2-digit for the extended sample and HS 2-digit, HS 4-digit, or HS 6-digit for the core sample.

Importance of the intensive margin

Let X_{ij} , N_{ij} and $x_{ij} \equiv X_{ij}/N_{ij}$ denote total exports, total number of exporting firms, and average exports per firm from country i to country j , respectively.⁹ In Figure 1 we plot the intensive margin ($\ln x_{ij}$) and extensive margin ($\ln N_{ij}$) vs. total exports ($\ln X_{ij}$) for the extended sample of countries. We restrict the sample to the origin-destination pairs with more than 100 exporting firms (i.e., ij pairs for which $N_{ij} > 100$) to reduce noise associated with country pairs with few exporting firms.¹⁰ All variables plotted are demeaned of origin-year and destination-year fixed effects. Each dot corresponds to $(\ln x_{ij}, \ln X_{ij})$ (Panel A) or $(\ln N_{ij}, \ln X_{ij})$ (Panel B). The lines can be ignored for now.

A key statistic that we use to summarize the pattern observed in Figure 1 is the *intensive margin elasticity* (IME), which is the slope of the (not shown) regression line in Panel A. In a given year, the IME can be obtained from an OLS regression of $\ln x_{ij}$ on $\ln X_{ij}$ with origin and destination fixed effects:

$$\ln x_{ij} = FE_i^o + FE_j^d + \alpha \ln X_{ij} + \varepsilon_{ij}. \quad (1)$$

The IME is the estimated regression coefficient

$$\hat{\alpha} = \frac{\text{cov}(\ln \tilde{x}_{ij}, \ln \tilde{X}_{ij})}{\text{var}(\ln \tilde{X}_{ij})}, \quad (2)$$

where we write $\ln \tilde{z}_{ij}$ to denote variable $\ln z_{ij}$ demeaned by origin-year and destination-year fixed effects. The complement of the IME is the *extensive margin elasticity*, defined as $\text{EME} \equiv \frac{\text{cov}(\ln \tilde{N}_{ij}, \ln \tilde{X}_{ij})}{\text{var}(\ln \tilde{X}_{ij})}$. The EME is the slope of the (not shown) regression line in Panel B of Figure 1 and satisfies $\text{EME} = 1 - \text{IME}$.

Figure 1 demonstrates that both the IME and the EME are positive and large. As shown in Panel A

⁸Concording the ISIC rev. 3 and HS 6-digit classifications, we consider only exports of HS 6-digit products corresponding to ISIC manufacturing sub-sectors 15-37.

⁹While there is variation in our data over time, for simplicity, we suppress the time subscript in our variables.

¹⁰The core sample includes 1,305 unique country pairs with $N_{ij} > 100$ while the extended sample includes 2,087 unique country pairs with $N_{ij} > 100$. The total number of unique country-pairs is 8,401 in the core and 10,663 in the extended sample.

of Table 2, depending on the type of fixed effects included, the IME ranges from 0.4 to 0.46 in the core sample that we will use for the analysis in the next two sections. Our preferred estimate of the IME is 0.4 based on the inclusion of origin-year and destination-year fixed effects (as in Figure 1).¹¹ In this estimate, the intensive margin accounts for approximately 40% of the variation in total exports across country pairs, while 60% is accounted for by the extensive margin. As the focus has so far been on accounting for the variation in bilateral trade flows while controlling for origin-year and destination-year fixed effects, it is natural to wonder how much of that variation is absorbed by the fixed effects alone. The results in Table 2 show that this is never more than 59 percent, implying that a large share of the variation in bilateral trade flows comes from the forces behind the estimated IME.¹²

Robustness

The finding of a positive and large IME is robust to considering different samples, adding industry controls, dealing with measurement error, and looking only at variation over time. Here we provide an overview of the results, with details left to Online Appendix B.

First, we obtain the IME for (i) a sample including all country pairs and the estimate is 0.58 when origin-year and destination-year fixed effects are included (Panel B of Table 2); (ii) the extended sample and the estimate is 0.38; (iii) a sample that excludes firms whose annual exports fell below \$1,000 in any year to ensure the IME is not driven by small exporting firms and the estimate is 0.4 for the core sample and 0.38 for the extended sample; (iv) each year from 2003 to 2013 separately and the estimates range from 0.55 to 0.60; and (v) separately for higher-income and lower-income countries or for different continents and the estimates are similar.¹³

Second, we address the possibility that the IME estimates could be affected by country differences in industry composition of exports combined with industry differences in average exports per firm. The IME actually increases when moving to industry-level data: at the lowest level of aggregation (HS 6-digit), the IME is 0.51 for the core sample with origin-year-industry and destination-year-industry fixed effects, and it is 0.52 for the extended sample at the HS 2-digit level.¹⁴

Third, measurement error could be a concern for our IME estimates. Since total exports is the sum of firm-level exports, classical measurement error in exports per firm x would bias the IME upward, but classical measurement error in the number of firms N would bias the IME downward. If measurement

¹¹We estimate $\ln x_{ijt} = FE_{it}^o + FE_{jt}^d + \alpha \ln X_{ijt} + \varepsilon_{ijt}$ using all available years of data for the core sample country pairs.

¹²Based on the R^2 of OLS regressions of log bilateral total exports ($\ln X_{ijt}$) on origin-year and destination-year fixed effects.

¹³The IMEs in (iii) are obtained for origin-destination pairs with at least 100 exporting firms and with origin-year and destination-year fixed effects. The corresponding IMEs for all country pairs are 0.57 using the core sample and 0.51 using the extended sample.

¹⁴The presence of large trading firms could increase both exports per firm and total exports and explain our IME estimates. While we are unable to identify large trading firms in the EDD data, we estimate the IME for a sample including only HS 2-digit industries with low shares of firms exporting via intermediaries, as defined in Chan (2019). The results barely change.

error is serially uncorrelated, then instrumenting total exports with its leads and/or lags should yield an unbiased estimate of the IME. The instrumented IMEs are very close to the OLS IMEs.

Finally, as an alternative to the use of cross-sectional variation in bilateral trade flows to estimate the IME, we can exploit only time-series variation in bilateral export flows. The results from regressions that include origin-destination fixed effects or regressions in first-differences (where the IME is identified only off the panel dimension) for the core and the extended sample show significantly larger IMEs (around 0.85) than those obtained exploiting cross-sectional variation in Table 2. This time series evidence shows clearly that the intensive margin plays an important role for changes in bilateral export flows (see Online Appendix B).

IME by percentiles

A positive IME could be due to the presence of export superstars that increase both average exports per firm and total exports for some country pairs, as discussed in Freund and Pierola (2015). We study this possibility by considering separate IME regressions for each exporter size percentile. For each origin-destination-year combination we distribute the exporting firms into percentiles based on the value of their exports. Denoting average exports per firm in percentile pct as x_{ij}^{pct} , we run regressions based on the following specification:

$$\ln x_{ij}^{pct} = FE_i^o + FE_j^d + \alpha^{pct} \ln X_{ij} + \epsilon_{ij}.$$

We define the IME for each percentile as $IME^{pct} \equiv \hat{\alpha}^{pct}$.¹⁵ We plot the IME^{pct} for each percentile (with confidence intervals) in Figure 2 along with the horizontal line at the overall IME of 0.4. The IME is 0.5 for the highest percentile. But the positive overall IME is not coming exclusively from the export superstars: the IME^{pct} rises steadily from 0.2 at the 50th percentile to 0.3 at the 80th percentile.¹⁶ The IME by percentile is robust to country differences in industry composition of exports and industry differences in average exports per firm across percentiles, as shown in Online Appendix B.

IME for multi-product firms

We can dig deeper and study whether average exports per firm can be explained by the number of products exported per firm or by exports per product per firm.¹⁷ Let O_{ij} be the total number of firm-product observations with positive exports from i to j and let $x_{ij}^p \equiv X_{ij}/O_{ij}$ be the average exports per product per firm exporting from i to j . We define the IME at the product level as $IME^p \equiv cov(\ln \tilde{x}_{ij}^p, \ln \tilde{X}_{ij}) / var(\ln \tilde{X}_{ij})$.

¹⁵For exporter percentiles to be well-defined we focus on country pairs for which $N_{ij} > 100$.

¹⁶Only the top percentile, in which there is a large share of exports, lies above the overall IME.

¹⁷Details on the multi-product extension of the Melitz-Pareto model are given in Online Appendix D.

Let $m_{ij} = O_{ij}/N_{ij}$ be the average number of products per firm exporting from i to j . Then, the IME is equal to the IME^P plus the extensive product margin elasticity,

$$\text{IME} = \text{IME}^P + \frac{\text{cov}(\ln \tilde{m}_{ij}, \ln \tilde{X}_{ij})}{\text{var}(\ln \tilde{X}_{ij})}.$$

Results from a regression of log average exports per product on log total exports including origin-year and destination-year fixed effects reveal an IME^P of 0.29 for the core sample as shown in Online Appendix B. This implies that most of the IME is explained by the systematic variation in average exports per product per firm, rather than in the average number of products exported.¹⁸

Taking stock: the IME in the EDD

Summarizing the results so far, we find the intensive margin elasticity to be positive and significant, both statistically and economically. This finding is robust to the inclusion of a variety of fixed effects, various samples, exclusion of small firms, and disaggregation by industry. The IME is positive and monotonically increasing across the whole distribution of exporter size. The systematic cross-country-pair variation of average exports per firm comes primarily from the behavior of average exports per product per firm.

Correlation between the intensive and extensive margins, and their relation with distance

We now move beyond the intensive margin elasticities and report additional stylized facts on the correlations between the intensive margin, the extensive margin, and distance. There is a positive and significant correlation between average exports per firm and the number of exporting firms (0.25, standard error 0.01) after taking out origin-year and destination-year effects. Table 3 shows how these margins vary with log distance with alternative sets of fixed effects. The elasticities are all negative and significant when controlling for origin-year and destination-year fixed effects: average exports per firm, the number of firms, average number of products exported per firm, and average exports per product per firm all decline with distance between trade partners. Average exports per firm declines with distance even when disaggregated at the HS 2-digit, 4-digit, or 6-digit levels controlling for origin-year-industry and destination-year-industry fixed effects (see Online Appendix B).

Relation to previous empirical results

We finish this section by relating our stylized facts to those of EKK, EKS, Bernard et al. (2007) and Bernard et al. (2009). EKK use firm-level export data for a single origin (France) and show that average exports per

¹⁸Bernard et al. (2009) present a similar decomposition for U.S. exports. We compare their results to ours below.

firm increase with market size of the destination (measured as manufacturing absorption) with an elasticity of 1/3. In our case, a regression of average exports per firm on destination market size, including origin and year fixed effects reveals that average exports per firm increase with destination market size with an elasticity of 0.19, a bit lower than the result in EKK.

EKK also show that firms exporting to more destinations exhibit higher sales in the domestic (French) market. Our data does not include domestic sales, but we can look at sales in the most popular destination market for each origin. Let $x_{il|j}$ denote average exports to destination l computed across firms from i that sell in markets l and j and let $l^*(i) \equiv \operatorname{argmax}_k N_{ik}$ be the most popular destination market for each origin country i (e.g., the United States for Mexico). We regress $\log \frac{x_{il^*(i)|j}}{x_{il^*(i)|l^*(i)}}$ on $\log \frac{N_{ij}}{N_{il^*(i)}}$ for all i and j for the core sample with origin-year and destination-year fixed effects. The pattern shown in EKK for French firms extends to our data with many origin countries: firms that sell in more markets are more productive as proxied by their sales in their origin country's most popular destination market.¹⁹

EKS find that average exports per firm are similar across four origin countries (Brazil, Denmark, France and Uruguay). They regress average exports per firm on origin and destination fixed effects and find that the origin fixed effects differ little across their four origins. Running the same regression in our dataset (but pooling across years and including year fixed effects), we find that origin fixed effects do vary significantly across countries (the coefficient of variation in the estimated origin fixed effects ranges from 0.81 to 2.56, depending on the sample used) and are higher for countries with higher GDP per capita and higher total exports.²⁰ Moreover, origin-year and destination-year fixed effects are not enough to capture the variation in $\ln x_{ij}$: a regression of $\ln x_{ij}$ on origin-year and destination-year fixed effects yields an R-squared of 0.59 when only country pairs with $N_{ij} > 100$ are considered and only 0.5 when all country pairs are considered (see Table 2).

Using firm-level export data for the U.S., Bernard et al. (2009) present a similar decomposition to the one we present above for multi-product firms, except that they cannot allow for destination fixed effects because their data is for a single origin in a single year. They find that IME^p is around 0.23, which is not far from our estimate of 0.29. Using similar data, Bernard et al. (2012) show that average exports per product per firm increase with distance but the coefficient is insignificant. In contrast, as shown in the bottom half of Table 3, our regressions of $\ln x_{ij}^p$ on $\ln \text{dist}_{ij}$ with origin-year and destination-year fixed effects yield a negative and significant coefficient on distance. A negative coefficient on distance is also found when regressing $\ln x_{ij}$ on $\ln \text{dist}_{ij}$, as seen in the top half of Table 3.

¹⁹The EKK estimating sample includes only firms with sales in France. To implement an approach comparable to theirs, we drop all firms from country i that do not sell to $l^*(i)$, so the sample includes only $N_{il^*(i)}$ firms for country i . This implies that all firms that make up N_{ij} are also selling to $l^*(i)$. See Online Appendix F.

²⁰We regress the estimated origin fixed effects on population, GDP, GDP per capita, and total exports, jointly and separately.

3. The Intensive Margin in the Melitz Model

In this section we study the implications of the Melitz model for the intensive margin of trade. We find that if firm-level productivity is drawn from a Pareto distribution then the model generates predictions that are odds with the facts presented in the previous section. Allowing for multi-product firms or granularity does not help the Melitz-Pareto model better match the intensive margin facts, while moving from a Pareto to a lognormal distribution does.

3.1. Preliminaries

As this is a well-known model, we will be brief in the presentation of the main assumptions. There are many countries indexed by i, j . Labor is the only factor of production available in fixed supply L_i in country i and the wage is w_i . Preferences across varieties are constant elasticity of substitution (CES) with elasticity of substitution across varieties $\sigma > 1$. In each country i there is a large pool of firms of measure N_i , each producing a single variety sold under monopolistic competition with productivity φ distributed according to cumulative distribution function (CDF) $G_i(\varphi)$ and probability density function (PDF) $g_i(\varphi)$. For convenience we will assume that $g_i(\varphi) > 0$ for all φ , so that $G_i(\varphi)$ is everywhere increasing and hence invertible. Firms from country i incur in fixed trade costs F_{ij} (in units of the numeraire) and iceberg trade costs τ_{ij} to sell in country j . We do not need here to close the model and characterize the full equilibrium. Instead, we derive a few equilibrium relationships that will be useful to understand the model's implications for the intensive margin of trade.

Sales in destination j by a firm from origin i with productivity φ are

$$x_{ij}(\varphi) = \left(\bar{\sigma} \frac{w_i \tau_{ij}}{\varphi} \right)^{1-\sigma} P_j^{1-\sigma} w_j L_j, \quad (3)$$

where

$$P_j = \left(\sum_i N_i \int_{\varphi \geq \varphi_{ij}^*} \left(\bar{\sigma} \frac{w_i \tau_{ij}}{\varphi} \right)^{1-\sigma} dG_i(\varphi) \right)^{1/(1-\sigma)}$$

is the price index in j , $\bar{\sigma} \equiv \sigma / (\sigma - 1)$ is the markup, and φ_{ij}^* is the productivity cutoff for exports from i to j , which is determined implicitly by

$$x_{ij}(\varphi_{ij}^*) = \sigma F_{ij}. \quad (4)$$

The value of exports and the measure of firms in i that export to j are then as follows:

$$X_{ij} = N_i \int_{\varphi \geq \varphi_{ij}^*} x_{ij}(\varphi) dG_i(\varphi) \quad (5)$$

$$N_{ij} = N_i \left(1 - G_i(\varphi_{ij}^*)\right), \quad (6)$$

respectively. We will use $n_{ij} \equiv N_{ij}/N_j$ to denote the share of firms in i that export to j and $x_{ij} \equiv X_{ij}/N_{ij}$ to denote the associated average exports per firm.

The ratio of average to minimum exports per firm for each country pair can be written as

$$\frac{x_{ij}}{x_{ij}(\varphi_{ij}^*)} = \left(\frac{\tilde{\varphi}_i(\varphi_{ij}^*)}{\varphi_{ij}^*} \right)^{\sigma-1}, \quad (7)$$

where $\tilde{\varphi}_i(\varphi^*)$ is an average productivity level defined in Melitz (2003) as

$$\tilde{\varphi}_i(\varphi^*) \equiv \left(\frac{1}{1 - G_i(\varphi^*)} \int_{\varphi^*}^{\infty} \varphi^{\sigma-1} g_i(\varphi) d\varphi \right)^{\frac{1}{\sigma-1}}.$$

Equations (4), (6) and (7) imply that

$$x_{ij} = \Omega_i(n_{ij}) \sigma F_{ij}, \quad (8)$$

where $\Omega_i(n)$ is a function defined as

$$\Omega_i(n) \equiv \left(\frac{\tilde{\varphi}_i(G_i^{-1}(1-n))}{G_i^{-1}(1-n)} \right)^{\sigma-1}. \quad (9)$$

We make two observations here that will be important below. First, the function $\Omega_i(n)$ is completely determined by σ and the CDF $G_i(\cdot)$. Second, given that $G_i^{-1}(1-n)$ is increasing, if $\frac{\tilde{\varphi}_i(\varphi^*)}{\varphi^*}$ is decreasing in φ^* then $\Omega_i(n)$ would be an increasing function.

3.2. The Intensive Margin

Without loss of generality, we can write variable and fixed trade costs as $\tau_{ij} = \tau_i^o \tau_j^d \tilde{\tau}_{ij}$ and $F_{ij} = F_i^o F_j^d \tilde{F}_{ij}$. Taking logs on both sides of equation (8) then yields

$$\ln x_{ij} = \ln(\sigma F_i^o) + \ln F_j^d + \ln \Omega_i(n_{ij}) + \ln \tilde{F}_{ij}. \quad (10)$$

Similarly, from equations (3), (4), (6) and (8) we get

$$\ln x_{ij} = \ln(\bar{\sigma} w_i / \tau^o)^{1-\sigma} + \ln(P_j^{1-\sigma} w_j L_j / \tau^d)^{1-\sigma} + \ln\left(\Omega_i(n_{ij}) (G_i^{-1}(1-n_{ij}))^{\sigma-1}\right) + \ln \tilde{\tau}_{ij}^{1-\sigma}. \quad (11)$$

Imagine for now that $\text{var}(\ln \tilde{F}_{ij}) = 0$, so that all variation in fixed trade costs comes from origin and destination fixed effects with no country-pair component, for example because $F_{ij} \propto w_i^\gamma w_j^{1-\gamma}$, as in Arkolakis (2010). If firm productivity is distributed Pareto, $G_i(\varphi) = 1 - (\varphi/b_i)^{-\theta}$, with $\theta > \sigma - 1$ and

$b_i \leq \varphi_{ij}^*$ for all i, j , then $\left(\frac{\tilde{\varphi}(\varphi_{ij}^*)}{\varphi_{ij}^*}\right)^{\sigma-1} = \frac{\bar{\theta}}{\bar{\theta}-1}$ for all φ_{ij}^* , where $\bar{\theta} \equiv \theta/(\sigma-1)$. This implies that $\Omega_i(n)$ does not vary with n and hence $\text{IME} = 0$ while $\text{EME} = 1$. This is captured in Figure 1 by the horizontal line for the model-implied intensive margin (panel A) and the line with unit slope for the model-implied extensive margin (panel B). These implications of the model stand in sharp contrast to what is seen in the data, which reveals an IME of 0.4 or higher (see Figure 1 and Table 2).

One could certainly abandon the assumption that $\text{var}(\ln \tilde{F}_{ij}) = 0$, while retaining the Pareto assumption and allowing fixed trade costs to vary so as to match the findings in Section 2. However, the required variation in fixed costs seems unreasonable. As discussed in the previous section, average exports per firm decline with distance (see Table 3), so equation (10) implies that \tilde{F}_{ij} would need to systematically *fall* with distance. Setting $\theta = 5$ from Head and Mayer (2014) and $\sigma = 5$ from Bas et al. (2017), we can use equations (10) and (11) to compute the model-implied \tilde{F}_{ij} and $\tilde{\tau}_{ij}$ and relate these to distance. Table 4 shows the implied elasticity of \tilde{F}_{ij} with respect to distance would be -0.28, while the corresponding elasticity for $\tilde{\tau}_{ij}$ would be 0.27. As suggested by these results, \tilde{F}_{ij} would need to be *negatively* correlated with $\tilde{\tau}_{ij}$ to generate a positive IME, as shown formally in Online Appendix C. Intuitively, a positive IME implies a positive covariance between average and total sales (disregarding origin and destination fixed effects). Higher average sales go along with a higher \tilde{F}_{ij} and so a positive IME would need a positive covariance between \tilde{F}_{ij} and total sales, which would require a negative covariance between \tilde{F}_{ij} and $\tilde{\tau}_{ij}$. To the best of our knowledge, there are no models that would microfound a negative correlation between fixed and variable trade costs while also generating a positive correlation between fixed trade costs and aggregate trade flows (as we see when we project on distance).²¹ Finally, variation in \tilde{F}_{ij} does not prevent another implication of the Pareto distribution that is at odds with the findings of Section 2, namely that the IME calculated separately for each exporter size percentile would be the same as the overall IME: $\text{IME}^{\text{pct}} = \text{IME}$, for all pct, in contrast to what is shown in Figure 2.²²

These sharp and contrafactual implications of the Melitz-Pareto model all come from the fact that a Pareto distribution implies that $\tilde{\varphi}(\varphi_{ij}^*)/\varphi_{ij}^*$ does not vary with φ_{ij}^* . In contrast, as argued in footnote 15 of

²¹Allowing for tariffs in addition to iceberg trade costs would naturally lead to a *positive* correlation between model-implied variable and fixed trade costs. This is because a tariff affects trade flows both by increasing the price of the affected good, as with iceberg trade costs, and by decreasing the net profits conditional on the quantity sold, as with fixed trade costs. See Costinot and Rodríguez-Clare (2014), Felbermayr, Jung and Larch (2015), and Caliendo, Feenstra, Romalis and Taylor (2015). Alternatively, one may consider models with endogenous transportation costs. In Anderson and Yotov (2020), capital specific to transportation for each country pair implies that bilateral trade costs are increasing in bilateral trade flows (although this effect would be weaker in the long run when capital can be adjusted to demand). By increasing trade flows, low fixed trade costs could then lead to higher variable trade costs. A similar result may arise in Brancaccio, Kalouptsi and Papageorgiou (2020), where markups charged by shippers to exporters may increase with the number of exporters along any route. Contrary to our results, however, these mechanisms would imply a positive correlation between fixed trade costs and trade flows.

²²Exports of a firm in the p^{th} percentile of the exporter size distribution are $\sigma F_{ij} \left(\varphi^p / \varphi_{ij}^* \right)^{\sigma-1}$, where φ^p is such that $\Pr \left[\varphi < \varphi^p | \varphi > \varphi_{ij}^* \right] = p$. Since productivity is distributed Pareto, the ratio $\varphi^p / \varphi_{ij}^*$ and thus average exports per firm in each percentile should be the same for all ij pairs.

Melitz (2003), $\tilde{\varphi}(\varphi_{ij}^*)/\varphi_{ij}^*$ is decreasing in φ_{ij}^* if the distribution $g_i(\varphi)$ “belongs to one of several common families of distributions: lognormal, exponential, gamma, Weibul, or truncations on $(0, +\infty)$ of the normal, logistic, extreme value, or Laplace distributions. (A sufficient condition is that $g_i(\varphi)\varphi/(1 - G_i(\varphi))$ be increasing to infinity on $(0, +\infty)$.)” To understand the implication of this property, consider a decline in τ_{ij} , so that φ_{ij}^* decreases with no effect on minimum sales (which remain equal to σF_{ij}). The decline in τ_{ij} leads to an increase in exports of incumbent firms (which increases average exports per firm) and entry of low productivity firms (which decreases average exports per firm). Under Pareto these two effects exactly offset each other so there is no change in average exports per firm. In contrast, if productivity is distributed in such a way that $\frac{\tilde{\varphi}(\varphi_{ij}^*)}{\varphi_{ij}^*}$ is decreasing then the second effect does not fully offset the first, and average exports per firm would increase with a decline in τ_{ij} . More directly, a decline in τ_{ij} leads to an increase in n_{ij} and since $\Omega_i(n)$ is increasing then average exports also increase, implying a positive correlation between total and average exports, and hence a positive IME even with $\text{var}(\ln \tilde{F}_{ij}) = 0$.

A particularly convenient distribution in the family highlighted by Melitz (2003) is the lognormal distribution. In the next section we consider a Melitz model with a lognormal distribution but extended to allow for additional dimensions of firm heterogeneity that are important to match the microdata. We then provide a rigorous estimation of the extended model and explore its implications for the IME. For now, however, we can offer a quick preview of how simply moving from Pareto to lognormal can significantly improve the fit of the Melitz model with the empirical patterns presented in Section 2.

Given values of the mean and standard deviation of productivity, $\mu_{\varphi,i}$ and $\bar{\sigma}_{\varphi}$, as well as a value of N_i for every country, we can use our data on N_{ij} to compute $n_{ij} = N_{ij}/N_i$ and $\Omega(n_{ij})$ for all country pairs and then explore the implications of the Melitz-lognormal model for the intensive margin, as well as the implied trade costs. We use the QQ-estimation proposed by Head et al. (2014) to obtain estimates of σ_{φ} and $\mu_{\varphi,i}$ for every i and Bento and Restuccia (2017) data to estimate a value for N_i for all the countries in our sample (see Online Appendix H for a detailed description).^{23,24} For our estimate of the shape parameter, $\bar{\sigma}_{\varphi} = 4.02$, even with $\text{var}(\ln \tilde{F}_{ij}) = 0$, the model’s implied IME is 0.28, which is not far from what we found in the previous section.²⁵

We can again use equations (10) and (11) to compute model-implied fixed and variable trade costs

²³Using census, surveys and registry data, Bento and Restuccia (2017) compiled a dataset with the number of manufacturing firms for a set of countries. Unfortunately, their sample has missing observations for a number of countries in the EDD. We impute missing values projecting the log number of firms on log population. There is a tight positive relationship between log number of firms in their dataset and log population with an elasticity of 0.945. We acknowledge slippage between theory and data in that we obviously do not have a measure of the entry level N_i , but (at best) only the number of existing firms, which in theory would correspond to $(1 - G_i(\varphi_{ii}^*))N_i$. We avoid this problem in the analysis of the next Section.

²⁴We compute three sets of QQ-estimates of $\bar{\sigma}_{\varphi}$: for the full sample, the largest 50% of firms and the largest 25% of firms for each origin-destination pair in each year. These estimates are higher than the estimate obtained by Head et al. (2014), so we use the minimum among them, $\bar{\sigma}_{\varphi} = 4.02$, which corresponds to the subsample with the largest 25% of firms. See Online Appendix H for a discussion of these estimates and their relation to the estimate in Head et al. (2014).

²⁵If we instead use the estimate in Head et al. (2014) of $\bar{\sigma}_{\varphi} = 2.4$ then we get an IME of around 0.12.

but now under the assumption that productivity is distributed lognormal. The correlations between the implied values of \tilde{F}_{ij} and $\tilde{\tau}_{ij}^{\sigma-1}$ and distance are reported in Table 4. In contrast to our results with Pareto, under lognormal both the model-implied variable and fixed trade costs are increasing with distance, with elasticities of 0.3 and 0.15, respectively.

Finally, we can explore the implications of the Melitz-lognormal model for the IME by percentiles. As shown in Figure 3, and in line with our findings in Section 2, the IME is positive and increasing across percentiles, with most but not all the action at the highest percentiles.

3.3. Multi-Product Firms and Granularity

In this section we discuss whether maintaining the Pareto assumption but moving beyond the Melitz model in other ways can improve its fit with the findings for the intensive margin of trade in Section 2. Specifically, we consider two extensions of the Melitz-Pareto model: multi-product firms and granularity.

With multi-product firms as in Bernard et al. (2011), average exports per firm may fall along with total exports (thereby creating a positive IME) as firms facing higher product-level fixed trade costs export fewer products (even though they export more per product). Roughly speaking, allowing for multi-product firms implies that part of the extensive margin in the basic Melitz-Pareto model now operates inside the firm and appears as part of the firm intensive margin in the data. As shown in Online Appendix D, however, under the Pareto assumption the effect of higher product-level fixed trade costs on the number of products exported per firm is exactly offset by higher average exports per product, and so the contrafactual results described above remain valid in this extension.

Dropping the assumption of a continuum of firms and allowing for granularity, as in Eaton et al. (2012) or Gaubert and Itskhoki (2021), may generate a positive covariance between average exports and total exports that could in theory explain our empirical findings for the IME. Intuitively, large exports by a particular firm from country i to market j could lead to both high average exports and high total exports from i to j . We explored this formally using the extension of the Melitz-Pareto model to allow for granularity in Eaton et al. (2012), following two approaches described in Online Appendix E. First, we estimate the elasticity of model-implied fixed trade costs with respect to distance taking into account granularity. Although the distance elasticities are significantly lower than those estimated ignoring granularity, they remain negative, so the fixed trade costs implied by the Melitz-Pareto model are still decreasing with distance. Second, we simulate exports of N_{ij} firms for each of the country pairs in the sample and then compute the implied IME under $\text{var}(\tilde{F}_{ij}) = 0$. Consistent with the intuition above, now the IME is positive but –under plausible values for the key parameters– an order of magnitude lower than that reported in Section 2. Moreover, even if we push parameters to extreme values to get a realistic IME, all the action explaining the positive IME would come from the superstar firms (with the IME being close to zero for

small percentiles), a result contrary to Figure 2.

4. The Intensive Margin in an Extended Melitz-Lognormal Model

In Section 3 we found that the Melitz model does better in matching the facts presented in Section 2 when we assume that firm productivity is distributed lognormal than when we assume that it is distributed Pareto. Encouraged by those results, in this section we first present a Melitz model with lognormally distributed firm productivity, destination-specific fixed costs and demand shocks.²⁶ We then describe a maximum-likelihood approach to estimate the model using our firm-level microdata and finally study the implications of the estimated model for the IME as well as for the model-implied trade costs.

4.1. Model

Our extended Melitz model is similar to that in Eaton et al. (2011) in that it allows for firm-specific fixed trade costs and demand shocks that vary by destination. The main difference is that we assume that firm productivity φ , demand shocks, α , and fixed trade costs, f , are distributed jointly lognormal. A firm from origin i with productivity φ , demand shocks $(\alpha_1, \dots, \alpha_J)$ and fixed trade costs (f_1, \dots, f_J) would have net profits from selling in market j given by $\alpha_j D_{ij} \varphi^{\sigma-1} / \sigma - f_j$, where $D_{ij} \equiv (\bar{\sigma} w_i \tau_{ij})^{1-\sigma} P_j^{\sigma-1} w_j L_j$.

Without loss of generality, we allow mean log productivity to be origin-specific while imposing that the mean of demand shocks be the same across origin-destination pairs (we cannot separately identify these parameters). Mean fixed trade costs are allowed to vary across origin-destination pairs and can be correlated with demand shocks within destinations. In line with these assumptions, we let $\mu_{\varphi,i}$, μ_{α} , $\mu_{f,ij}$ denote log averages of productivity, demand shocks, and fixed trade costs for firms from origin i selling to destination j . In turn, we allow the dispersion of log productivity, log demand shocks and log fixed trade costs to differ across origins, but assume that – for a given origin – the dispersion of log demand shocks and log fixed trade costs do not vary across destinations. Thus, we let $\sigma_{\varphi,i}^2$, $\sigma_{\alpha,i}^2$, $\sigma_{f,i}^2$, $\sigma_{\alpha f,i}$ denote the variance of log productivity, demand shocks, fixed trade cost and covariance between demand shocks and fixed trade costs for a firm from i . To further clarify these assumptions, it is useful to consider the case of firms from i with only two destinations labeled 1 and 2.²⁷ The joint distribution of productivity,

²⁶Eaton et al. (2011) also allow for lognormally distributed fixed cost and demand shocks, but retain the Pareto assumption for productivity, and so, as explained in Online Appendix G, the behaviour of the intensive margin is the same as in the Melitz-Pareto model. Nigai (2017) also combines Pareto with lognormal, but assuming that productivity is lognormal for most of the distribution and then becomes Pareto in the right tail. We used Nigai's Matlab code on our data to estimate the point of truncation (percentile) where the lognormal ends and the Pareto begins: for 75% of country pairs with more than 100 exporters the truncation point occurs after the 99th percentile and for the median country pair the truncation point is at the 99.9th percentile. In light of these results, in the rest of the paper we consider a fully lognormal distribution for productivity.

²⁷The general formulation of the joint probability distribution is reported in Online Appendix I.

demand shocks, and fixed trade costs in this case would be

$$\begin{bmatrix} \ln \varphi \\ \ln \alpha_1 \\ \ln \alpha_2 \\ \ln f_1 \\ \ln f_2 \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \mu_{\varphi,i} \\ \mu_{\alpha} \\ \mu_{\alpha} \\ \mu_{f,i1} \\ \mu_{f,i2} \end{bmatrix}, \begin{bmatrix} \sigma_{\varphi,i}^2 & 0 & 0 & 0 & 0 \\ 0 & \sigma_{\alpha,i}^2 & 0 & \sigma_{\alpha f,i} & 0 \\ 0 & 0 & \sigma_{\alpha,i}^2 & 0 & \sigma_{\alpha f,i} \\ 0 & \sigma_{\alpha f,i} & 0 & \sigma_{f,i}^2 & 0 \\ 0 & 0 & \sigma_{\alpha f,i} & 0 & \sigma_{f,i}^2 \end{bmatrix} \right). \quad (12)$$

Without risk of confusion, we change notation in this section and use $X_i \equiv (X_{i1}, \dots, X_{iJ})$ to denote the random variable representing log sales of a firm from i in each of the J destinations, with $x_i \equiv (x_{i1}, \dots, x_{iJ})$ being a realization of X_i , and with $g_{X_i}(x_i)$ being the associated probability density function. According to the model, a firm does not export to destination j if it has a large fixed trade cost draw f_j relative to its productivity and its demand shock for that destination. Let $Z_{ij} \equiv \ln D_{ij} + \ln \alpha_j + (\sigma - 1) \ln \varphi$ be sales in destination j by a firm from i with productivity φ and demand shock α_j . This is a latent variable that we observe only if a firm actually exports,

$$X_{ij} = \begin{cases} Z_{ij} & \text{if } \ln \sigma + \ln f_{ij} \leq Z_{ij} \\ \emptyset & \text{otherwise} \end{cases},$$

with $Z_i \equiv (Z_{i1}, \dots, Z_{iJ})$ distributed according to

$$\begin{bmatrix} Z_{i1} \\ \vdots \\ Z_{iJ} \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} d_{i1} \\ \vdots \\ d_{iJ} \end{bmatrix}, \begin{bmatrix} \bar{\sigma}_{\varphi,i}^2 + \sigma_{\alpha,i}^2 & \cdots & \bar{\sigma}_{\varphi,i}^2 \\ \vdots & \ddots & \vdots \\ \bar{\sigma}_{\varphi,i}^2 & \cdots & \bar{\sigma}_{\varphi,i}^2 + \sigma_{\alpha,i}^2 \end{bmatrix} \right), \quad (13)$$

where $d_{ij} \equiv \ln D_{ij} + \mu_{\alpha} + (\sigma - 1) \mu_{\varphi,i}$ and $\bar{\sigma}_{\varphi,i} \equiv (\sigma - 1) \sigma_{\varphi,i}$.

4.2. Estimation Approach

Using firm-level data from the EDD and China, we can estimate the parameters in (13) as well as mean log fixed trade costs (up to a constant) and their dispersion using maximum likelihood methods. Online Appendix J shows how to derive the density function $g_{X_{i1}, \dots, X_{iJ}}(x_{i1}, \dots, x_{iJ})$ for the case when we observe sales to J destinations. We simplify the analysis by considering only data for 15 destinations (USA, Germany, Japan, France and the 11 largest destinations by export value for each origin), which we label

$j = 1, \dots, 15$ for year 2007 for each of 39 origins. We compute $g_{X_{i1}, \dots, X_{iJ}}(x_{i1}, \dots, x_{iJ})$ for each observation in our dataset (which is a realization of $\{X_{i1}, \dots, X_{iJ}\}$ that we observe). Since all random variables are independent across firms, we can compute the log-likelihood function as a sum of log-densities,

$$\ln L(\Theta_i | \{x_{i1}(k_i), \dots, x_{iJ}(k_i)\}_{i, k_i}) = \sum_{k_i=1}^{\tilde{N}_i} \ln \left[g_{(X_{i1}, \dots, X_{iJ})}(x_{i1}(k_i), \dots, x_{iJ}(k_i)) \right], \quad (14)$$

where \tilde{N}_i is the number of firms from i that sell to either of the 15 destinations we consider and k_i is an index for a particular observation in our dataset (for origin i it takes values in $1, \dots, N_i$) and Θ_i is an origin-specific vector of parameters that we want to estimate,

$$\Theta_i = \left\{ \left\{ d_{ij}, \bar{\mu}_{f,ij} \right\}_{i,j}, \bar{\sigma}_{\varphi,i}, \sigma_{\alpha,i}, \sigma_{f,i}, \rho_i \right\}, \quad (15)$$

where $\bar{\mu}_{f,ij} \equiv \ln \sigma + \mu_{f,ij}$ and $\rho \equiv \frac{\sigma_{\alpha f, i}}{\sigma_{\alpha, i} \sigma_{f, i}}$. As the likelihood is potentially not concave in θ_i and because there are 34 parameters to estimate per origin, we rely on the estimation methodology proposed by [Chernozhukov and Hong \(2003\)](#).²⁸ We use the Metropolis-Hastings MCMC algorithm to construct a chain of estimates $\Theta_i^{(n)}$ for each origin country. [Chernozhukov and Hong \(2003\)](#) show that $\bar{\Theta} \equiv \frac{1}{N} \sum_{n=1}^N \theta_i^{(n)}$ is a consistent estimator of Θ_i , while the covariance matrix of $\bar{\Theta}_i$ is given by the variance of $\Theta_i^{(n)}$, so we use this to construct confidence intervals for $\bar{\theta}_i$. For each origin, we run 5 different chains that start at a different random starting value $\theta_i^{(i)}$. We then explore whether the different parameters in θ_i converged to the same values across different chains and discuss the convergence of the chains in Online Appendix L.

Loosely speaking, identification works as follows. First, data on export flows and the number of exporters across country pairs helps in identifying d_{ij} and $\bar{\mu}_{f,ij}$. Second, the variance of firm sales within each ij pair helps in identifying the sum of the dispersion parameters for productivity and demand shocks, $\bar{\sigma}_{\varphi,i} + \sigma_{\alpha,i}$. Third, the extent of correlation of firm sales from a particular origin across different destinations helps in identifying $\sigma_{\varphi,i}$ separately from $\sigma_{\alpha,i}$: the more correlated firm sales are across destinations, the larger is $\sigma_{\varphi,i}$ relative to $\sigma_{\alpha,i}$. Fourth, the correlation between fixed costs and demand shocks can be inferred from the distribution of sales of small firms. Intuitively, if the correlation is negative, then a firm with a bad demand shock would likely draw a high fixed trade cost and thus would not export, hence, we would not see many small firms in the data. Finally, to understand how $\sigma_{f,i}$ is

²⁸This methodology uses the likelihood function as a probability distribution of the set of parameters to be estimated, and then via simulation finds the expectation of this distribution. This expectation is used as the estimator for the parameters. Note that this is not maximum likelihood estimation, since we are not selecting the point where the density is maximized. A detailed description of the estimation procedure can be found in Online Appendix K.

identified, imagine for simplicity that there is only one destination. We then have

$$g_{X_{i1}}(x_{i1}) = \frac{g_{Z_{i1}}(x_{i1}) \times \Pr\{\ln \sigma + \ln f_{i1} \leq x_{i1} | Z_{i1} = x_{i1}\}}{C_i}$$

where $C_i \equiv \Pr\{\ln \sigma + \ln f_{i1} \leq Z_{i1}\}$ and $g_{Z_{i1}}(\cdot)$ is the probability density function of the latent sales Z_{i1} . This implies that we can get the density of X_{i1} by applying weights $\frac{\Pr\{\ln \sigma + \ln f_{i1} \leq x_{i1} | Z_{i1} = x_{i1}\}}{C_i}$ to the density of Z_{i1} . The parameter $\sigma_{f,i}$ regulates how these weights behave with x_{i1} . In the extreme case in which $\sigma_{f,i} = 0$ then the weights are 0 for $x_{i1} \leq \mu_{f_{i1}}$ and $1/C$ for $x_{i1} > \mu_{f_{i1}}$, while in the other extreme with $\sigma_{f,i} = \infty$ the weights are all equal to 1. For intermediate cases the density of X_{i1} will be somewhere in the middle, with the left tail becoming fatter and the right tail becoming thinner as $\sigma_{f,i}$ increases. This suggests that we can identify $\sigma_{f,i}$ from the shape of the density of sales.

We will use the results of the estimation to conduct exercises similar to those in the previous sections. First, we will compute the IME for all firms and for each percentile using the estimated model. Second, after removing origin and destination fixed effects, we will compute the correlation across the estimated values of d_{ij} and $\bar{\mu}_{f,ij}$, and between them and distance.

4.3. Estimation Results

To estimate the parameters of the full Melitz-lognormal model we use firm-level data from the EDD and China for year 2007 for 37 origins (of 39 origins possible 2 were dropped due to convergence issues discussed in the Online Appendix L).²⁹ We use $\sigma = 5$ based on the estimates of [Bas et al. \(2017\)](#).

Figure 4 reveals the goodness of fit of the estimated model relative to the data. Panel a plots the density function for standardized firm-level log sales pooled across multiple origins and destinations.³⁰ The model generates a distribution that closely fits the one in the data. We next look at deviations from the strict hierarchy of firms sales across destinations (for each origin) in the data and in the estimated model. If there were no demand and fixed cost shocks across firms, then all firms from a given origin that export to less popular destinations would also export to the most popular destination. The share of firms that only sell in the less popular destinations is then a measure of the extent to which this strict hierarchy predicted by the simplest model is violated. Panel B shows that the share predicted by the estimated model is quite close to the one in the data. Finally, for each origin and any two destinations among the three most popular ones, Panel C shows the correlation in export value across all firms that sell in those two destinations. The estimated model mostly implies a positive correlation driven by firm-level productivity shocks, while in the data this correlation exhibits more dispersion.

Table 5 shows the estimates of the variance-covariance parameters $(\bar{\sigma}_{\varphi,i}, \sigma_{\alpha,i}, \sigma_{f,i}, \rho_i)$. The median

²⁹For computational reasons, for China we considered only a random sample consisting of 5% of exporters.

³⁰Standardized firm-level log sales for each origin-destination cell subtract the mean and divide by the standard deviation.

estimated values for $\bar{\sigma}_{\varphi,i}$ and $\sigma_{\alpha,i}$ across 37 origins are 3.18 and 2.67, respectively, and for $\sigma_{f,i}$ and ρ_i are 2.39 and 0.50. Even though the variance-covariance parameters were precisely estimated for each of the origins, the parameters vary quite a bit across different origins.³¹ In general, there is a positive correlation between demand and fixed costs shocks, but some origins exhibit a negative correlation.

Table 6 and Panel D of Figure 4 show the implications of the estimated model for the IME. We compute the IME implied by the estimated model by drawing one million firms for each origin (this implies one million latent log sales and log fixed costs for each destination), computing average sales (taking into account selection), and then multiplying average sales by N_{ij} in the data to compute total exports.³² The IME implied by the model is 0.63. This is actually higher than our preferred IME estimate of 0.4 in Section 2, but the gap comes in large part from the different sample of origin-destination pairs used here. Using the same sample of 37 origins and 4 destinations for year 2007 we estimate an IME of 0.67 (with a standard error of 0.03) that is statistically indistinguishable from the one implied by our estimated lognormal model.³³ The associated IME for each percentile is plotted in Panel D – the pattern of the IME across percentiles is remarkably close to what we see in the data.

Table 7 shows the elasticity of estimated variable and fixed trade costs with respect to distance (controlling for origin and destination fixed effects). Now both types of trade costs are strongly increasing in distance. Surprisingly, however, we still get a negative correlation between fixed and variable trade costs.

Overall, our estimated full Melitz-lognormal model does a very good job in fitting the EDD data and in solving the puzzles associated with the Pareto model. The lognormal model generates an IME that is close to the one we see in the EDD and implies fixed trade costs that increase with distance. The implied pattern for the IME across different percentiles is also very similar to what we see in the data.

We also estimated the full model with Pareto-distributed productivity with the CDF given by $\Pr[\varphi_i \leq \varphi] = 1 - \left(\frac{\varphi}{b_i}\right)^{-\theta_i}$, $\forall \varphi_i \geq b_i$. This model is similar to Eaton et al. (2011) without the requirement that $b_i \rightarrow 0$. To have finite price indices, $\frac{\theta_i}{\sigma-1} > 1$ should hold, and thus we imposed the restriction $\frac{\theta_i}{\sigma-1} \in [1.05, \infty]$.³⁴ The rest of the model is similar to the Melitz-lognormal model in having lognormally-distributed demand shocks and fixed costs. We recomputed the IME implied by the estimated model, correlations of trade costs with distance, as well as goodness of fit measures, in the same way as for the full Melitz-lognormal model. The Melitz-Pareto model does a good job in several respects: it fits the standardized distribution of sales, it generates similar patterns of hierarchy and correlation of firm sales across destinations as the Melitz-lognormal model (Figure 5), it implies an IME that is close to the one in the data (Table 6), and it yields positive correlations between trade costs and distance (Table 7). The estimated Melitz-Pareto

³¹The estimates and confidence bands for each of the parameters are reported in Online Appendix L.

³²We pick one million as a numerical approximation to the case with a continuum of firms.

³³The confidence interval in Table 6 comes from 1,000 random realizations of the parameters in our Markov chains.

³⁴The unconstrained model yields very similar results, see Online Appendix M.

model cannot, however, reproduce the upward-sloping IME across percentiles that we see in the data. Panel D of Figure 5 shows that it implies a downward-sloping pattern of IME across percentiles, in contrast to the upward-sloping pattern in the data. Furthermore, without $b_i \rightarrow 0$, the Melitz-Pareto model loses the tractability that gives rise to the explicit aggregate expressions in Eaton et al. (2011).

5. Counterfactual Analysis

In this section we study whether the counterfactual implications of the Melitz-lognormal model estimated in the previous section differ from those of the standard Melitz-Pareto model. We start by presenting an extension of the “exact hat algebra” approach popularized by Dekle et al. (2008) to accommodate any distribution of productivity, demand, and fixed cost shocks in the Melitz model. We then use this approach to quantify how trade flows and welfare respond to changes in trade costs in the Melitz-lognormal model, and compare these responses to those in the standard Melitz-Pareto model.

To conduct counterfactual analysis, we need to close the model. We do so in standard fashion by assuming that labor is the only factor of production, with wage w_i and perfectly inelastic labor supply L_i in country i , by assuming that entry costs are in terms of labor, and that fixed exporting costs are in terms of labor of the exporting country. To make the model be perfectly consistent with the data, we allow for trade imbalances via exogenous international transfers, as in Dekle et al. (2008). Formally, letting $X_i = \sum_l X_{li}$ denote total sales by country i and $Y_i = \sum_j X_{ij}$ denote total expenditure, trade imbalances are equal to international transfers Δ_i – that is, $\Delta_i = X_i - Y_i$.

5.1. Exact Hat Algebra in the Generalized Melitz Model

Here we show how to extend the “exact hat algebra” for counterfactual analysis in Dekle et al. (2008) to the Melitz model with a general productivity distribution (not necessarily lognormal) and allowing for firm-level demand and fixed-cost shocks.

We start by introducing some notation. Let $\tilde{\varphi}_{ij} \equiv (\sigma - 1)(\ln \varphi_i - \mu_{\varphi,i}) + \ln \alpha - \mu_\alpha$ and $\tilde{f}_{ij} \equiv \ln f_{ij} - \mu_{f,ij}$ be mean-adjusted productivity and fixed costs, respectively, and let h_{ij} be an endogenous cutoff such that firms from i serve market j if and only if $\tilde{f}_{ij} - \tilde{\varphi}_{ij} \leq h_{ij}$. The price index in market j is then

$$P_j^{1-\sigma} = \sum_i P_{ij}^{1-\sigma}, \quad (16)$$

with

$$P_{ij}^{1-\sigma} = N_i (\bar{\sigma} w_i \tau_{ij})^{1-\sigma} e^{(\sigma-1)\mu_{\varphi,i} + \mu_\alpha} \int_{-\infty}^{\infty} e^{\tilde{\varphi}} \int_{-\infty}^{h_{ij} + \tilde{\varphi}} g_{ij}(\tilde{\varphi}, \tilde{f}) d\tilde{f} d\tilde{\varphi}, \quad (17)$$

where g_{ij} is the joint PDF of \tilde{f}_{ij} and $\tilde{\varphi}_{ij}$, trade shares are

$$\lambda_{ij} = \frac{P_{ij}^{1-\sigma}}{P_j^{1-\sigma}}, \quad (18)$$

and the share of firms from origin i that sell in destination j is given by

$$n_{ij} \equiv N_{ij}/N_i = \int_{-\infty}^{\infty} \int_{-\infty}^{h_{ij}+\tilde{\varphi}} g_{ij}(\tilde{\varphi}, \tilde{f}) d\tilde{f} d\tilde{\varphi}. \quad (19)$$

In equilibrium, the cutoff variable h_{ij} must be such that log profits derived in market j by a firm from i with productivity $\tilde{\varphi}_{ij}$ and fixed cost \tilde{f}_{ij} be equal to $h_{ij} + \tilde{\varphi}_{ij} - \tilde{f}_{ij}$. Thus, h_{ij} must satisfy:

$$h_{ij} = \ln \left(\frac{(\tilde{\sigma} w_i \tau_{ij})^{1-\sigma} P_j^{\sigma-1} X_j}{\sigma w_i} \right) + (\sigma - 1) \mu_{\varphi,i} + \mu_{\alpha} - \mu_{f,ij}. \quad (20)$$

In turn, free entry implies that profits net of fixed costs of exporting are equal to entry costs. As shown in Online Appendix I, this can be written as

$$\sigma F^e w_i N_i = \sum_j \lambda_{ij} X_j \left(1 - \frac{\int_{-\infty}^{\infty} \int_{-\infty}^{h_{ij}+\tilde{\varphi}} e^{\tilde{f}} g_{ij}(\tilde{\varphi}, \tilde{f}) d\tilde{f} d\tilde{\varphi}}{e^{h_{ij}} \int_{-\infty}^{\infty} e^{\tilde{\varphi}} \int_{-\infty}^{h_{ij}+\tilde{\varphi}} g_{ij}(\tilde{\varphi}, \tilde{f}) d\tilde{f} d\tilde{\varphi}} \right). \quad (21)$$

An equilibrium is defined as variables $\{h_{ij}, \lambda_{ij}, P_{ij}\}$ and $\{X_j, P_j, w_i\}$ such that equations (16) - (21) are satisfied for all i, j , and in addition for all i

$$w_i L_i = \sum_j \lambda_{ij} X_j \quad (22)$$

$$X_j = w_j L_j + \Delta_j. \quad (23)$$

We now consider the analogous system of equations in hat changes rather than in levels, with standard hat notation $\hat{x} = x'/x$, where we use primes to denote counterfactual values. To quantify the effect of changes in trade costs and trade imbalances, we take $\{\hat{\tau}_{ij}\}$ and $\{\hat{\Delta}_j\}$ as exogenous and solve for hat changes in endogenous variables, $\{\hat{h}_{ij}, \hat{\lambda}_{ij}, \hat{P}_{ij}\}$ and $\{\hat{X}_j, \hat{P}_j, \hat{w}_i\}$, given parameter σ , functions $\{g_{ij}\}$, and data $\{\lambda_{ij}, h_{ij}\}$ and $\{X_j, Y_i\}$. The corresponding system of equations is relegated to Online Appendix I. It is important to note that, in contrast to the hat-algebra in the Melitz-Pareto model, in the more general case considered here we also need data on n_{ij} so that we can compute the implied h_{ij} . And of course, instead of simply knowing the Pareto shape parameter (or trade elasticity), here we need to know the functions $\{g_{ij}\}$ for all ij pairs.

5.2. Counterfactual Analysis in the Estimated Full Melitz-Lognormal Model

For our counterfactual analysis we need a set of countries for which we have $\{X_j, Y_i\}$ as well as $\{\lambda_{ij}, h_{ij}\}$ and $\{\lambda_{ij}, h_{ij}\}$ for all i and j in that set. Since we assume that the variances of φ_i and f_{ij} differ by origin but not by destination (see Section 4), then $g_{ij} = g_i$ for all i and j . We have estimated g_i for a set of 37 EDD countries and we can infer the implied N_i for all those countries, so we can include any subset of those countries in our analysis.³⁵ We construct h_{ij} using data for N_i , N_{ij} and equation (19). Finally, we also need X_{ij} and N_{ij} for $i = j$. Following the approach proposed by Ossa (2015), we construct X_{ii} as manufacturing value-added in country i from the World Development Indicators divided by 0.25, which is close to the average share of manufacturing value-added in gross production from the World Input-Output Database (WIOD) for the set of covered EDD countries in 2007. We set $N_{ii} = N_i$, which would be true if there are no fixed costs for domestic sales.

We conduct our counterfactual analysis for a world composed of the 12 Latin American countries and China, for which we have estimated the full Melitz-lognormal model.³⁶ We do not consider the whole EDD dataset for computational reasons. Some of the country pairs in the whole EDD dataset trade very little and this would make our welfare calculations imprecise (since we need to use numerical approximation to compute some of the integrals).

To compare the counterfactual implications of the full Melitz-lognormal model with those of the Melitz-Pareto model, we need a value for the Pareto shape parameter, θ . Following ACR, we set this parameter equal to an estimate of the trade elasticity, which we obtain as follows. Our estimated parameters $\ln \hat{d}_{ij}$ are a sum of origin and destination specific components, a constant, and the term $(1 - \sigma) \ln \tau_{ij}$. Hence, we can combine these estimated values of $\ln \hat{d}_{ij}$ with actual trade flows X_{ij} to estimate the trade elasticity from the following regression:

$$\ln X_{ij} = \gamma_i^o + \gamma_j^d - \frac{\beta}{\sigma - 1} \ln \hat{d}_{ij} + \zeta_{ij}. \quad (24)$$

Using again $\sigma = 5$, this yields an estimate of the trade elasticity equal to $\hat{\beta} = 4.03$. In the Melitz-Pareto model, this implies that $\theta = 4.03$.

We consider four different trade costs shocks: 1%, 5%, 10%, and 25% uniform reductions in international trade costs – formally, $\hat{\tau}_{ij} = \hat{\tau} \in \{0.99, 0.95, 0.9, 0.75\}$ if $i \neq j$, while $\hat{\tau}_{ii} = 1$. For each trade cost shock $\hat{\tau}$ we compute the counterfactual implications in both the Melitz-lognormal and Melitz-Pareto models. We show the results of this exercise in Figures 6 and 7. We use $\hat{W}_i^m \equiv \hat{w}_i^m / \hat{p}_i^m$ and \hat{X}_{ij}^m to denote the

³⁵To see how we get N_i , note that the estimated model provides us with a probability that a random firm from some origin is selling to at least one of the 15 destinations we consider. As we observe the total number of exporters to those destinations, we can infer the total number of firms from which they are drawn. We provide details on this procedure in Online Appendix J.

³⁶We implicitly assigns trade flows between these countries and countries outside of this group to domestic transactions.

that changes in welfare and trade flows for the Melitz-lognormal model ($m = LN$) and the Melitz-Pareto model ($m = P$). Figure 6 plots $\hat{W}_i^{LN} - 1$ (horizontal axis) against $\hat{W}_i^P - 1$ (vertical axis) in response to the four different trade costs shocks. It is evident that both models yield very similar results. As is well known from Arkolakis et al. (2012) and Melitz and Redding (2015), the welfare effects of trade liberalization depend critically on the behavior of the trade elasticity, which is qualitatively different across the two models: while the trade elasticity in the Melitz-Pareto model is common across country pairs and invariant to shocks, this is no longer true in the Melitz-lognormal. We can use our estimated g_i and n_{ij} to compute the local trade elasticity in the Melitz-lognormal model for each country pair using the formula derived by Bas et al. (2017). The resulting elasticity ranges from 4 to 6.8 with a standard deviation of 0.57 and the higher values occurring for country pairs with a low n_{ij} , as shown in Online Appendix N. However, this variation in trade elasticities across country pairs matters little for the gains from a uniform decline in trade costs, because the larger gains obtained with partners for which the trade elasticity is higher are compensated by the lower gains with partners for which the trade elasticity is lower. Loosely speaking, for a uniform trade cost shock, what matters is the average trade elasticity, and so the Melitz-Pareto model yields a good approximation for the gains from uniform trade liberalization.

Even though gains from trade liberalization are similar in the two models, the Melitz-Pareto and the Melitz-lognormal models differ in their implications for the counterfactual changes in bilateral trade flows. Figure 7 plots the ratio of the difference between the counterfactual changes in trade flows in the Melitz-lognormal model (\hat{X}_{ij}^{LN}) and in the Melitz-Pareto model (\hat{X}_{ij}^P) against the trade elasticity implied by the Melitz-lognormal model. We can see that Melitz-Pareto model can significantly over- or under-predict changes in trade flows depending on the actual trade elasticity. Naturally, a higher trade elasticity in the Melitz-lognormal model leads to larger changes in trade flows relative to the Melitz-Pareto model.

What happens if trade liberalization is asymmetric? We consider an extreme case in which, for each origin, trade costs decrease only for exports to the destination with the largest number of exporters – formally, we consider 13 separate shocks, one for each Latin American country and China as an origin, with the shock for origin i being that $\hat{\tau}_{ij} = 0.25$ if $j = \arg\max_l N_{il}$ and $\hat{\tau}_{ij} = 1$ otherwise. Since the trade elasticity should be low for the affected pairs, we expect the Melitz-lognormal model to deliver smaller welfare gains than the Melitz-Pareto model. This is confirmed in Figure 8. However, the differences in welfare gains between the two models are small. Again, as in the analysis with symmetric trade cost declines, we see bigger differences across models in the effects on trade flows, as shown by Figure 9.

Finally, it is interesting to compare our results to those in Melitz and Redding (2015). They find that the ACR ex-post formula for welfare evaluation does a poor job of capturing the true welfare changes from a decline in trade costs in a symmetric Melitz model with a truncated Pareto distribution. In contrast, we find that the ACR formula does a good job in approximating welfare changes in the estimated

Melitz-lognormal model. The difference comes from how much the trade elasticity varies in the two models: whereas it falls from 15 to 5 as trade costs decline in the Melitz-Redding exercise, the trade elasticity shows little variation in our Melitz-lognormal model. In particular, three quarters of bilateral trade elasticities lie between 4 and 4.2, and they cannot fall below $\sigma - 1 = 4$. We discuss this further in Online Appendix O, where we show that we can reproduce the [Melitz and Redding \(2015\)](#) results but only by setting parameters to values far from those we estimate.

6. Conclusion

The canonical Melitz model of trade with Pareto-distributed firm productivity has a stark prediction: conditional on the fixed costs of exporting, all variation in exports across partners should be due to the number of exporting firms (the extensive margin) and there should be no variation along the intensive margin (exports per exporting firm). We use the World Bank's *Exporter Dynamics Database* plus China to test this prediction. Compared to existing studies, this data allows us to look for systematic variation in the intensive and extensive margins of trade — allowing for year, origin, and destination components of fixed trading costs. We find that at least 40 percent of the variation in exports occurs along the intensive margin. That is, when exports from a given origin to a given destination are high, exports per firm are responsible for, on average, at least 40 percent of the high exports. When we look at average exports by percentile of exporting firms (rather than average exports per firm), we find the intensive margin is more important the higher the size percentile.

Although variation in fixed trade costs across country pairs can make the Melitz-Pareto model fit the intensive margin in the data, such fixed trade costs would need to be negatively correlated with distance. Moreover, variation in fixed trade costs does not reproduce the pattern of a steadily rising intensive margin across exporter size percentiles. Allowing firms to export multiple products or taking into account granularity does not reverse these implications.

In contrast, moving away from a Pareto distribution and assuming that the productivity distribution is lognormal resolves the puzzles. Specifically, we estimate a Melitz model with lognormally distributed firm productivity and idiosyncratic firm-destination demand shifters and fixed trade costs using likelihood methods on the EDD firm-level data. Our estimated Melitz-lognormal model is consistent with the positive intensive margin overall and with the intensive margin rising across exporter size percentiles. This estimated model also implies fixed trade costs that increase with distance.

Since the trade elasticity is no longer a constant in the full Melitz-lognormal model, one would expect from the analysis in [Arkolakis et al. \(2012\)](#) that the welfare effects of a trade cost reduction would be different from those in the Melitz-Pareto model (see [Melitz and Redding, 2015](#)). Extending the exact

that algebra approach popularized by Dekle et al. (2008) to our estimated full Melitz-lognormal model, however, we find that the Melitz-Pareto model provides a remarkably good approximation for the welfare effects of trade liberalization.

Looking ahead, moving from Pareto to lognormal firm productivity may matter more when taking into account how domestic firms can learn from firms selling or producing in the domestic market. The size of this dynamic learning gain from trade should depend on whether the distribution of firm productivity is Pareto versus lognormal, as it interacts with how trade alters the distribution of producer and seller productivity. For example, trade liberalization induces more entry of marginal exporters under Pareto than under lognormal — as illustrated by the unchanging exports per exporter under Pareto (zero intensive margin elasticity, unit extensive margin elasticity) versus the sizable intensive margin and weaker extensive margin under lognormal.

References

- Adao, Rodrigo, Arnaud Costinot, and Dave Donaldson, “Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade,” *American Economic Review*, 2017, 107 (3), 633–89.
- Anderson, James E and Yoto V Yotov, “Short run gravity,” *Journal of International Economics*, 2020, 126, 103341.
- Arkolakis, Costas, “Market Penetration Costs and the New Consumers Margin in International Trade,” *Journal of Political Economy*, 2010, 118 (6), 1151–1199.
- , Arnaud Costinot, and Andrés Rodríguez-Clare, “New Trade Models, Same Old Gains?,” *American Economic Review*, 2012, 102 (1), 94–130.
- , —, Dave Donaldson, and Andrés Rodríguez-Clare, “The Elusive Pro-competitive Effects of Trade,” *Review of Economic Studies*, 2019, 86 (1), 46–80.
- , Natalia Ramondo, Andrés Rodríguez-Clare, and Stephen Yeaple, “Innovation and Production in the Global Economy,” *American Economic Review*, 2018, 108 (8), 2128–73.
- , Sharat Ganapati, and Marc-Andreas Muendler, “The Extensive Margin of Exporting Products: A Firm-Level Analysis,” *American Economic Journal: Macroeconomics*, October 2021, 13 (4), 182–245.
- Bas, Maria, Thierry Mayer, and Mathias Thoenig, “From Micro to Macro: Demand, Supply, and Heterogeneity in the Trade Elasticity,” *Journal of International Economics*, 2017, 108, 1–19.
- Bento, Pedro and Diego Restuccia, “Misallocation, Establishment Size, and Productivity,” *American Economic Journal: Macroeconomics*, 2017, 9 (3), 267–303.

- Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott, "Firms in International Trade," *Journal of Economic Perspectives*, 2007, 21 (3), 105–130.
- , —, —, and —, "The Margins of U.S. Trade," *American Economic Review*, 2009, 99 (2), 487–493.
- Bernard, Andrew B, Jonathan Eaton, J Bradford Jensen, and Samuel Kortum, "Plants and Productivity in International Trade," *American Economic Review*, 2003, 93 (4), 1268–1290.
- Bernard, Andrew B., Stephen J. Redding, and Peter K. Schott, "Multiproduct Firms and Trade Liberalization*," *Quarterly Journal of Economics*, 2011, 126 (3), 1271–1318.
- Brancaccio, Giulia, Myrto Kalouptsi, and Theodore Papageorgiou, "Geography, Transportation, and Endogenous Trade Costs," *Econometrica*, 2020, 88 (2), 657–691.
- Caliendo, Lorenzo, Robert C Feenstra, John Romalis, and Alan M Taylor, "Tariff Reductions, Entry, and Welfare: Theory and Evidence for the Last Two Decades," Technical Report, National Bureau of Economic Research 2015.
- Chan, Jackie M.L., "Financial Frictions and Trade Intermediation: Theory and Evidence," *European Economic Review*, 2019, 119, 567–593.
- Chaney, Thomas, "Distorted Gravity: the Intensive and Extensive Margins of International Trade," *American Economic Review*, 2008, 98 (4), 1707–1721.
- Chernozhukov, Victor and Han Hong, "An MCMC approach to Classical Estimation," *Journal of Econometrics*, 2003, 115 (2), 293–346.
- Costinot, Arnaud and Andrés Rodríguez-Clare, "Trade Theory with Numbers: Quantifying the Consequences of Globalization," *Handbook of International Economics*, 2014, 4, 197.
- Dekle, Robert, Jonathan Eaton, and Samuel Kortum, "Global Rebalancing with Gravity: Measuring the Burden of Adjustment," *IMF Staff Papers*, 2008, 55 (3), 511–540.
- Eaton, Jonathan, Marcela Eslava, Maurice Kugler, and James Tybout, *The Margins of Entry into Export Markets: Evidence from Colombia*, The Organization of Firms in a Global Economy, Cambridge, MA: Harvard University Press, 2008.
- , Samuel Kortum, and Francis Kramarz, "An Anatomy of International Trade: Evidence From French Firms," *Econometrica*, 2011, 79 (5), 1453–1498.
- , —, and Sebastian Sotelo, "International Trade: Linking Micro and Macro," NBER Working Papers 17864, National Bureau of Economic Research, Inc 2012.

- Feenstra, Robert C., “Restoring the Product Variety and Pro-competitive Gains from Trade with Heterogeneous Firms and Bounded Productivity,” *Journal of International Economics*, 2018, 110, 16–27.
- Felbermayr, Gabriel, Benjamin Jung, and Mario Larch, “The welfare consequences of import tariffs: A quantitative perspective,” *Journal of International Economics*, 2015, 97 (2), 295–309.
- Fernandes, Ana M., Caroline Freund, and Martha Denisse Pierola, “Exporter Behavior, country Size and Stage of Development: Evidence from the Exporter Dynamics Database,” *Journal of Development Economics*, 2016, 119, 121–137.
- Freund, Caroline and Martha Denisse Pierola, “Export Superstars,” *Review of Economics and Statistics*, 2015, 97 (5), 1023–1032.
- Gaubert, Cecile and Oleg Itskhoki, “Granular Comparative Advantage,” *Journal of Political Economy*, 2021, 129 (3), 871–939.
- Head, Keith and Thierry Mayer, “Gravity Equations: Workhorse, Toolkit, and Cookbook,” *Handbook of International Economics*, 2014, 4, 131–196.
- , —, and Mathias Thoenig, “Welfare and Trade without Pareto,” *American Economic Review*, 2014, 104 (5), 310–316.
- , — et al., *Poor Substitutes?: Counterfactual Methods in IO and Trade Compared*, Centre for Economic Policy Research, 2021.
- Krugman, Paul, “Scale Economies, Product Differentiation, and the Pattern of Trade,” *American Economic Review*, 1980, 70 (5), 950–959.
- Manova, Kalina and Zhiwei Zhang, “Export Prices across Firms and Destinations,” *Quarterly Journal of Economics*, 2012, 127 (1), 379–436.
- Mayer, Thierry and Soledad Zignago, “Notes on CEPII’s distances measures: The GeoDist database,” Working Papers 2011-25, CEPII research center December 2011.
- Melitz, Marc J., “The Impact of Trade on Intra-industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 2003, 71 (6), 1695–1725.
- and Stephen J. Redding, “New Trade Models, New Welfare Implications,” *American Economic Review*, 2015, 105 (3), 1105–46.
- Nigai, Sergey, “A Tale of Two Tails: Productivity Distribution and the Gains from Trade,” *Journal of International Economics*, 2017, 104 (C), 44–62.

Ossa, Ralph, “Why Trade Matters After All,” *Journal of International Economics*, 2015, 97 (2), 266–277.

Tables and Figures

Table 1: Core Sample of EDD countries+China, years firm-level data is available

ISO3	Country name	1st year	Last year	ISO3	Country name	1st year	Last year
ALB	Albania	2004	2012	KHM	Cambodia	2003	2009
BFA	Burkina Faso	2005	2012	LAO	Laos	2006	2010
BGD	Bangladesh	2005	2013	LBN	Lebanon	2008	2012
BGR	Bulgaria	2003	2006	MAR	Morocco	2003	2013
BOL	Bolivia	2006	2012	MDG	Madagascar	2007	2012
BWA	Botswana	2003	2013	MEX	Mexico	2003	2012
CHL	Chile	2003	2012	MKD	Macedonia	2003	2010
CHN	China	2003	2008	MMR	Myanmar	2011	2013
CIV	Cote d'Ivoire	2009	2012	MUS	Mauritius	2003	2012
CMR	Cameroon	2003	2013	MWI	Malawi	2009	2012
COL	Colombia	2007	2013	NIC	Nicaragua	2003	2013
CRI	Costa Rica	2003	2012	NPL	Nepal	2011	2013
DOM	Dominican Republic	2003	2013	PAK	Pakistan	2003	2010
ECU	Ecuador	2003	2013	PRY	Paraguay	2007	2012
EGY	Egypt	2006	2012	PER	Peru	2003	2013
ETH	Ethiopia	2008	2012	QOS	Kosovo	2011	2013
GAB	Gabon	2003	2008	ROU	Romania	2005	2011
GEO	Georgia	2003	2012	RWA	Rwanda	2003	2012
GIN	Guinea	2009	2012	THA	Thailand	2012	2013
GTM	Guatemala	2005	2013	TZA	Tanzania	2003	2012
HRV	Croatia	2007	2012	UGA	Uganda *	2003	2010
IRN	Iran	2006	2010	URY	Uruguay	2003	2012
JOR	Jordan	2003	2012	YEM	Yemen	2008	2012
KEN	Kenya	2006	2013	ZAF	South Africa	2003	2012
KGZ	Kyrgyzstan	2006	2012	ZMB	Zambia	2003	2011

* Uganda does not have data for 2006

Table 2: IME regressions, core sample

	Coefficient from $\ln x_{ij}$ on $\ln X_{ij}$		
	(1)	(2)	(3)
<i>Panel a: country pairs with $N_{ij} \geq 100$</i>			
IM elasticity	0.438***	0.459***	0.400***
Standard error	[0.0058]	[0.0041]	[0.0055]
R^2	0.55	0.74	0.85
Variation in $\ln X_{ij}$ explained by FE, %	0.01	0.20	0.59
Observations	7,781	7,768	7,324
<i>Panel b: all country pairs</i>			
IM elasticity	0.503***	0.530***	0.579***
Standard error	[0.0018]	[0.0017]	[0.0023]
R^2	0.77	0.81	0.85
Variation in $\ln X_{ij}$ explained by FE, %	0.00	0.20	0.50
Observations	47,129	47,129	47,037
Year FE	Yes		
Origin \times year FE		Yes	Yes
Destination \times year FE			Yes

Note: The table presents the estimated coefficients of the regression of log average exports per firm on log total exports. The data are aggregated at the origin-destination-year level for a set of origin-years listed in Table 1. Panel a) represents the regression on the sample of country-pairs with at least 100 exporters. Panel b) represents the regression on the full sample. Robust standard errors are reported in brackets. *, **, and *** represent the 5%, 1%, and 0.1% significance levels, respectively.

Table 3: Margins of trade and distance

Elasticity with respect to distance		
x_{ij}	0.128***	-0.276***
Standard error	[0.0158]	[0.0164]
N_{ij}	-0.419***	-1.012***
Standard error	[0.0141]	[0.0152]
Observations	7,437	7,019
x_{ij}^p	0.302***	-0.0644***
Standard error	[0.0165]	[0.0172]
m_{ij}	-0.174***	-0.212***
Standard error	[0.0063]	[0.006]
Observations	7,437	7,019
Origin \times year FE	Yes	Yes
Destination \times year FE		Yes

Note: the table presents the estimated coefficients of the regression of log average exports per firm, number of firms, average exports per product per firm (total exports divided by the number of firm-HS6 product observations with positive exports from origin i to destination j in a given year), and number of products on log distance between origins and destinations. The data are aggregated at the origin-destination-year level for a set of origin-years listed in Table 1. Population-weighted distance between origins and destinations is taken from Mayer and Zignago (2011). The sample is restricted to the origin-destination pairs with at least 100 exporters. Egypt is not included in the sample since its data does not include HS 6-digit product level disaggregation. Robust standard errors are reported in brackets. *, **, and *** represent the 5%, 1%, and 0.1% significance levels, respectively.

Table 4: Trade costs and distance

	$\ln \tilde{F}_{ij}$	$\ln \tilde{\tau}_{ij}$	$\ln \tilde{f}_{ij}$
Panel a: Melitz-Pareto model			
$\ln dist_{ij}$	-0.280***	0.272***	-0.071***
Standard error	[0.0143]	[0.0048]	[0.0153]
Observations	7,320	7,320	7,320
Panel b: Melitz-lognormal model			
$\ln dist$	0.151***	0.306***	
Standard error	[0.0148]	[0.0051]	
Observations	7320	7320	

Note: Panel a) of the table presents the estimated coefficients of the regression of the implied firm-level log fixed trade costs (column 1), log variable trade costs (column 2), and product-level log fixed trade costs (column 3) on log distance between origins and destinations. We calculate trade costs by inverting equations (10) and (11) with $\theta = 5$. The data are aggregated at the origin-destination-year level for a set of origin-years listed in Table 1. Population-weighted distance between origins and destinations is taken from Mayer and Zignago (2011). Panel b) of the table presents the estimated coefficients of the regression of the log fixed and variable trade costs implied by the Melitz model with lognormal distribution of productivity as discussed in Section 4.1. on log distance. The sample is restricted to the origin-destination pairs with at least 100 exporters. All regressions include origin-year and destination-year fixed effects. Robust standard errors are reported in brackets. *, **, and *** represent the 5%, 1%, and 0.1% significance levels, respectively.

Table 5: Estimates of dispersion, full Melitz-lognormal model

	mean	median	min	max
$\bar{\sigma}_\varphi$	3.32	3.18	0.93	5.82
σ_α	2.72	2.67	1.94	3.64
σ_f	2.39	2.39	1.64	3.11
ρ	0.47	0.50	-0.33	0.90

Note: The table presents the estimates of the full Melitz-lognormal model. The estimation procedure is discussed in Section 4.1. The sample includes 37 origin countries for which our estimates converged and 15 destinations per origin. The mean, median, min, and max statistics are calculated across origins.

Table 6: Implied IME in full Melitz-Pareto models

	IME	95% CI
Data	0.67	[0.61, 0.73]
Full Melitz-lognormal model	0.63	[0.59, 0.67]
Melitz-Pareto model, constrained	0.64	[0.57, 0.71]

Note: The table presents the coefficient from the regression of log average exports per firm on log total exports with origin and destination fixed effects implied by the simulated full Melitz-lognormal model, Melitz-Pareto constrained model. The sample includes 37 origins and 4 main destinations (USA, Germany, France, and Japan) in 2007. The IME in the data is estimated for the same sample. The point estimates and 95% confidence intervals are calculated based on 1,000 simulations based on random parameter draws from the generated Monte-Carlo Markov chain.

Table 7: Implied trade costs in simulated models

	Estimate	95% CI	Estimate	95% CI
	Full Melitz lognormal		Melitz-Pareto constrained	
$corr(\tilde{F}_{ij}, \tilde{\tau}_{ij})$	-0.31	[-0.45, -0.1]	-0.70	[-0.52, 0.08]
Distance elasticity:				
Fixed costs	0.31	[0.18, 0.41]	0.49	[0.37, 0.65]
Variable costs	0.34	[0.30, 0.37]	0.16	[0.13, 0.19]

Note: The table presents the coefficients from the regression of log fixed and variable trade costs on distance, origin, and destination fixed effects implied by the simulated full Melitz-lognormal and Melitz-Pareto constrained models. The sample includes 37 origins and 4 main destinations (USA, Germany, France, and Japan) in 2007. The point estimates and 95% confidence intervals are calculated based on 1,000 simulations based on random parameter draws from the generated Monte-Carlo Markov chain.


Figure 1: Intensive and Extensive margins of exporting

Panel A: Average size of exporters and total exports *Panel B: Number of exporters and total exports*




Source: Exporter Dynamics Database, extended sample. The x-axis represents log total exports at the origin-destination-year level demeaned by origin-year, and destination-year fixed effects. Only origin-destination pairs with more than 100 exporters included. The dots represent the raw measures. The line is the slope predicted by the Melitz-Pareto model.

Figure 2: IME for each percentile, data



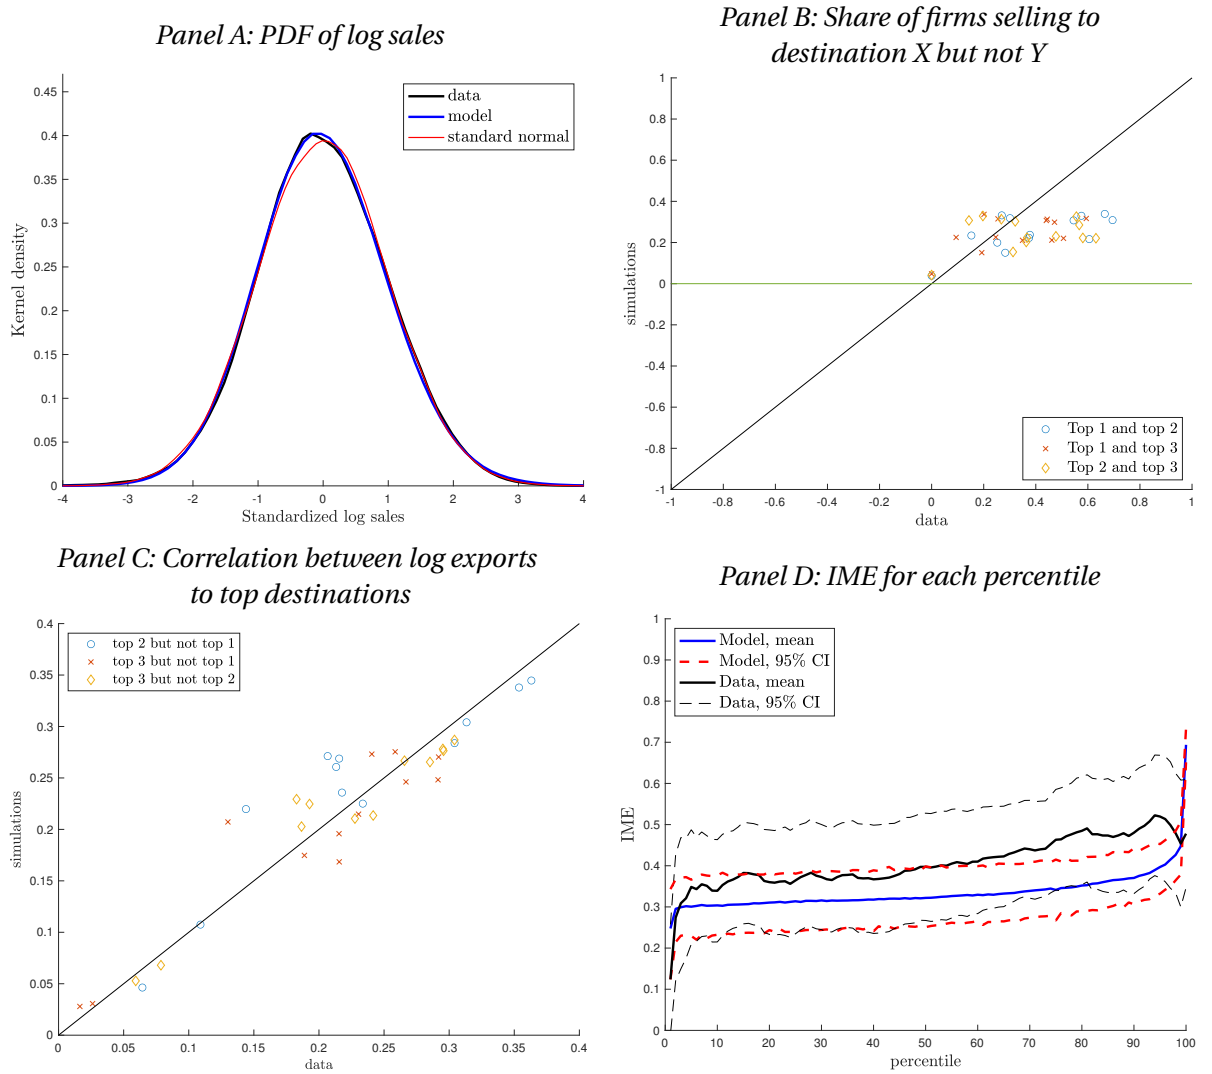
Source: Exporter Dynamics Database, core sample of countries. The x-axis represents percentiles of the average exporter size distribution. Each dot represents the coefficient from the regression of log average exports per firm in an exporter size percentile on log total exports. The data is demeaned by origin-year and destination-year fixed effects.

Figure 3: IME for each percentile, lognormal



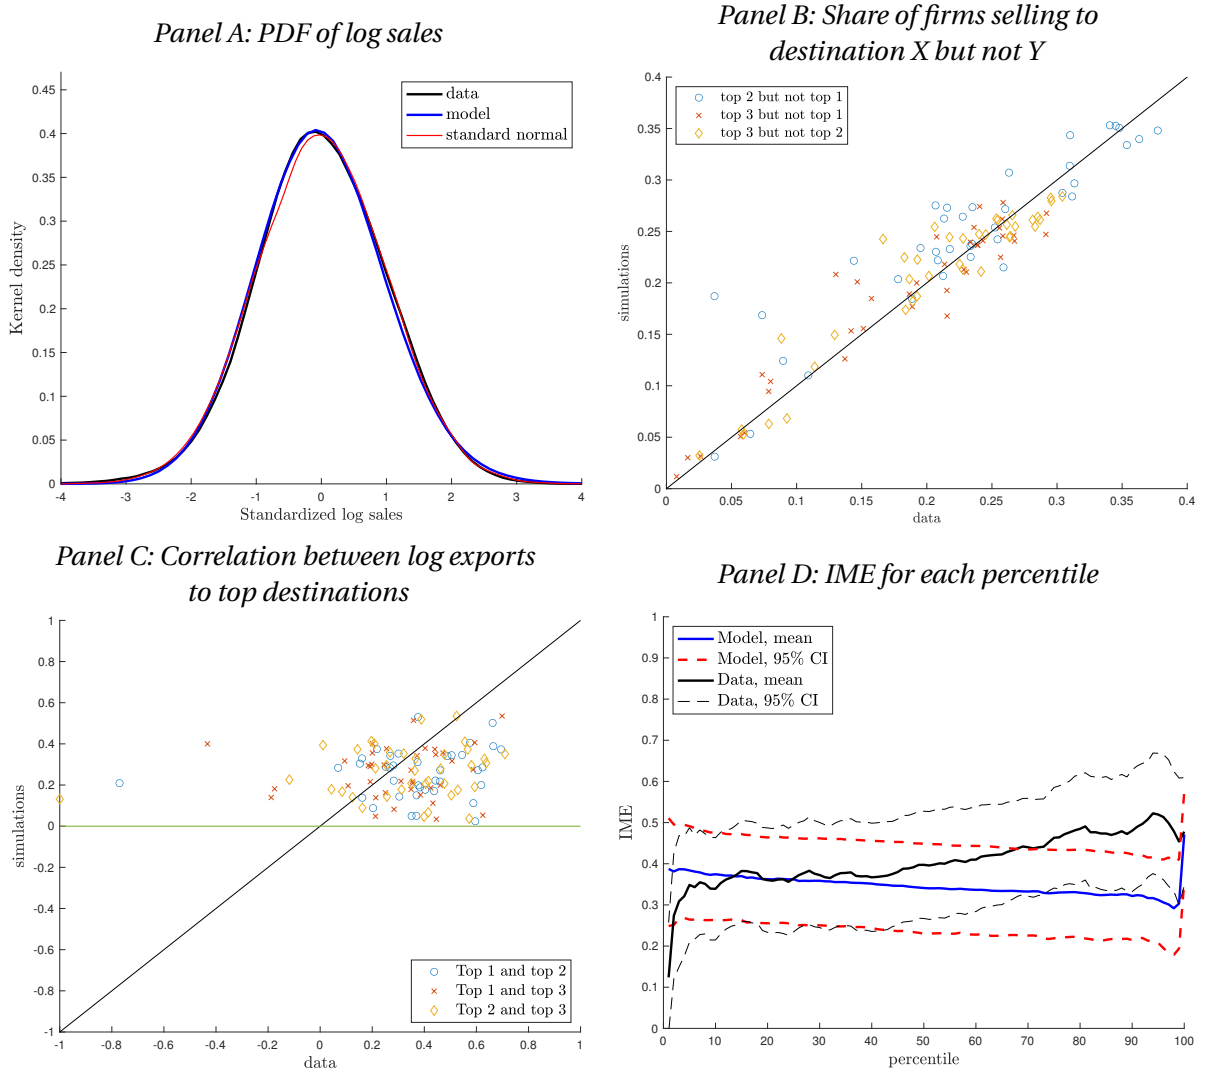
Source: Exporter Dynamics Database. The darker solid line corresponds to IME for each percentile estimated using EDD and four main destinations: France, Germany, Japan and the U.S.. Dashed lines indicate 95% confidence intervals. The lighter solid line is IME for each percentile implied by the model with lognormal distribution of productivity, $\bar{\sigma}_\varphi = 4.02$ (our estimate) and $\sigma = 5$ from [Bas et al. \(2017\)](#). The level of bilateral fixed trade costs was chosen to match overall IME in the data. The total number of firms was imputed from [Bento and Restuccia \(2017\)](#).

Figure 4: Full Melitz-lognormal model, goodness of fit



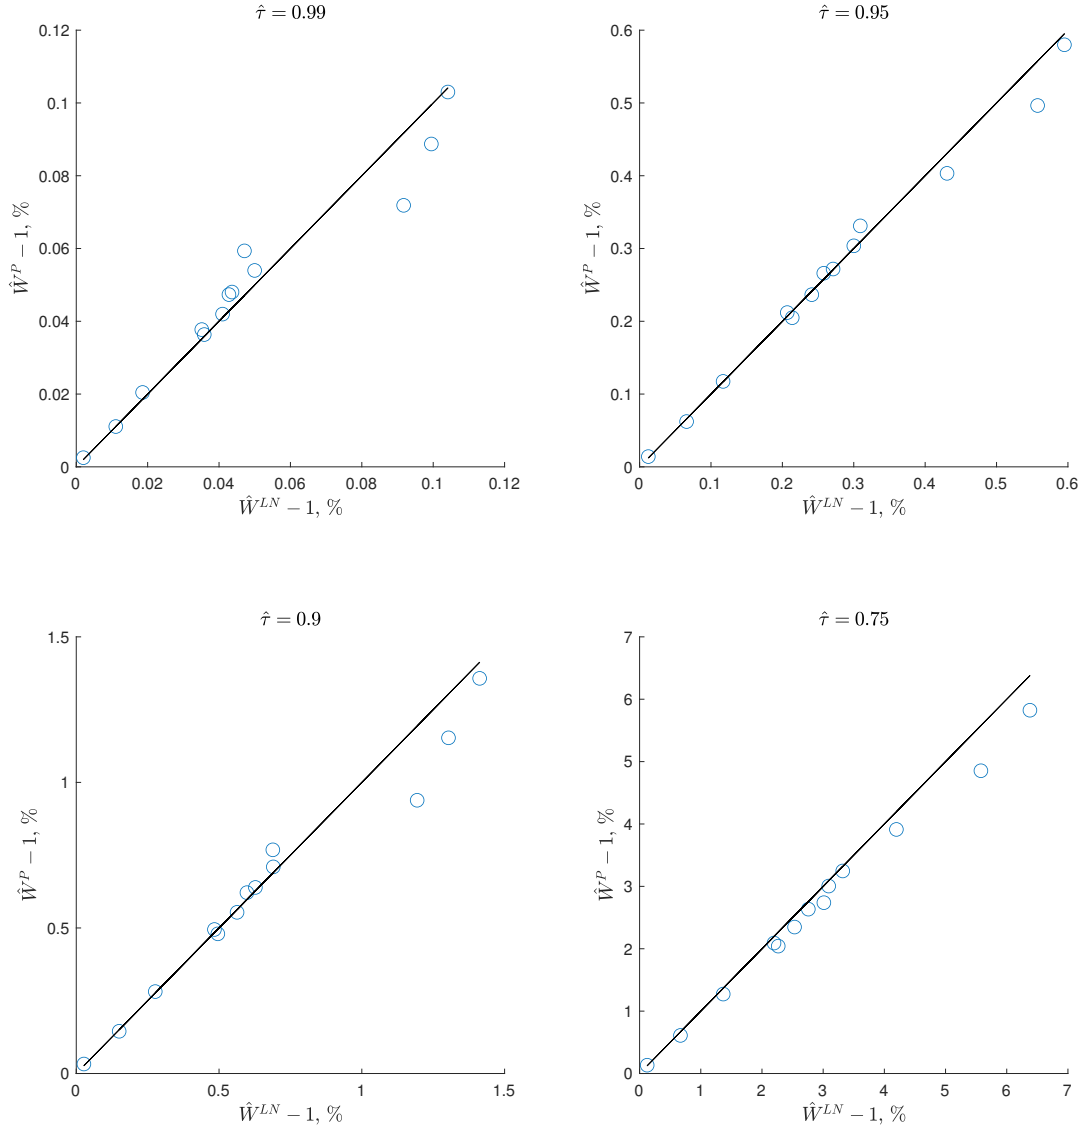
Source: Exporter Dynamics Database and authors' calculations. Panel A: The black line corresponds to the standardized log sales (demeaned, divided by standard deviation) pooled across different origin-destination cells. The blue line corresponds to the standardized log sales pooled across different cells in the model. Panel B: Each point corresponds to the share of firms exporting only to less popular markets in the data (horizontal axis) and according to the estimated model (vertical axis) for each origin. Panel C: each point corresponds for each origin and any two destinations among the three most popular ones, the correlation in export value across all firms that sell in those two destinations in the data (horizontal axis) and according to the estimated model (vertical axis). Panel D: the x-axis represent percentiles; the blue solid line represents coefficient from the regression of log average exports in each percentile on log total exports in the model; the dashed red lines represent 95% confidence interval; the black solid line represents coefficient from the regression of log average exports in each percentile on log total exports in the data; the dashed black lines represent the 95% confidence interval.

Figure 5: Melitz-Pareto model (constrained), goodness of fit




Source: Exporter Dynamics Database and authors' calculations. Panel A: The black line corresponds to the standardized log sales (demeaned, divided by standard deviation) pooled across different origin-destination cells. The blue line corresponds to the standardized log sales pooled across different cells in the model. Panel B: Each point corresponds to the share of firms exporting only to less popular markets in the data (horizontal axis) and according to the estimated model (vertical axis) for each origin. Panel C: each point corresponds for each origin and any two destinations among the three most popular ones, the correlation in export value across all firms that sell in those two destinations in the data (horizontal axis) and according to the estimated model (vertical axis). Panel D: the x-axis represent percentiles; the blue solid line represents coefficient from the regression of log average exports in each percentile on log total exports in the model; the dashed red lines represent the 95% confidence interval; the black solid line represents coefficient from the regression of log average exports in each percentile on log total exports in the data; the dashed black lines represent the 95% confidence interval.

Figure 6: Gains from trade liberalization




Note: The figure represents the change in welfare in response to a variable trade costs shock in the full Melitz-lognormal model and the Melitz-Pareto model. To calculate welfare gains in the full Melitz-lognormal model we used the parameter estimates from the Monte-Carlo Markov chain. Section 5.1. describes the procedure to calculate gains from trade liberalization in the full Melitz-lognormal model. We used the Dekle et al. (2008) ‘exact hat’ algebra to calculate changes in trade shares in the Melitz-Pareto model and the Arkolakis et al. (2012) formula to calculate the gains from trade liberalization. The x-axis represents gains in the full Melitz-lognormal model. The y-axis represents gains in the Melitz-Pareto model. Each of the four panels reports the results for a different change in trade costs (1%, 5%, 10%, 25%). In the Melitz-Pareto model when we use the trade elasticity estimated from equation (24).

Figure 7: Counterfactual changes in trade flows




Note: The figure represents the difference between changes in trade flows in the full Melitz-lognormal model and the Melitz-Pareto model in response to a reduction of variable trade costs on the vertical axis, and the trade elasticity implied by the full Melitz-lognormal model on the horizontal axis. To calculate changes in trade flows in the full Melitz-lognormal model we used parameter estimates from the Monte-Carlo Markov chain. Section 5.1. describes the procedure to calculate changes in trade flows in the full Melitz-lognormal model.

Figure 8: Gains from asymmetric trade liberalization



Note: the figure represents the change in welfare in response to a variable trade costs shock in the full lognormal Melitz model and the Melitz-Pareto model. Each dot represents a change in welfare in a given origin in response to a 25% decline in costs of exporting to its biggest market. To calculate welfare gains in the full Melitz-lognormal model we used the parameter estimates from the Monte-Carlo Markov chain. Section 5.1. describes the procedure to calculate gains from trade liberalization in the full Melitz-lognormal model. We used the [Dekle et al. \(2008\)](#) 'exact hat' algebra to calculate changes in trade shares in the Melitz-Pareto model and the [Arkolakis et al. \(2012\)](#) formula to calculate the gains from trade liberalization. The x-axis represents gains in the full Melitz-lognormal model. The y-axis represents gains in the Melitz-Pareto model.

Figure 9: Counterfactual changes in trade flows, asymmetric trade liberalization



Note: the figure represents changes in trade flows in response to a variable trade costs shock in the full lognormal Melitz model and the Melitz-Pareto model. Each dot represents a change in welfare in a given origin in response to a 25% decline in costs of exporting to its biggest market. To calculate changes in trade flows in the full Melitz-lognormal model we used the parameter estimates from the Monte-Carlo Markov chain. Section 5.1. describes the procedure to changes in trade flows after trade liberalization in the full Melitz-lognormal model. We used the Dekle et al. (2008) 'exact hat' algebra to calculate changes in trade shares in the Melitz-Pareto model. The x-axis represents changes in trade flows in the full Melitz-lognormal model. The y-axis represents changes in trade flows in the Melitz-Pareto model.