

The intensive margin in trade

*Ana M. Fernandes, Peter J. Klenow, Sergii Meleshchuk, Martha Denisse Pierola,
Andrés Rodríguez-Clare*

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

www.cesifo-group.org/wp

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

The intensive margin in trade

Abstract

Is the variation in bilateral trade flows across countries primarily due to differences in the number of exporting firms (the extensive margin) or in the average size of an exporter (the intensive margin)? And how does this affect the estimation and quantitative implications of the Melitz (2003) trade model? The benchmark Melitz model with Pareto-distributed firm productivity and fixed costs of exporting, predicts that, conditional on the fixed costs of exporting, all variation in exports across trading partners should occur on the extensive margin. We subject this theoretical prediction to a reality check drawing upon the World Banks Exporter Dynamics Database (EDD) which has firm-level exports from 50 developing countries to all destinations. We find that around 50 percent of the variation in exports across trading partners is along the intensive margin, contradicting the benchmark Melitz-Pareto model. We find that moving from a Pareto to a lognormal distribution of firm productivity allows the Melitz model to successfully match the role of the intensive margin evident in the EDD. We then study the implications of our findings for quantitative trade theory. Using likelihood methods and the EDD, we estimate a generalized Melitz model with a joint lognormal distribution for firm productivity, fixed costs and demand shifters, and use exact hat algebra to quantify the counterfactual effects of a decline in trade costs on trade flows and welfare in the estimated model. Finally, we compare these effects to those that would be predicted by the Melitz-Pareto model, with the Pareto shape parameter chosen to match the average trade elasticity implied by our estimated Melitz-lognormal model. We find that the effects on welfare turn out to be quite close to those in the standard Melitz-Pareto model even though the effects on trade flows remain different.

JEL-Codes: F100, F120.

Keywords: intensive margin of trade, extensive margin of trade, productivity distribution, trade costs, welfare, Pareto.

Ana M. Fernandes
World Bank / Washington DC / USA
afernandes@worldbank.org

Peter J. Klenow
Stanford University / CA / USA
pete.klenow@gmail.com

Sergii Meleshchuk
IMF / Washington DC / USA
meleshchuk@berkeley.edu

Martha Denisse Pierola
Inter-American Development Bank /
Washington DC / USA
mpierola@iadb.org

Andrés Rodríguez-Clare
UC Berkeley / CA / USA
andres1000@gmail.com

October 26, 2018

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

Across trading partners, exports can vary along the extensive margin (number of exporting firms) and along the intensive margin (average exports per exporting firm). The classic Krugman (1980) model predicts all export variation will be on the intensive margin because all firms export to every destination. Melitz (2003) brings the extensive margin to life with fixed costs of exporting, and emphasizes the importance of selection of firms into exporting.

In this paper we document the importance of the two margins empirically and then study the implications of our findings for quantitative trade theory. Exploiting data on a much larger set of exporting countries than in the existing literature,¹ we find that the intensive margin is at least as important as the extensive margin in driving bilateral trade flows. The importance of the intensive margin is at odds with the Melitz model if one assumes that firm productivity is distributed Pareto, as is customary in the quantitative trade literature. In the benchmark Melitz-Pareto model, the extensive margin is dominant. In contrast, a lognormal distribution of productivities allows the Melitz model to successfully match the empirical prominence of the intensive margin. We finish by examining whether moving from a Pareto to a lognormal firm productivity distribution alters how trade costs affect trade flows and welfare.

To elaborate, we use the World Bank's *Exporter Dynamics Database* (hereafter EDD) to systematically examine the importance of the extensive and intensive margins. The EDD covers firm-level exports from 59 (mostly developing) countries to all destination countries in most years from 2003 to 2013. For 49 of the countries, every exporting firm's exports to each destination in a given year can be broken down into HS 6-digit products.² We add China to the 49 EDD countries to arrive at 50 countries for our analysis. Having many origin and destination countries enables us to study the role of the intensive and extensive margins while allowing for origin-year and destination-year fixed effects that control for differences in population, wages, and other country characteristics that affect firm entry into exporting and exports per firm.

We find that between 40 and 60 percent of the variation in overall exports across origin-destination pairs is accounted for by the intensive margin, with the rest accounted for by the extensive margin. For reasons that will become clear in Section 2, we refer to these two shares as the intensive margin and extensive margin elasticities. This breakdown into the intensive and extensive margin elasticities is robust to using different country samples or sets of fixed effects,

¹Most firm-level empirical trade studies have only 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 et al. (2012a) for Denmark and France, Manova and Zhang (2012) for China, and Arkolakis and Muendler (2013) for Brazil, Chile, Denmark and Norway.

²See Fernandes, Freund and Pierola (2016) for a detailed description of the dataset.

excluding country pairs with few exporters or tiny exporters, and looking within industries.

We interpret the finding that up to 60 percent of the variation in bilateral trade flows are explained by the extensive margin as providing support for the Melitz (2003) model. But the finding that at least 40 percent of that variation is explained by the intensive margin — even while allowing for origin-year and destination-year fixed effects — contradicts an important special case of the Melitz model, namely the one with Pareto-distributed firm productivity and fixed trade costs that vary only because of separate origin and destination components. In this case all variation in exports across trading partners should occur through the number of exporters (the extensive margin). Lower variable trade costs should stimulate sales of a given exporting firm, but draw in marginal exporting firms to the point that average exports per exporter (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. 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, Costinot and Rodríguez-Clare, 2012). We want to make sure our empirical finding cannot be made compatible with Melitz-Pareto before exploring an alternative productivity distribution.

First, we consider the possibility that fixed trade costs vary by origin-destination pair. Higher fixed trade costs raise average exports per exporter, but *lower* total exports. For average exports to be increasing in total exports, therefore, one needs variable trade costs to be very negatively correlated with fixed trade costs. A corollary is that, whereas variable trade costs rise decisively with distance between trade partners, fixed trade costs would need to fall with distance. We find this rather implausible. More importantly, this model would imply that the intensive margin elasticity should be equally important among the smallest and largest exporting firms; in the data, the intensive margin elasticity rises steadily with exporting firm size when we put exporting firms into size percentiles.

Second, we explore the role of multi-product firms. If the typical firm exports a higher number of products to destinations with larger overall exports, this could account for the importance of the intensive margin for exports. We find that the number of HS 6-digit products per exporting firm does indeed account for about 12 percent of the variation in overall exports, or about

³This 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, 2015); non-monopolistic competition (Bernard, Eaton, Jensen and Kortum, 2003), and multinational production (Arkolakis, Ramondo, Rodríguez-Clare and Yeaple, 2014).

one-fourth of the intensive margin elasticity we estimate. In the context of the multi-product Melitz-Pareto model developed by Bernard, Redding and Schott (2011), however, this explanation still requires *firm-level* fixed costs to fall with the distance between trading partners. And the intensive margin elasticity per firm-product requires that fixed costs of exporting *per product* also fall with distance.

A third hypothesis we investigate is granularity. With a finite number of firms, the intensive margin (and overall exports) can be high because of favorable productivity draws from the Pareto distribution within a country. We develop an estimator for the elasticity of fixed trade costs to distance that is valid under granularity as in Eaton, Kortum and Sotelo (2012a), and continue to find that fixed trade costs must fall with distance to explain a positive intensive margin elasticity. Using simulations of finite draws from a Pareto distribution, we find that granularity generates only a modest intensive margin elasticity, and — in contrast to the data — almost entirely in the right tail of the exporter size distribution.

After these attempts to rescue the Melitz-Pareto special case, we entertain a lognormal distribution of firm productivity. 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 (2015) show how the trade elasticity varies with a lognormal distribution. Both papers marshal evidence from firms in France and China pointing to the empirical relevance of the lognormal distribution.

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 as well as an idiosyncratic demand shifter and 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 is potentially not concave as a function of the parameters, and since we have a large number of parameters to estimate (means, variances, covariance, and trade costs), we rely on the estimation methodology proposed by Chernozhukov and Hong (2003).

Our estimation shows that a lognormal distribution for firm productivity can indeed 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 Shifting to lognormal

⁴The result holds under other thin-tailed productivity distributions, such as bounded Pareto as in Feenstra (2014).

productivity also changes our inference about fixed trade costs, so that now they are *rising* with distance. 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 quantitative impact of changes in trade costs. We show how to extend the “exact hat algebra” used in Dekle, Eaton and Kortum (2008) 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 counterfactual changes in trade costs on trade flows and welfare in our estimated full Melitz-lognormal model. We compare these effects to those that obtain 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 in this Melitz-Pareto approximation are very close to those in the Melitz-lognormal model, although there are significant differences regarding the effects on trade flows.⁵

Our findings for the intensive margin and the importance of assuming a particular productivity distribution have implications beyond those from these counterfactual exercises. The Melitz model has no frictions to reallocating inputs across firms, and takes firm productivities as exogenous. In the presence of reallocation frictions, whether the adjustment takes place along the intensive or extensive margin could alter how trade liberalization impacts trade flows and welfare. The distribution of firm productivity, meanwhile, affects technology diffusion through trade in studies such as Buera and Oberfield (2015) and Perla, Tonetti and Waugh (2015), which assume fat-tailed productivity distributions such as Pareto. In Perla et al. (2015), the gains from trade are larger when trade liberalization raises average firm size relative to cutoff firm size — which is closely tied to the intensive margin elasticity.

Our counterfactual analysis is related to Head et al. (2014) and Melitz and Redding (2015). Using lognormal and bounded Pareto distributions, 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 if 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

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 and Mayer (2018). 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 counterfactual implications are different 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.

and Pareto models. In contrast, Melitz and Redding (2015) show that the formula proposed in Arkolakis et al. (2012) to compute welfare 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 Melitz-Pareto 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.

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 ad hoc fixed trade costs, multi-product firms and granularity. 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 and the EDD to estimate a Melitz model with a lognormal distribution for productivity plus idiosyncratic demand shocks and fixed costs. Finally, we extend the exact hat algebra approach to a generalized Melitz model and use it with the estimated Melitz-lognormal model to explore counterfactual trade flow and welfare implications relative to those of the Melitz-Pareto model.

The rest of the paper is organized as follows. Section 2 describes the EDD data and document the importance of the extensive and intensive margins in accounting for cross-country variation in exports. Section 3 contrasts the predictions of the Melitz-Pareto model (with a continuum of single product firms, multi-product firms or a finite number of firms) to the EDD facts. Section 4 shows how the implications of the Melitz model change when we drop the Pareto assumption and instead assume that the firm productivity distribution is lognormal. Section 5 describes the quantitative impact of trade cost shocks estimated based on “exact hat algebra.” Section 6 concludes.

2. The Intensive Margin in the Data

The Exporter Dynamics Database

We use the Exporter Dynamics Database (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). For each country, the raw firm-level customs data con-

tains annual export flows (in current values) disaggregated by firm, destination and Harmonized System (HS) 6-digit product. Oil exports are excluded from the customs data due to lack of accurate firm-level 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. More than 100 statistics from the EDD are publicly available at the origin-year, origin-product-year, origin-destination-year, or origin-product-destination-year levels. These include average exports per firm as well as the number of exporting firms.

For the descriptive analysis in this section as well as for the regression and simulation 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.⁶ However, to use the most comprehensive sample of countries available we rely for the motivating plots below on an *extended sample* that includes the 59 origin countries from the EDD plus China. Both samples cover a subset of years from 2003 and 2013 — see Table 1 and Table II in the Online Appendix.

We focus on EDD statistics based on products belonging only to the manufacturing sector. Specifically, relying on a concordance across the ISIC rev. 3 classification and the HS 6-digit classification, we consider only exports of HS 6-digit products that correspond to ISIC manufacturing sub-sectors 15-37. Using these data 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.⁷ 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.⁸ 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.

⁶China is not included in the EDD due to confidentiality concerns.

⁷There is also variation in our data over time, so in principle we should add a time subscript as well, as in X_{ijt} . For simplicity, we suppress the time subscript.

⁸The 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.

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 corresponds to 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 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).⁹ 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.¹⁰

Robustness

The finding of a positive and large IME is robust to considering different samples. In Panel B of Table 2 we estimate the IME including all country pairs — even those with less than 100 exporting firms. The IME in this case reaches 0.58 when origin-year and destination-year fixed effects are included. In the Online Appendix Table 12 we reproduce the regressions in Table 2

⁹To be specific, the equation estimated in this case is $\ln x_{ijt} = FE_{it}^o + FE_{jt}^d + \alpha \ln X_{ijt} + \varepsilon_{ijt}$ using all years of available data for the country pairs included in the core sample.

¹⁰This percentage comes from the R-squared of an OLS regression of bilateral total exports in logs ($\ln X_{ijt}$) on origin-year and destination-year fixed effects.

but now for the extended sample of countries. In the preferred specification with origin-year and destination-year fixed effects, the IME is 0.38 among origin-destination pairs with at least 100 exporting firms and 0.52 among all origin-destination pairs.¹¹ To make sure the IME is not driven by small exporting firms, we re-estimate it after excluding firms whose annual exports fell below \$1,000 in any year. The corresponding IME estimates in Table 3 (core sample) and Online Appendix Table I4 (extended sample) change only slightly.

A separate concern is measurement error. Since total exports is the sum of firm-level exports, classical measurement error in exports per exporter x would bias the IME upward, but classical measurement error in the number of exporters N would bias the IME downward. Depending on their relative importance compared to the true IME, classical measurement error could bias the IME upward or downward. If the measurement error is serially uncorrelated, then instrumenting total exports with its leads and/or lags should yield an unbiased estimate of the IME. As shown in Online Appendix Table I5, the instrumented IMEs are very close to the OLS IME, both economically and statistically.

Our results for the IME could be coming from country differences in industry composition of exports combined with industry differences in average exports per firm. In Figure 2 we plot the (demeaned) intensive and extensive margins against total exports at the origin-industry-destination-year level using HS 2-digit industries. The pattern here is similar to that in Figure 1. Table 4 shows that the IME actually increases when moving to industry-level data. At the lowest level of aggregation available (HS 6-digit), for the core sample of countries the IME is 0.51 with origin-year-industry and destination-year-industry fixed effects. The results also hold in the extended sample, for which we calculated IME disaggregated at HS2 product level. As reported in the Online Appendix Table I6, this IME is also close to 0.52.¹²

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

¹¹When we allow the IME to differ across origin countries depending on their GDP per capita, we find that IME estimates are close to or larger than 0.4 for any group of countries, as shown in Online Appendix Table I3.

¹²The presence of large trading firms could lead to both high exports per firm and high total exports and explain our IME estimates. While we are unable to identify large trading firms in the EDD data, we estimate the IME based on a sample including only HS 2-digit industries with low shares of firms exporting via intermediaries, as defined in Chan (2017). The results in Online Appendix Table I7 show an almost unchanged IME at 0.53.

pct as x_{ij}^{pct} , we run the regressions:

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

We define the IME for each percentile as $IME^{pct} \equiv \hat{\alpha}^{pct}$.

We plot the IME^{pct} for each percentile (with confidence intervals) in Figure 3 along with the horizontal line at the overall IME of 0.4.¹³ 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.

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. Let m_{ij} be the average number of products exported from i to j by firms exporting from i to j , and let $x_{ij}^p \equiv x_{ij}/m_{ij}$ be the average exports per product per firm exporting from i to j . We define the IME *at the product level* as $IME^p \equiv cov(\ln \tilde{x}_{ij}^p, \ln \tilde{X}_{ij})/var(\ln \tilde{X}_{ij})$. Since $x_{ij} = x_{ij}^p m_{ij}$, the IME is equal to the IME^p plus the extensive product margin elasticity,

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

Table 5 reports the results for the IME^p for the core sample.¹⁴ 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 by firm.

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.

¹³For exporter percentiles to be well-defined we focus on country pairs for which $N_{ij} > 100$.

¹⁴Bernard et al. (2009) present a similar decomposition for U.S. exports. We compare their results to ours below.

Correlation between intensive and extensive margin, and 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 6 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.

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 firm increase with market size of the destination (measured as manufacturing absorption) with an elasticity of 1/3. In Figure 4 we plot market size (horizontal axis) against our estimated destination fixed effects (vertical axis) from a regression of average exports per firm on origin, destination and year fixed effects based on the core sample and country pairs for which $N_{ij} > 100$. A regression line (not shown) through the points in the plot implies that average exports per firm increase with destination market size with an elasticity of 0.19, a bit lower than the result in EKK.¹⁵

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 instead 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 $l^*(i) \equiv \arg \max_k N_{ik}$ be the most popular destination market for each origin country i (e.g., the United States (U.S.) for Mexico). In Figure 5 we plot $\log \frac{x_{il^*(i)|j}}{x_{il^*(i)|l^*(i)}}$ (vertical axis) against $\log \frac{N_{ij}}{N_{il^*(i)}}$ (horizontal axis) for all i and j for the core sample.¹⁶ It is very clear that the results derived by EKK for French firms remains valid for 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.

¹⁵Similar findings are obtained in unreported plots where the destination fixed effects are based on the extended sample and all country pairs or based on the core sample and either country pairs for which $N_{ij} > 100$ or all country pairs.

¹⁶The 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 $l^*(i)$.

EKS find that average exports per firm are very similar across four origin countries (Brazil, Denmark, France and Uruguay) for which they have customs data. 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.¹⁷ 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 0.65 when only country pairs with $N_{ij} > 100$ are considered and only 0.37 when all country pairs are considered.

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. They find that IME^p is around 0.23, which is not far from our finding of around 0.29. On the other hand, contrary to our results, Bernard et al. (2007) find that average exports per product per firm increase with distance. We believe that the difference arises from the fact that, by having data for multiple origins, we are able to control for destination fixed effects. In fact, Table 6 shows that regressing $\ln x_{ij}^p$ on $\ln \text{dist}_{ij}$ with only origin and year fixed effects but without destination fixed effects yields a positive and significant coefficient as in Bernard et al. (2007), whereas the coefficient becomes negative and significant when destination fixed effects are added. The same happens when regressing $\ln x_{ij}$ on $\ln \text{dist}_{ij}$.

3. The Intensive Margin in the Melitz-Pareto Model

In this section we ask how the Melitz model with Pareto distributed firm productivity stacks up relative to the findings of the previous section. We focus on the implications of this model for the intensive margin elasticity. We start with the simplest model, which entails a continuum of single-product firms with a Pareto distribution for productivity as in Chaney (2008) and Arkolakis et al. (2008). We derive a series of properties of this model, and then explore their robustness to allowing for destination-specific demand and fixed trade cost shocks at the firm level as in EKK, for multi-product firms, and for granularity.

¹⁷For this purpose, we run regressions of the estimated origin fixed effects on population, GDP, GDP per capita, and total exports, jointly and separately.

3.1. The Basic Melitz-Pareto Model

Theory

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 are constant elasticity of substitution (CES) with elasticity of substitution across varieties $\sigma > 1$. Each firm produces one variety under monopolistic competition. In each country i there is a large pool of firms of measure N_i with productivity φ distributed Pareto with shape parameter $\theta > \sigma - 1$ and scale parameter b_i , $\Pr(\varphi \leq \varphi_0) = G_i(\varphi_0) = 1 - (\varphi_0/b_i)^{-\theta}$. Firms from country i also incur fixed trade costs F_{ij} as well as iceberg trade costs τ_{ij} to sell in country j .¹⁸

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

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

where $A_j \equiv P_j^{1-\sigma} w_j L_j$, $P_j^{1-\sigma} = \sum_i N_i \int_{\varphi \geq \varphi_{ij}^*} \left(\bar{\sigma} \frac{w_i \tau_{ij}}{\varphi} \right)^{1-\sigma} dG_i(\varphi)$ 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 overall exports and the number of firms that export from i to j are then $X_{ij} = N_i \int_{\varphi \geq \varphi_{ij}^*} x_{ij}(\varphi) dG_i(\varphi)$ and $N_{ij} = N_i \int_{\varphi \geq \varphi_{ij}^*} dG_i(\varphi)$, respectively. Using again the fact that $G_i(\varphi)$ is Pareto and assuming that $\varphi_{ij}^* > b_i$ for all i, j , we get that

$$X_{ij} = \left(\frac{\theta}{\theta - (\sigma - 1)} \right) A_j (w_i \tau_{ij})^{1-\sigma} b_i^\theta N_i (\varphi_{ij}^*)^{\sigma-\theta-1} \quad (5)$$

and

$$N_{ij} = b_i^\theta N_i (\varphi_{ij}^*)^{-\theta}. \quad (6)$$

Combining Equations (3) - (6), the extensive margin is

$$N_{ij} = N_i \left(\frac{w_i}{b_i} \right)^{-\theta} \left(\frac{\sigma}{A_j} \right)^{-\theta/(\sigma-1)} \tau_{ij}^{-\theta} F_{ij}^{-\theta/(\sigma-1)}, \quad (7)$$

¹⁸ F_{ij} is in units of the numeraire. Since we focus on cross-section properties of the equilibrium, we do not need to specify whether the fixed trade cost entails hiring labor in the origin or the destination country.

while the intensive margin is

$$x_{ij} \equiv \frac{X_{ij}}{N_{ij}} = \left(\frac{\theta\sigma}{\theta - (\sigma - 1)} \right) F_{ij}. \quad (8)$$

We can always decompose variable and fixed trade costs as follows: $\tau_{ij} = \tau_i^o \tau_j^d \tilde{\tau}_{ij}$ and $F_{ij} = F_i^o F_j^d \tilde{F}_{ij}$. Taking logs in (7) and (8), and defining variables appropriately, we have

$$\ln N_{ij} = \mu_i^{N,o} + \mu_j^{N,d} - \theta \ln \tilde{\tau}_{ij} - \bar{\theta} \ln \tilde{F}_{ij} \quad (9)$$

and

$$\ln x_{ij} = \mu_i^{x,o} + \mu_j^{x,d} + \ln \tilde{F}_{ij}, \quad (10)$$

where $\bar{\theta} \equiv \frac{\theta}{\sigma-1}$. These are the two key equations that we use to derive the results in the rest of this section.

Combining the definition of the intensive margin elasticity given in the previous section (i.e., $\text{IME} = \frac{\text{cov}(\ln \tilde{x}_{ij}, \ln \tilde{X}_{ij})}{\text{var}(\ln \tilde{X}_{ij})}$) with equations (9) and (10), the model implies that

$$\text{IME} = \frac{-(\bar{\theta} - 1) \text{var}(\ln \tilde{F}_{ij}) - \theta \text{cov}(\ln \tilde{\tau}_{ij}, \ln \tilde{F}_{ij})}{\text{var}(-\theta \ln \tilde{\tau}_{ij} - (\bar{\theta} - 1) \ln \tilde{F}_{ij})}. \quad (11)$$

This result can be used to extract several implications of the model, which we present in the form of four observations in the rest of this section.

Our first observation says that if 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)), then the model implies that the intensive margin elasticity is zero.

Observation 1: If $\text{var}(\ln \tilde{F}_{ij}) = 0$ then $\text{IME} = 0$.

Since this is a key result, it is worth understanding it in more detail. Using Equation (3) together with the definition of x_{ij} , taking logs and differentiating w.r.t. $\ln \tau_{ij}$ we get

$$\frac{d \ln x_{ij}}{d \ln \tau_{ij}} = 1 - \sigma - \frac{d \ln (1 - G_i(\varphi_{ij}^*))}{d \ln \varphi_{ij}^*} \left(1 - \frac{x_{ij}(\varphi_{ij}^*)}{x_{ij}} \right).$$

The first term is the direct effect on incumbent firms, while the second term captures selection. In turn, selection is the product of $-\frac{d \ln (1 - G_i(\varphi_{ij}^*))}{d \ln \varphi_{ij}^*}$, which captures the effect of τ_{ij} (and hence φ_{ij}^*) on average exports per firm through its impact on the share of firms that export, and $\left(1 - \frac{x_{ij}(\varphi_{ij}^*)}{x_{ij}} \right)$, which captures how much less firms export at the cutoff relative to the average.

Obviously, if $x_{ij} = x_{ij}(\varphi_{ij}^*)$ then there is no selection, while the effect of selection is maximized if $x_{ij}(\varphi_{ij}^*)/x_{ij} = 0$. With a Pareto distribution for productivity we have $-\frac{d \ln(1-G_i(\varphi_{ij}^*))}{d \ln \varphi_{ij}^*} = \theta$ and $\frac{x_{ij}}{x_{ij}(\varphi_{ij}^*)} = \frac{\theta}{\theta-(\sigma-1)}$, therefore $\frac{d \ln x_{ij}}{d \ln \tau_{ij}} = 0$.

Combined with the assumption that $\bar{\theta} > 1$, the result in equation (11) also implies that if the intensive margin elasticity is positive then there must be a negative correlation between the variable and fixed trade costs (ignoring origin and destination fixed costs).

Observation 2: If $\text{IME} > 0$ then $\text{corr}(\ln \tilde{F}_{ij}, \ln \tilde{\tau}_{ij}) < 0$.

Ignoring origin and destination fixed effects, equation (10) implies that

$$\text{cov}(\ln \tilde{F}_{ij}, \ln \widetilde{\text{dist}}_{ij}) = \text{cov}(\ln \tilde{x}_{ij}, \ln \widetilde{\text{dist}}_{ij}).$$

Thus, if average exports per firm fall with distance then fixed trade costs must also fall with distance. This is captured formally by our third observation which is related to the fixed trade costs elasticity with respect to distance.

Observation 3: $\frac{\text{cov}(\ln \tilde{x}_{ij}, \ln \widetilde{\text{dist}}_{ij})}{\text{var}(\ln \widetilde{\text{dist}}_{ij})} = \frac{\text{cov}(\ln \tilde{F}_{ij}, \ln \widetilde{\text{dist}}_{ij})}{\text{var}(\ln \widetilde{\text{dist}}_{ij})}$.

We can go beyond the previous qualitative observations and derive the fixed and variable trade costs implied by the model to compute values for $\text{corr}(\ln \tilde{F}_{ij}, \ln \tilde{\tau}_{ij})$ and $\frac{\text{cov}(\ln \tilde{F}_{ij}, \ln \widetilde{\text{dist}}_{ij})}{\text{var}(\ln \widetilde{\text{dist}}_{ij})}$. Combining equations (9) and (10) to solve for $\ln \tilde{F}_{ij}$ and $\ln \tilde{\tau}_{ij}$ in terms of $\ln x_{ij}$ and $\ln N_{ij}$ yields

$$\ln \tilde{F}_{ij} = \delta_i^{F,o} + \delta_j^{F,d} + \ln x_{ij} \quad (12)$$

and

$$\theta \ln \tilde{\tau}_{ij} = \delta_i^{\tau,o} + \delta_j^{\tau,d} - \bar{\theta} \ln x_{ij} - \ln N_{ij}. \quad (13)$$

Model-implied values for $\ln \tilde{F}_{ij}$ are (ignoring origin and destination fixed effects) directly given by $\ln x_{ij}$, but for $\ln \tilde{\tau}_{ij}$ a value for $\bar{\theta}$ is required to go from $\ln x_{ij}$ and $\ln N_{ij}$ in the data to model-implied values for $\theta \ln \tilde{\tau}_{ij}$.

Exports 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. This implies that the intensive margin elasticity calculated separately for each exporter size percentile is the same as the overall intensive margin elasticity.

Observation 4: $\text{IME}^{pct} = \text{IME}$, for all pct .

Data

We now use Observations 1 – 4 above to relate the simple Melitz-Pareto model to the data as described in Section 2.

Observation 1 indicates that if fixed trade costs vary by origin and destination but not across country pairs, i.e., $var(\tilde{F}_{ij}) = 0$, then the IME should be equal to zero while the EME should be equal to one. This is captured in Figures 1 and 2 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, both in Figures 1 and 2 and in Tables 2, 3, and 4, which reveal an IME of 0.4 or higher.

For the simple Melitz-Pareto model to be consistent with the data, we need to move away from $var(\tilde{F}_{ij}) = 0$. As per Observation 2, however, the positive IME seen in the data implies a *negative* correlation between model-implied fixed and variable trade costs. Moreover, Observation 3 combined with the result in Table 6 of a negative distance elasticity of average exports per firm implies that model-implied fixed trade costs *fall* with distance.

We explore these results further by using equations (12) and (13) to compute model-implied fixed and variable trade costs.¹⁹ The correlation between the resulting fixed and variable trade costs is -0.786 (with a standard error of 0.007). Figure 6 plots these trade costs against distance. The Figure shows that model-implied fixed trade costs are decreasing with distance, while model-implied variable trade costs are increasing with distance.²⁰ The distance elasticities corresponding to Figure 6 are reported in Table 7. For fixed trade costs this elasticity is -0.28 (as per Observation 3, this is equal to the distance elasticity of average exports reported in Table 6) while for variable trade costs the distance elasticity is 0.272 , both statistically significant.

Finally, according to Observation 4, the simple Melitz-Pareto model implies that $IME^{pct} = IME$ for all *pct*. This theoretical prediction of a common elasticity across percentiles is captured by the horizontal line red in Figure 3. This is at odds with the data.

To conclude, the simple version of the Melitz-Pareto model with fixed trade costs varying only because of origin and destination fixed effects is clearly at odds with the data. One can of course allow a richer pattern of variation in fixed trade costs across country pairs to make the model perfectly consistent with the data, but then the positive IME has further puzzling implications for fixed trade costs, which should fall with distance and be very negatively correlated with variable trade costs. To the best of our knowledge, there are no models that would microfound

¹⁹To compute model-implied variable trade costs as in equation (13), values for θ and $\bar{\theta}$ are required. We set $\theta = 5$ from Head and Mayer (2014) and $\sigma = 5$ from Bas et al. (2015), which jointly imply $\bar{\theta} = 1.25$.

²⁰Variable trade costs must increase with distance so that total exports fall with distance, as implied by the results in Table 6.

such a strong and negative correlation between the two types of trade costs and a negative fixed trade costs elasticity with respect to distance.²¹ The data is also at odds with the implication from the Melitz-Pareto model of a constant IME across exporter size percentiles.

3.2. Multi-Product Extension of Melitz-Pareto

In this section we explore whether the puzzling implications for trade costs arising from the Melitz-Pareto model can be avoided by extending the model to multi-product firms. The idea would be that 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 an intensive margin. We will see, however, that 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, so that Observation 1 in the basic model will remain valid in this extension.

Theory

We consider an extension of the Melitz-Pareto model due to Bernard, Redding and Schott (2011). Each firm can produce a differentiated variety of each of a continuum of products in the interval $[0,1]$ with productivity $\varphi\lambda$, where φ is common across products and λ is product-specific. The firm component φ is drawn from a Pareto distribution $G^f(\varphi)$ with shape parameter θ^f , while the firm-product component λ is drawn from a Pareto distribution $G^p(\lambda)$ with shape parameter θ^p . To have well-defined terms given a continuum of firms, we impose $\theta^f > \theta^p > \sigma - 1$. To sell any products in market j , firms from country i have to pay a fixed cost F_{ij} , and to sell each individual product requires an additional fixed cost of f_{ij} . Variable trade costs are still τ_{ij} .

The cutoff λ for a firm from country i with productivity φ that wants to export to market j , $\lambda_{ij}^*(\varphi)$, is given implicitly by

$$A_j \left(\frac{w_i \tau_{ij}}{\varphi \lambda_{ij}^*(\varphi)} \right)^{1-\sigma} = \sigma f_{ij}. \quad (14)$$

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 the online appendix of Costinot and Rodríguez-Clare (2014), Felbermayr et al. (2015), and Caliendo et al. (2015).

We can then write the profits in market j for a firm from country i with productivity φ as

$$\pi_{ij}(\varphi) \equiv \int_{\lambda_{ij}^*(\varphi)}^{\infty} \left[\left(\frac{\lambda}{\lambda_{ij}^*(\varphi)} \right)^{\sigma-1} - 1 \right] f_{ij} dG^p(\lambda). \quad (15)$$

The cutoff productivity for firms from i to sell in j is implicitly $\pi_{ij}(\varphi_{ij}^*) = F_{ij}$. As in the canonical model, the number of firms from country i that export to market j is $N_{ij} = \left[1 - G^f(\varphi_{ij}^*) \right] N_i$, while the number of products sold by firms from i in j is $M_{ij} = N_i \int_{\varphi_{ij}^*}^{\infty} \left[1 - G^p(\lambda_{ij}^*(\varphi)) \right] dG^f(\varphi)$. Combining the previous expressions, using the fact that $G^p(\lambda)$ and $G^f(\varphi)$ are Pareto, writing $f_{ij} = f_i^o f_j^d \tilde{f}_{ij}$, $F_{ij} = F_i^o F_j^d \tilde{F}_{ij}$, and $\tau_{ij} = \tau_i^o \tau_j^d \tilde{\tau}_{ij}$, and defining variables appropriately we get

$$\ln X_{ij} = \mu_i^{X,o} + \mu_j^{X,d} - \theta^f \ln \tilde{\tau}_{ij} - \left(\frac{\theta^f}{\sigma-1} - \frac{\theta^f}{\theta^p} \right) \ln \tilde{f}_{ij} - \left(\frac{\theta^f}{\theta^p} - 1 \right) \ln \tilde{F}_{ij}, \quad (16)$$

$$\ln x_{ij}^p \equiv \ln X_{ij} - \ln M_{ij} = \mu_i^{x^p,o} + \mu_j^{x^p,d} + \ln \tilde{f}_{ij}, \quad (17)$$

$$\ln x_{ij} \equiv \ln X_{ij} - \ln N_{ij} = \mu_i^{x^f,d} + \mu_j^{x^f,d} + \ln \tilde{F}_{ij}. \quad (18)$$

It is easy to verify that if $f_{ij} = 0$ for all i, j then this model collapses to the canonical model with single-product firms.

Recalling our definition of the intensive margin elasticity at the firm and product level introduced in Section 2 and letting $\bar{\theta} \equiv \theta^f / (\sigma - 1)$ and $\chi \equiv \theta^f / \theta^p$, then from equations (16) to (18) we have

$$\text{IME} = - \frac{(\chi - 1) \text{var}(\ln \tilde{F}_{ij}) + (\bar{\theta} - \chi) \text{cov}(\ln \tilde{f}_{ij}, \ln \tilde{F}_{ij}) + \theta \text{cov}(\ln \tilde{F}_{ij}, \ln \tilde{\tau}_{ij})}{\text{var}(\ln \tilde{X}_{ij})} \quad (19)$$

and

$$\text{IME}^p = - \frac{(\bar{\theta} - \chi) \text{var}(\ln \tilde{f}_{ij}) + (\chi - 1) \text{cov}(\ln \tilde{f}_{ij}, \ln \tilde{F}_{ij}) + \theta \text{cov}(\ln \tilde{f}_{ij}, \ln \tilde{\tau}_{ij})}{\text{var}(\ln \tilde{X}_{ij})}. \quad (20)$$

Observation 1 in the single-product firm model remains valid in the multi-product firm model, while we now have an analogous observation for the product-level intensive margin elasticity:

Observation 5: If $\text{var}(\ln \tilde{f}_{ij}) = 0$ then $\text{IME}^p = 0$.

The assumption $\theta^f > \theta^p > \sigma - 1$ implies that $\chi > 1$ and $\bar{\theta} > \chi > 1$ and in turn this leads to the following extensions of observation 2:

Observation 6: If $\text{IME} > 0$ then either $\text{cov}(\ln \tilde{f}_{ij}, \ln \tilde{F}_{ij}) < 0$ or $\text{cov}(\ln \tilde{F}_{ij}, \ln \tilde{\tau}_{ij}) < 0$ (or both).

Observation 7: If $\text{IME}^p > 0$ then either $\text{cov}(\ln \tilde{f}_{ij}, \ln \tilde{F}_{ij}) < 0$ or $\text{cov}(\ln \tilde{f}_{ij}, \ln \tilde{\tau}_{ij}) < 0$ (or both).

Observation 3 remains valid in the multi-product firm model, and we now also have an analogous observation for product-level fixed trade costs:

Observation 8:
$$\frac{\text{cov}(\ln \tilde{x}_{ij}^p, \ln \widetilde{\text{dist}}_{ij})}{\text{var}(\ln \widetilde{\text{dist}}_{ij})} = \frac{\text{cov}(\ln \tilde{f}_{ij}, \ln \widetilde{\text{dist}}_{ij})}{\text{var}(\ln \widetilde{\text{dist}}_{ij})}.$$

As in the single-product case, we can use the model to back out the implied trade costs. Equation (18) can be used to obtain a model-implied \tilde{F}_{ij} (which would be the same as the one derived in the single-product model) while Equation (17) can be used to obtain a model-implied \tilde{f}_{ij} , and Equation (16) can then be used to obtain a model-implied $\tilde{\tau}_{ij}$.

Data

Since Observations 1 and 3 remain valid when the basic model is extended to allow for multi-product firms, the conclusions regarding the necessity of having fixed trade costs decrease with distance remain valid. Turning to the implications for product-level fixed trade costs, the finding in Section 2 of a positive IME at the product level, $\text{IME}^p > 0$ in Table 5, combined with Observation 5 implies that, to be consistent with the data, the multi-product version of the Melitz-Pareto model presented above requires $\text{var}(\ln \tilde{f}_{ij}) > 0$. However, observations 6 and 7 imply that the two types of fixed trade costs would need to be negatively correlated, or that the covariances between those fixed trade costs and variable trade costs would have to be negative. Moreover, Observation 8 combined with the results in Table 6 implies that model-implied product-level fixed trade costs decrease with distance with an elasticity of -0.071 , as shown in the third column of Table 7 and illustrated in Figure 7. We conclude that the puzzling implications of the Melitz-Pareto model remain valid when the model is extended to allow for multi-product firms.

3.3. Firm-Level Demand and Fixed-Cost Shocks

EKK extend the basic Melitz-Pareto model presented in Section 3.1 to allow for (log-normally distributed) firm-level destination-specific demand and fixed-cost shocks. Except for constants that capture the net effects of these shocks, our equations (7) and (8) remain valid in the EKK environment, and hence so do observations 1-3.²²

It is important to note, however, that if productivity is distributed Pareto then the presence of log-normally distributed demand or fixed-cost shocks would imply that equations (7) and (8) no longer hold. The critical assumption in EKK that allows their model to be consistent with our equations (7) and (8) is that, loosely speaking, they consider the limit as the scale parameter of the Pareto distribution converges to zero.²³

²²This can be confirmed by simple manipulation of equations (20) and (28) in EKK.

²³More exactly, EKK specify a function for the measure of firms with productivity above some level, with that measure going to infinity as productivity goes to zero. This is equivalent to taking a limit with the (exogenous) measure of firms going to infinity and the scale parameter of the Pareto distribution going to zero. Although equations (7) and (8) do not hold anywhere in this sequence, they do hold in the limit.

To formally establish this result, recall that to get equations (7) and (8) we assumed that $\varphi_{ij}^* > b_i$. If instead $\varphi_{ij}^* \leq b_i$ then $N_{ij} = N_i$ and $x_{ij} = \left(\frac{\theta}{\theta - (\sigma - 1)}\right) A_j \left(\frac{w_i \tau_{ij}}{b_i}\right)^{1 - \sigma}$. In the extreme, if $\varphi_{ij}^* \leq b_i$ holds for all i, j pairs, then we would have $\text{IME} = 1$ rather than $\text{IME} = 0$. Now think about the case with firm-specific demand and fixed-cost shocks. Specifically, assume that each firm is characterized by a productivity level φ as well as a demand shock α_j and a fixed cost shock f_j in each destination j , with φ drawn from a Pareto distribution (with scale parameter b_i and shape parameter θ) and α_j and f_j drawn iid from some distribution. Let $x_{ij}(\varphi, \alpha_j) \equiv A_j \alpha_j \left(\bar{\sigma} \frac{w_i \tau_{ij}}{\varphi}\right)^{(1 - \sigma)}$ and let $\varphi_{ij}^*(\alpha_j, f_j)$ be implicitly defined by $x_{ij}(\varphi_{ij}^*, \alpha_j) = \sigma f_j$. By the same argument we used in Section 3.1, if for all i, j and all possible (α_j, f_j) we have $\varphi_{ij}^*(\alpha_j, f_j) > b_i$, we can easily show that we still have $\text{IME} = 0$.²⁴ However, if α_j and f_j are lognormally distributed, then for $b_i > 0$ for all i there must be a positive mass of firms for which $\varphi_{ij}^*(\alpha_j, f_j) < b_i$, and for those firms there would be a positive intensive margin elasticity. EKK essentially avoid this by taking the limit with $b_i \rightarrow 0$ for all i .

In principle, one could use this result to argue that a Melitz model with Pareto distributed productivity but extended to allow for log-normally distributed demand and fixed-cost shocks could match the positive IME that we see in the data. However, such a model would not exhibit any of the convenient features of the canonical Melitz-Pareto model: the sales distribution is not distributed Pareto, the trade elasticity is not common across country pairs and fixed, and the gains from trade are not given by the ACR formula. Given that, our approach in this paper is to move all the way to a model where productivity as well as destination-specific demand and fixed-cost shocks are lognormally distributed. Such a model has at least the advantage that it is computationally tractable, and amenable to likelihood estimation methods, as we show in Section 4.

3.4. Granularity

The previous sections have considered a model with a continuum of firms. With a discrete and finite number of firms it may be possible to generate a positive covariance between the intensive margin and total exports that could in principle explain our empirical findings. To state this formally, we rely on the extension of the Melitz-Pareto model to allow for granularity in Eaton et

²⁴Consider the group of firms from country i that have some given draw $\{(\alpha_j, f_j), j = 1, \dots, n\}$. The exact same argument used in Section 3.1 can be used to show that the sample of firms obtained by combining such firms across all origins i satisfies $\text{IME} = 0$. One can then simply integrate across all possible draws $\{(\alpha_j, f_j), j = 1, \dots, n\}$ to show that $\text{IME} = 0$ for the whole set of firms.

al. (2012b). Equations (9) and (10) now become

$$\ln N_{ij} = \mu_i^{N,o} + \mu_j^{N,d} - \theta \ln \tilde{\tau}_{ij} - \bar{\theta} \ln \tilde{F}_{ij} + \xi_{ij} \quad (21)$$

and

$$\ln x_{ij} = \mu_i^{x,o} + \mu_j^{x,d} + \ln \tilde{F}_{ij} + \varepsilon_{ij}, \quad (22)$$

where ξ_{ij} and ε_{ij} are error terms arising from the fact that now the number of firms is discrete and random. Using the same definition for the intensive margin elasticity as in Section 3, the previous equations imply that

$$\text{IME} = \frac{-(\bar{\theta} - 1) \text{var}(\ln \tilde{F}_{ij}) - \theta \text{cov}(\ln \tilde{\tau}_{ij}, \ln \tilde{F}_{ij}) + \text{var}(\varepsilon_{ij}) + \text{COV}}{\text{var}(-\theta \ln \tilde{\tau}_{ij} - (\bar{\theta} - 1) \ln \tilde{F}_{ij} + \varepsilon_{ij} + \xi_{ij})}, \quad (23)$$

where $\text{COV} \equiv \text{cov}(\ln \tilde{F}_{ij} + \varepsilon_{ij}, \xi_{ij}) + \text{cov}(\ln \tilde{F}_{ij} + \ln \tilde{\tau}_{ij}, \varepsilon_{ij})$. If $\text{var}(\varepsilon_{ij})$ is large relative to $-\text{COV}$, this could explain $\text{IME} > 0$ even with $\text{cov}(\ln \tilde{F}_{ij}, \ln \tilde{\tau}_{ij}) > 0$. Thus, in theory, granularity could explain the positive intensive margin elasticity that we find in the data without relying on implausible patterns for fixed trade costs.

To check whether granularity is a plausible explanation for the positive IME in the data we conduct two tests, whose details are show in the Online Appendix B. First, we estimate the elasticities of model-implied firm-level and product-level fixed trade costs with respect to distance taking into account granularity. The results shown in Table 8 imply that although the distance elasticities are significantly lower than those estimated ignoring granularity in Table 7, they remain negative, implying that model-implied fixed trade costs are decreasing with distance. Hence granularity does not help to eliminate one of the puzzles emerging from the comparison between the Melitz-Pareto model and the data.

Second, to assess how well granularity can explain a positive IME, we simulate exports of N_{ij} firms for each of the country pairs in the sample. To study the IME generated by granularity by itself, we assume that there is no country-pair variation in fixed-trade costs, $\text{var}(\tilde{F}_{ij}) = 0$. In addition to the standard Melitz model, where firm sales are perfectly correlated across destination markets, we also allow for a case in which all heterogeneity comes from destination specific demand shocks, so that sales are independent across destinations. Finally, we simulate productivities for 3 values of $\bar{\theta}$: an estimate of θ using the procedure in Eaton et al. (2011), as outlined in the Online Appendix, A, which yields $\bar{\theta} = 2.4$; the value that can be inferred from standard estimates of θ and σ in the literature (i.e., $\theta = 5$, the central estimate of the trade elasticity in Head

and Mayer (2014), and $\sigma = 5$ from Bas et al. (2015), so $\bar{\theta} = 1.25$ (see footnote 19); and finally and $\bar{\theta} = 1$ (as in Zipf's Law).

Table 9 reports the results of this simulation exercise. Two broad patterns emerge from the table. First, the simulated IME decreases with $\bar{\theta}$. This is because the effect of granularity on the IME is stronger when there is more dispersion in productivity levels. Second, the simulated IME is highest when productivity is less correlated across destinations, again because this gives granularity more room to generate a covariance between average exports per firm and total exports. For our estimate of $\bar{\theta}$ ($\bar{\theta} = 2.4$) and with no demand shocks (so there is perfect correlation in firm-level productivity across destinations), the simulated IME of 0.001 is quite low. The highest simulated IME occurs for the case in which $\bar{\theta} = 1$ and there is no correlation between the product of demand shocks and productivity across destinations. In this case the simulated IME is 0.33, not too far from our preferred estimate based on the data of 0.4. But we think of this as an extreme case because $\bar{\theta} = 1$ is far from the estimates that come out of trade data, and because of the implausible assumption that firm-level exports are completely uncorrelated across destinations. A low $\bar{\theta}$ also implies that, in contrast to the data, virtually all of the action behind a positive IME comes from the superstar firms. To see that, we calculate average simulated exports per firm in each percentile and use those to estimate an IME per percentile. We plot the resulting 100 IME estimates in Figure 8 along with the corresponding IME estimates based on the actual data. The IME based on the actual data is increasing with a spike at the top percentile. Granularity and the Pareto distribution fail to reproduce this pattern in the simulated data, since the corresponding IME is much smaller than in the data for most percentiles. The IME in the simulated data is almost zero for small percentiles and is relatively high for a small number of top percentiles. We conclude that granularity does not offer a plausible explanation for the positive estimated IME in the data.

4. The Intensive Margin in the Melitz-Lognormal Model

In this section we depart from the assumption of a common Pareto distribution of firm-level productivity and instead assume a lognormal distribution.²⁵ In the theory section we start by showing how this can lead to a positive IME in a simple Melitz model, and then propose a maximum-

²⁵One could consider combining a lognormal distribution with a Pareto distribution on the right tail, as in Nigai (2017). We have used Nigai's Matlab code on our data to estimate the point of truncation (percentile) where the lognormal ends and the Pareto begins. We find that for 75% of country pairs with more than a hundred exporters the point of truncation occurs after the 99th percentile, and for the median country pair the truncation point is at the 99.9%. In light of these results, in the rest of the paper we focus on the case in which productivity is described by a fully lognormal distribution.

likelihood estimation procedure for a richer Melitz model with heterogeneous fixed costs and demand shocks. The data section presents the results from the estimation and the implications for the IME as well as for the model-implied trade costs.

We also develop a flexible methodology to calculate counterfactual changes in trade flows in response to a change in trade costs. We improve Dekle et al. (2008) “exact hat algebra” approach that can accommodate any distribution of productivity, demand, and fixed cost shocks. Using this approach we emphasize important differences in counterfactual trade flows responses in the Melitz-Pareto and lognormal models.

4.1. Theory

Melitz model with a lognormal distribution

Consider a model as that presented in Section 2 but with general CDF and PDF $G_i(\varphi)$ and $g_i(\varphi)$. The ratio of average to minimum exports per firm for each country pair can be written as

$$\frac{x_{ij}}{x_{ij}(\varphi_{ij}^*)} = \left(\frac{\tilde{\varphi}_i(\varphi_{ij}^*)}{\varphi_{ij}^*} \right)^{\sigma-1}, \quad (24)$$

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

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

If firm productivity is distributed Pareto with parameter $\theta > \sigma - 1$ (as in Section 2) then $\frac{\tilde{\varphi}_i(\varphi_{ij}^*)}{\varphi_{ij}^*} = \frac{\theta}{\theta - (\sigma - 1)}$ for all φ_{ij}^* , (see equation 8), so that the average to minimum ratio does not depend on selection. This property only holds with a Pareto distribution of productivity.

As argued in footnote 15 of Melitz (2003), $\frac{\tilde{\varphi}_i(\varphi_{ij}^*)}{\varphi_{ij}^*}$ is decreasing in the productivity cutoff 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 at 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. If productivity is distributed in such a way that $\frac{\tilde{\varphi}_i(\varphi_{ij}^*)}{\varphi_{ij}^*}$ is decreasing then the second effect does not fully offset the

first, and hence average exports per firm increase with a decline in τ_{ij} . Since this also increases the number of firms that export (and hence total exports), the result is a positive IME.

We now explore the magnitude of the implied IME in the case where $g_i(\varphi)$ is lognormal for every origin country i . Assuming that $g_i(\varphi)$ is lognormal with location parameter $\mu_{\varphi,i}$ and scale parameter σ_φ , and letting $\Phi(\cdot)$ be the CDF of the standard normal distribution, then

$$G_i(\varphi) = \Phi\left(\frac{\ln \varphi - \mu_{\varphi,i}}{\sigma_\varphi}\right). \quad (25)$$

Letting $h(x) \equiv \Phi'(x)/\Phi(x)$ be the ratio of the PDF to the CDF of the standard normal, Bas et al. (2015) (henceforth BMT), show that

$$\left(\frac{\tilde{\varphi}_i(\varphi_{ij}^*)}{\varphi_{ij}^*}\right)^{\sigma-1} = \frac{h\left[-(\ln \varphi_{ij}^* - \mu_{\varphi,i})/\sigma_\varphi\right]}{h\left[-(\ln \varphi_{ij}^* - \mu_{\varphi,i})/\sigma_\varphi + \bar{\sigma}_\varphi\right]}, \quad (26)$$

where $\bar{\sigma}_\varphi \equiv (\sigma - 1)\sigma_\varphi$. Combined with $1 - G_i(\varphi_{ij}^*) = N_{ij}/N_i$, we have

$$\frac{x_{ij}}{x_{ij}(\varphi_{ij}^*)} = \Omega\left(\frac{N_{ij}}{N_i}\right) \equiv \frac{h\left(\Phi^{-1}\left(\frac{N_{ij}}{N_i}\right)\right)}{h\left(\Phi^{-1}\left(\frac{N_{ij}}{N_i}\right) + \bar{\sigma}_\varphi\right)}. \quad (27)$$

Thus, the average to minimum ratio of exports per firm for country pair ij only depends on the share of total firms in i that export to j , with the relationship given by the function $\Omega(\cdot)$. As argued by BMT, $\Omega(\cdot)$ is an increasing function, which is consistent with the observation by Melitz (2003) above that $\frac{\tilde{\varphi}_i(\varphi_{ij}^*)}{\varphi_{ij}^*}$ is decreasing in φ_{ij}^* since N_{ij}/N_i is decreasing in φ_{ij}^* .

Given values of $\bar{\sigma}_\varphi$ as well as N_i for every country, we can use our data on N_{ij} to compute $\Omega(N_{ij}/N_i)$ for all country pairs. Combined with $x_{ij}(\varphi_{ij}^*) = \sigma F_{ij}$ and imposing $F_{ij} = F_i^o F_j^d$, we can use equation (27) to get the model-implied average exports per firm (in logs),

$$\ln x_{ij} = \mu_i^{x,o} + \mu_j^{x,d} + \ln \Omega(N_{ij}/N_i). \quad (28)$$

In contrast to Observation 1 for the Melitz-Pareto model, under under lognormality we will have a positive IME even with $\text{var}(\tilde{F}_{ij}) = 0$.

We can also compute model-implied fixed and variable trade costs similarly to what we did under the assumption of Pareto-distributed productivity. First, we obtain \tilde{F}_{ij} from

$$\ln \tilde{F}_{ij} = \delta_i^{F,o} + \delta_j^{F,d} + \ln x_{ij} - \ln \Omega\left(\frac{N_{ij}}{N_i}\right). \quad (29)$$

Second, to compute $\tilde{\tau}_{ij}$, we combine equations (3), (4), (25) and (27) to get (with appropriately defined fixed effects)

$$(\sigma - 1) \ln \tilde{\tau}_{ij} = \delta_i^{\tau,o} + \delta_j^{\tau,d} - \ln x_{ij} + \ln \Omega_i \left(\frac{N_{ij}}{N_i} \right) + \bar{\sigma}_\varphi \Phi^{-1} \left(1 - \frac{N_{ij}}{N_i} \right). \quad (30)$$

Armed with estimates of \tilde{F}_{ij} and $(\sigma - 1)\tilde{\tau}_{ij}$, we can compute their correlation and check whether \tilde{F}_{ij} increases or decreases with distance (demeaned by origin and destination fixed effects).

These empirical exercises require estimates for $\bar{\sigma}_\varphi$ as well as N_i for every country. We use Bento and Restuccia (2015) (henceforth BR) data to estimate a value for N_i for all the countries in our sample.²⁶ 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 for the number of existing firms, which in theory would correspond to $(1 - G_i(\varphi_{ii}^*)) N_i$ (our approach in the next subsection avoids this problem). We use the QQ-estimation proposed by Head et al. (2014) (henceforth HMT) to obtain estimates of σ_φ and $\mu_{\varphi,i}$ for every i (see Online Appendix C for a detailed description).

Full Melitz-lognormal model

The previous section has shown that a model with a lognormal distribution of firm productivity is capable of generating a positive intensive margin elasticity conditional on fixed costs. However, the model we considered had two very stark predictions. First, fixed trade costs that are common across firms lead to the prediction that sales of the least productive exporter from i to j are equal to σF_{ij} . In the data we observe many firms with very small export sales (sometimes as low as \$1) which implies unrealistic fixed trade costs. Second, as shown by Eaton et al. (2011), the model implies a perfect hierarchy of destination markets (i.e., destinations can be ranked according to profitability, with all firms that sell to a destination also selling to more profitable destinations) and perfect correlation of sales across firms that sell to multiple markets from one origin. None of these predictions holds in the data.

In this section we consider a richer model with firm-specific fixed trade costs and demand shocks that vary by destination. This is similar to the setup in Eaton et al. (2011). We assume that firm productivity, demand shocks (denoted by α_j) and fixed trade costs (denoted by f_j) are

²⁶Using census data as well as numerous surveys and registry data, BR compiled a dataset with the number of manufacturing firms for a set of countries. Unfortunately, the sample in BR 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 the BR dataset and log population with an elasticity of 0.945, as reported in Table 10 and in Figure 9.

distributed jointly lognormal, i.e., for each origin i :

$$\begin{bmatrix} \ln \varphi \\ \ln \alpha_1 \\ \vdots \\ \ln \alpha_J \\ \ln f_1 \\ \vdots \\ \ln f_J \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \mu_{\varphi,i} \\ \mu_{\alpha} \\ \vdots \\ \mu_{\alpha} \\ \mu_{f,i1} \\ \vdots \\ \mu_{f,iJ} \end{bmatrix}, \begin{bmatrix} \sigma_{\varphi,i}^2 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & \sigma_{\alpha,i}^2 & \dots & 0 & \sigma_{\alpha f,i} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_{\alpha,i}^2 & 0 & \dots & \sigma_{\alpha f,i} \\ 0 & \sigma_{\alpha f,i} & \dots & 0 & \sigma_{f,i}^2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_{\alpha f,i} & 0 & \dots & \sigma_{f,i}^2 \end{bmatrix} \right). \quad (31)$$

Note that we allow mean log productivity to be origin-specific while imposing that the mean of demand shocks be the same across origin-destination pairs (however, we cannot separately identify these parameters). Mean fixed costs are allowed to vary across origin-destination pairs and can be correlated with demand shocks within destinations. In our empirical estimation we will not be able to separately identify mean productivity from wages and variable trade costs – they will all be absorbed into an origin-destination fixed effect. Also, we allow the dispersion of log productivity, log demand shocks and log fixed trade costs to differ across all origins. We restrict the dispersion of log demand shocks and log fixed trade costs to be the same across destinations within a given origin.

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 $g_{X_i}(x_i)$ 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 $D_{ij} \equiv \ln [A_j (w_i \tau_{ij})^{1-\sigma}]$ and let $Z_{ij} \equiv D_{ij} + \ln \alpha_j + (\sigma - 1) \ln \varphi$ 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

$$\begin{bmatrix} Z_{i1} \\ \vdots \\ Z_{iJ} \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} d_{i1} \\ \vdots \\ d_{iJ} \end{bmatrix}, \begin{bmatrix} \bar{\sigma}_{\varphi,i}^2 + \sigma_{\alpha,i}^2 & \cdots & \bar{\sigma}_{\varphi,i}^2 \\ \vdots & \ddots & \vdots \\ \bar{\sigma}_{\varphi,i}^2 & \cdots & \bar{\sigma}_{\varphi,i}^2 + \sigma_{\alpha,i}^2 \end{bmatrix} \right), \quad (32)$$

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

Using firm-level data from the EDD and China across different origins and destinations, we can estimate the parameters in (32) as well as mean log fixed trade costs (up to a constant) and their dispersion using maximum likelihood methods. The Online Appendix D 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 fifteen destinations (USA, Germany, Japan, France and the 11 largest destinations by exports value for each origin), which we label $j = 1, \dots, 15$ for the 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,

$$\ln L \left(\Theta_i | \{x_{i1}(k_i), \dots, x_{iJ}(k_i)\}_{i,k_i} \right) = \sum_{k_i=1}^{\tilde{N}_i} \ln [g_{(X_{i1}, \dots, X_{iJ})}(x_{i1}(k_i), \dots, x_{iJ}(k_i))], \quad (33)$$

where \tilde{N}_i is the number of firms from i that sell to either of the fifteen destinations we consider, and where 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,

$$\Theta_i = \left\{ \{d_{ij}, \bar{\mu}_{f,ij}\}_{i,j}, \bar{\sigma}_{\varphi,i}, \sigma_{\alpha,i}, \sigma_{f,i}, \rho_i \right\}, \quad (34)$$

where $\bar{\mu}_{f,ij} = \ln \sigma + \mu_{f,ij}$ and $\rho = \frac{\sigma_{\alpha f,i}}{\sigma_{\alpha,i} \sigma_{f,i}}$. As the likelihood is potentially not concave in θ_i and because there are 34 parameters per origin to estimate, we rely on the estimation methodology proposed by Chernozhukov and Hong (2003).²⁷ We use the Metropolis-Hastings MCMC algorithm to construct a chain of estimates $\Theta_i^{(n)}$ for each origin country. Chernozhukov and Hong

²⁷Roughly speaking, this 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 the Online Appendix E.

(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 the Online Appendix F.

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 correlation is negative, then a firm with a bad demand shock would also likely draw a high fixed cost shock and thus will not export, hence, we would not see a lot of small firms in the data. Finally, to understand how $\sigma_{f,i}$ is 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.

4.2. Data

Simple Melitz model with lognormal distribution

Online Appendix Table 19 reports the QQ-estimate of $\bar{\sigma}_\varphi$. We report three sets of estimates: 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 on the high side relative to the estimate obtained by HMT, so we will use the minimum among them, $\bar{\sigma}_\varphi = 4.02$, which corresponds to the subsample with the largest 25% of firms.²⁸ We highlight three findings. First, a lognormal distribution allows the intensive margin elasticity to be positive even under the assumption of a continuum of firms. Second, for our estimate of the shape parameter ($\bar{\sigma}_\varphi = 4.02$), the implied IME is 0.28, which is close to that from the data.²⁹ Third, most of the action comes from the right tail of the exporter size distribution, as seen in Figure 10.

We use equations (29) and (30) to compute the model-implied fixed and variable trade costs. The correlations between those costs and distance are reported in Table 11 and plotted in Figure 11. In contrast to our results under Pareto, now under lognormal both the model-implied variable and fixed trade costs are increasing with distance.

Overall, the model does much better in fitting the data when we assume that firm productivity is distributed lognormal than when we assume that it is distributed Pareto. However, the IME for each percentile is not a perfect match to the data, and there is still a negative correlation between the model implied variable and fixed trade costs, although it is much closer to zero than with Pareto (-0.3 rather than -0.8). In any case, this is just a "proof of concept" that lognormally-distributed productivity can by itself improve the performance of the model relative to the data. In the next subsection we present the results obtained with the estimated full Melitz-lognormal model.

Full Melitz-lognormal model

To estimate the parameters of the full Melitz-lognormal model we use firm-level data from the EDD and China for the year 2007.³⁰ This entails 39 different origins, but we had to drop 2 origins from our analysis due to convergence issues discussed in the Online Appendix F. As mentioned above, we consider 15 destinations: USA, Germany, France, Japan, and the 11 largest destinations by export value for each origin. Finally, we use the estimates of Bas et al. (2015) and set $\sigma = 5$.

Before presenting the results of the estimation and discussing their implications for the IME,

²⁸See the section in the Appendix titled QQ-Estimation of $\bar{\sigma}_\varphi$ for a discussion of these estimates and their relation to the estimate in HMT.

²⁹Using Head et al. (2014) estimate of $\bar{\sigma}_\varphi = 2.4$ we get IME of around 0.12.

³⁰For computational reasons, for China we considered only a random sample consisting of 5% of exporters.

we show three figures revealing the fit of the estimated model with the data. Figure 12 shows a plot of the density function for standardized firm-level log sales pooled across multiple origin and destinations.³¹ We see that 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. According to Figure 13, the share predicted by the estimated model is quite close to the one in the data.

Finally, Figure 14 shows, 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. The estimated model mostly implies positive correlation driven by firm-level productivity shocks, while in the data this correlation exhibits more dispersion.

We now turn to our estimates of the variance-covariance parameters $(\bar{\sigma}_{\varphi,i}, \sigma_{\alpha,i}, \sigma_{f,i}, \rho_i)$. These are shown in Table 12. The median estimated values for $\bar{\sigma}_{\varphi,i}$ and $\sigma_{\alpha,i}$ across 37 origins are 3.18 and 2.67 respectively, while the median estimates 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.³² In general, there is a positive correlation between demand and fixed costs shocks, but some origins exhibit a negative correlation.

Table 13 and Figure 15 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.³³ 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 the year 2007 we estimate IME of 0.67 (with a standard error of 0.03) that is statistically indistinguishable from the one implied by our estimated lognormal model.³⁴ We plot

³¹Standardized firm-level log sales are obtained by, for each origin-destination cell, subtracting the mean and dividing by the standard deviation.

³²The estimates and confidence bands for each of the parameters are reported in the Online Appendix F.

³³We pick one million because at this point we are not interested in granularity – this is just a numerical approximation to the case with a continuum of firms.

³⁴The confidence interval in Table (13) comes 1000 random realizations of the parameters in our Markov chains.

the associated IME for each percentile in Figure 15 – the pattern of the IME across percentiles is remarkably close to what we see in the data.

Table 14 shows the elasticity of 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 lognormal-Melitz 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 decrease with distance. The implied pattern for the IME across different percentiles is also very similar to what we see in the data.

5. Counterfactual Analysis

In this section we study whether the counterfactual implications of the Melitz-lognormal model estimated as in the previous section differ from those of the Melitz-Pareto model estimated as described below. 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 $\mu_{\varphi,i}$, μ_α and $\mu_{f,ij}$ denote the mean of $\ln \varphi_i$, $\ln \alpha$, and $\ln f_{ij}$, respectively, and let $\tilde{f}_{ij} \equiv \ln f_{ij} - \mu_{f,ij}$, $\tilde{\varphi}_{ij} \equiv (\sigma - 1)(\ln \varphi_i - \mu_{\varphi,i}) + \ln \alpha - \mu_\alpha$ and

$$A_{ij} \equiv \ln \left(\frac{(\bar{\sigma} w_i \tau_{ij})^{1-\sigma} P_j^{\sigma-1} X_j}{\sigma w_i} \right) + (\sigma - 1)\mu_{\varphi,i} + \mu_\alpha - \mu_{f,ij}.$$

Using this notation, firms from i export to j if and only if $\tilde{f}_{ij} \leq A_{ij} + \tilde{\varphi}_{ij}$. The fraction of country i firms that export to country j is then

$$n_{ij} \equiv \frac{N_{ij}}{N_i} = \int_{-\infty}^{+\infty} \int_{-\infty}^{A_{ij} + \tilde{\varphi}} g_{ij}(\tilde{\varphi}, \tilde{f}) d\tilde{f} d\tilde{\varphi},$$

where g_{ij} is the joint pdf of \tilde{f}_{ij} and $\tilde{\varphi}_{ij}$. The previous equation implicitly defines a function $\mathcal{H}_{ij}(\cdot)$ such that $A_{ij} = \mathcal{H}_{ij}(n_{ij})$. Combining it with our definition of A_{ij} above, we then have the following equilibrium condition:

$$\mathcal{H}_{ij}(n_{ij}) = \ln \left(\frac{(\bar{\sigma} w_i \tau_{ij})^{1-\sigma} P_j^{\sigma-1} X_j}{\sigma w_i} \right) + (\sigma - 1) \mu_{\varphi, i} + \mu_{\alpha} - \mu_{f, ij}. \quad (35)$$

In turn, the price index is

$$P_j^{1-\sigma} = \sum_k P_{kj}^{1-\sigma}, \quad (36)$$

with

$$P_{ij}^{1-\sigma} = N_i (\bar{\sigma} w_i \tau_{ij})^{1-\sigma} e^{(\sigma-1)\mu_{\varphi, i} + \mu_{\alpha}} \int_{-\infty}^{\infty} e^{\tilde{\varphi}} \int_{-\infty}^{\mathcal{H}_{ij}(n_{ij}) + \tilde{\varphi}} g_{ij}(\tilde{\varphi}, \tilde{f}) d\tilde{f} d\tilde{\varphi}. \quad (37)$$

Free entry implies that profits net of fixed costs of exporting are equal to entry costs. Profits gross of fixed costs are equal to $\frac{1}{\sigma} \sum_j \lambda_{ij} X_j$ and total fixed costs of exporting per destination are

$$N_i w_i e^{\mu_{f, ij}} \int_{-\infty}^{\infty} \int_{-\infty}^{\mathcal{H}_{ij}(n_{ij}) + \tilde{\varphi}} e^{\tilde{f}} g(\tilde{\varphi}, \tilde{f}) d\tilde{f} d\tilde{\varphi},$$

hence the free entry condition is

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

But using Equation 4 together with

$$\lambda_{ij} = \frac{P_{ij}^{1-\sigma}}{P_j^{1-\sigma}} \quad (38)$$

we get

$$e^{\mu_{f, ij}} = \frac{\lambda_{ij} X_j}{e^{\mathcal{H}_{ij}(n_{ij})} N_i \sigma w_i \int_{-\infty}^{\infty} e^{\tilde{\varphi}} \int_{-\infty}^{\mathcal{H}_{ij}(n_{ij}) + \tilde{\varphi}} g_{ij}(\tilde{\varphi}, \tilde{f}) d\tilde{f} d\tilde{\varphi}},$$

and so we can rewrite the free entry condition as

$$\sigma F^e w_i N_i = \sum_j \lambda_{ij} X_j \left(1 - \frac{\int_{-\infty}^{\infty} \int_{-\infty}^{\mathcal{H}_{ij}(n_{ij}) + \tilde{\varphi}} e^{\tilde{f}} g_{ij}(\tilde{\varphi}, \tilde{f}) d\tilde{f} d\tilde{\varphi}}{e^{\mathcal{H}_{ij}(n_{ij})} \int_{-\infty}^{\infty} e^{\tilde{\varphi}} \int_{-\infty}^{\mathcal{H}_{ij}(n_{ij}) + \tilde{\varphi}} g_{ij}(\tilde{\varphi}, \tilde{f}) d\tilde{f} d\tilde{\varphi}} \right). \quad (39)$$

An equilibrium is defined as variables $\{n_{ij}, \lambda_{ij}, P_{ij}\}$ and $\{X_j, P_j, w_i\}$ such that Equations (4) - (39) are satisfied for all i, j , and in addition

$$w_i L_i = \sum_j \lambda_{ij} X_j, \quad (40)$$

and

$$X_j = w_j L_j + \Delta_j \quad (41)$$

are satisfied for all i .

We now transform this system of equations in levels into one in hat changes, 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 changes in endogenous variables. Since n_{ij} always appears as an argument in the $\mathcal{H}_{ij}(\cdot)$ function, it is convenient to focus instead on $h_{ij} \equiv \mathcal{H}_{ij}(n_{ij})$. The system of equations in hat changes is then:

$$h_{ij}(\hat{h}_{ij} - 1) = \ln \left(\frac{(\hat{w}_i \hat{\tau}_{ij})^{1-\sigma} \hat{P}_j^{\sigma-1} \hat{X}_j}{\hat{w}_i} \right) \quad (42)$$

$$\hat{P}_j^{1-\sigma} = \sum_k \lambda_{kj} \hat{P}_{kj}^{1-\sigma} \quad (43)$$

$$\hat{P}_{ij}^{1-\sigma} = \hat{N}_i (\hat{w}_i \hat{\tau}_{ij})^{1-\sigma} \frac{\int_{-\infty}^{\infty} e^{\tilde{\varphi}} \int_{-\infty}^{\hat{h}_{ij} h_{ij} + \tilde{\varphi}} g_{ij}(\tilde{\varphi}, \tilde{f}) d\tilde{f} d\tilde{\varphi}}{\int_{-\infty}^{\infty} e^{\tilde{\varphi}} \int_{-\infty}^{h_{ij} + \tilde{\varphi}} g_{ij}(\tilde{\varphi}, \tilde{f}) d\tilde{f} d\tilde{\varphi}} \quad (44)$$

$$\hat{\lambda}_{ij} = \frac{\hat{P}_{ij}^{1-\sigma}}{\hat{P}_j^{1-\sigma}} \quad (45)$$

$$\begin{aligned} \hat{w}_i \hat{N}_i \sum_j \lambda_{ij} X_j \left(1 - \frac{\int_{-\infty}^{\infty} \int_{-\infty}^{h_{ij} + \tilde{\varphi}} e^{\tilde{f}} g_{ij}(\tilde{\varphi}, \tilde{f}) d\tilde{f} d\tilde{\varphi}}{e^{h_{ij}} \int_{-\infty}^{\infty} e^{\tilde{\varphi}} \int_{-\infty}^{h(n_{ij}) + \tilde{\varphi}} g_{ij}(\tilde{\varphi}, \tilde{f}) d\tilde{f} d\tilde{\varphi}} \right) \\ = \sum_j \lambda_{ij} X_j \hat{\lambda}_{ij} \hat{X}_j \left(1 - \frac{\int_{-\infty}^{\infty} \int_{-\infty}^{\hat{h}_{ij} h_{ij} + \tilde{\varphi}} e^{\tilde{f}} g_{ij}(\tilde{\varphi}, \tilde{f}) d\tilde{f} d\tilde{\varphi}}{e^{\hat{h}_{ij}} \int_{-\infty}^{\infty} e^{\tilde{\varphi}} \int_{-\infty}^{\hat{h}_{ij} h_{ij} + \tilde{\varphi}} g_{ij}(\tilde{\varphi}, \tilde{f}) d\tilde{f} d\tilde{\varphi}} \right) \end{aligned} \quad (46)$$

$$\hat{w}_i Y_i = \sum_j \lambda_{ij} X_j \hat{\lambda}_{ij} \hat{X}_j \quad (47)$$

$$\hat{X}_j X_j = \hat{w}_j Y_j + \hat{\Delta}_j (X_j - Y_j) \quad (48)$$

Here σ is a known parameter and g_{ij} are known functions, λ_{ij} , h_{ij} , X_j and Y_i are data, $\hat{\tau}_{ij}$ and $\hat{\Delta}_j$ reflect exogenous shocks, and \hat{h}_{ij} , $\hat{\lambda}_{ij}$, \hat{w}_i , \hat{N}_i , \hat{P}_j , \hat{P}_{ij} and \hat{X}_j are the endogenous variables that we solve for.

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

For our counterfactual analysis we need a set of countries for which we have λ_{ij} , h_{ij} , X_j and Y_i , as well as g_{ij} for all i and j in that set. Since we assume (see Section 4) that the variances of φ_i and f_{ij} differ by origin but not by destination, 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 then construct h_{ij} using data for N_{ij} and $h_{ij} \equiv \mathcal{H}_{ij}(N_{ij}/N_i)$. Finally, we also need X_{ij} and N_{ij} for $i = j$. Following the approach proposed by Ossa (2012), 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 lognormal model.³⁶ We do not consider the whole EDD dataset for computational reasons. This is because some of the country pairs in the whole EDD dataset trade very little and this makes our welfare calculations imprecise (since we need to use numerical approximation to compute some of the integrals).

We are interested in comparing the counterfactual implications of the full Melitz-lognormal model with those of the Melitz-Pareto model. We proceed in three steps. First, we allow for some change in trade costs and compute the counterfactual implications using the estimated full Melitz-lognormal model. We consider four different trade costs shocks: a 1%, a 5%, a 10%, and a 25% uniform reduction in international trade costs – formally, $\hat{\tau}_{ij} = \hat{\tau} \in \{0.99, 0.95, 0.9, 0.75\}$ if $i \neq j$, while $\hat{\tau}_{ii} = 1$. Second, we use these results to estimate the trade elasticity by running the OLS regression

$$\ln \hat{X}_{ij} = \gamma_i^o + \gamma_j^d - \theta \ln \hat{\tau}_{ij} + \zeta_{ij}. \quad (49)$$

This leads to four values of the trade elasticity, one for each $\hat{\tau}$. Following ACR, we use these values as estimates of the shape parameter of the Pareto distribution in the Melitz-Pareto model, which we denote by $\theta_{\hat{\tau}}$ for $\hat{\tau} \in \{0.99, 0.95, 0.9, 0.75\}$.³⁷ Finally, for each trade cost shock $\hat{\tau}$ we

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. Since we observe the total number of exporters to those destinations, we can infer the total number of firms from which they are drawn. The more detailed procedure is described in the Online Appendix D.

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

³⁷We get $\theta_{0.99} = 4.5$; $\theta_{0.95} = 4.49$; $\theta_{0.9} = 4.47$; $\theta_{0.75} = 4.43$.

compute the counterfactual implications in the Melitz-Pareto model using both $\theta_{0.99}$ (a local approximation) and $\theta_{\hat{\tau}}$ (an average trade elasticity computed for the actual trade cost shock). We expect $\theta_{0.99}$ to differ from $\theta_{\hat{\tau}}$ because in the true model (i.e., the full Melitz-lognormal model) the trade elasticity is not a constant, as emphasized by Melitz and Redding (2015).

We show the results of this exercise in Figures 16 and 17. We use $\hat{W}_i^m \equiv \hat{w}_i^m / \hat{P}_i^m$ and \hat{X}_{ij}^m to denote the hat changes in welfare and trade flows for the lognormal model ($m = LN$) and the Pareto model ($m = P$). Figure 16 plots $\hat{W}_i^{LN} - 1$ (horizontal axis) versus $\hat{W}_i^P - 1$ (vertical axis) in response to the four different trade costs shocks, in each case reporting $\hat{W}_i^P - 1$ computed with both $\theta_{0.99}$ and $\theta_{\hat{\tau}}$. It is evident that both models yield very similar results even when we only use the local approximation $\theta_{0.99}$ for the trade elasticity in the Melitz-Pareto model. 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. However, it turns out that the (average) trade elasticity (as implied by the regression in 49) does not vary much in the estimated Melitz-lognormal model as we move from one equilibrium to the other, as can be seen by noting that $\hat{\theta}_{0.99} = 4.5$ while $\hat{\theta}_{0.75} = 4.43$. We explore this further by using our estimated g_i and n_{ij} to compute the local trade elasticity for each country pair using the formula derived by Bas et al. (2015) for the Melitz-lognormal model. The resulting elasticity ranges from 4 to 6.9 with a standard deviation of 0.59, with the higher values occurring for country pairs with a low n_{ij} , as shown in the Online Appendix G. 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 is 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 Pareto model yields a good approximation for the gains from uniform trade liberalization, and this holds even when using the local trade elasticity $\theta_{0.99}$.

Even though gains from trade liberalization are very close across the two models, the Pareto and the lognormal models differ in their implications for the counterfactual changes in bilateral trade flows. Figure 17 plots the ratio of the difference between the counterfactual changes in trade flows in the lognormal model, \hat{X}_{ij}^{LN} , and in the Pareto model, \hat{X}_{ij}^P , against the trade elasticity implied by the lognormal model. We can see that Melitz-Pareto model can significantly over- or underpredict changes in trade flows depending on the actual trade elasticity. Naturally, the higher trade elasticity in the 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 highest 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 18. 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 regarding the effects on trade flows, as shown on Figure 19.

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 4.8. Note that after trade liberalization, elasticities can never fall below 4, and thus the local trade elasticities are a good approximation to the average trade elasticities during the liberalization episodes. We discuss this further in Appendix H, in which 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 productivities 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). There should be no variation along the intensive margin (exports per exporting firm), again conditional on fixed costs.

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% 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% of the high exports. This finding is robust to looking at all destinations or only the largest destinations, including all firms or ignoring very small firms, including all country pairs or only ones for which more than 100 firms export, and to disaggregating across industries. 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 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 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 costs. We estimate this model 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 by exporting firm percentile. This specification 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). However, extending exact the hat algebra approach popularized by Dekle et al. (2008) to our estimated full Melitz-lognormal model, 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. In Alvarez et al. (2014), Buera and Oberfield (2015), and Perla et al. (2015), trade liberalization boosts the level or growth rate of technology through such learning spillovers. The size of this dynamic gain should depend on whether the distribution of firm productivity is Pareto vs. 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

- Alvarez, Fernando, Francisco J. Buera, and Robert E. Lucas Jr., “Idea Flows, Economic Growth, and Trade,” 2014.
- Arkolakis, Costas, “Market Penetration Costs and the New Consumers Margin in International Trade,” *Journal of Political Economy*, 2010, 118 (6), 1151–1199.
- and Marc-Andreas Muendler, “Exporters and their products: a collection of empirical regularities,” *CESifo Economic Studies*, 2013, 59 (2), 223–248.
- , 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,” *Unpublished, MIT*, 2015.
- , Natalia Ramondo, Andrés Rodríguez-Clare, and Stephen Yeaple, “Innovation and Production in the Global Economy,” 2014.
- , Svetlana Demidova, Peter J Klenow, and Andres Rodriguez-Clare, “Endogenous Variety and the Gains from Trade,” *American Economic Review*, 2008, 98 (2), 444–50.
- Bas, Maria, Thierry Mayer, and Mathias Thoenig, “From Micro to Macro: Demand, Supply, and Heterogeneity in the Trade Elasticity,” 2015.
- Bento, Pedro and Diego Restuccia, “Misallocation, Establishment Size, and Productivity,” Technical Report 2015.
- Bernard, Andrew B, J Bradford Jensen, Stephen J Redding, and Peter K Schott, “Firms in international trade,” *The Journal of Economic Perspectives*, 2007, 21 (3), 105–130.
- , —, —, and —, “The margins of US trade,” *The American Economic Review*, 2009, 99 (2), 487–493.
- , Jonathan Eaton, J Bradford Jensen, and Samuel Kortum, “Plants and Productivity in International Trade,” *American economic review*, 2003, 93 (4), 1268–1290.
- , Stephen J Redding, and Peter K Schott, “Multiproduct Firms and Trade Liberalization*,” *The Quarterly journal of economics*, 2011, 126 (3), 1271–1318.
- Buera, Francisco J. and Ezra Oberfield, “The Global Diffusion of Ideas,” 2015.

- 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 ML, "Financial Frictions and Trade Intermediation: Theory and Evidence," 2017.
- Chaney, Thomas, "Distorted gravity: the intensive and extensive margins of international trade," *The 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.
- , Samuel S Kortum, and Sebastian Sotelo, "International trade: Linking micro and macro," 2012.
- , —, and —, "International trade: Linking micro and macro," Technical Report, National bureau of economic research 2012.
- Feenstra, Robert C, "Restoring the product variety and pro-competitive gains from trade with heterogeneous firms and bounded productivity," 2014.
- 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.
- Head, Keith and Thierry Mayer, “Gravity Equations: Workhorse, Toolkit, and Cookbook,” *Handbook of International Economics*, 2014, 4, 131.
- and — , “Poor Substitutes? Counterfactual Methods in IO and Trade compared.,” 2018.
- , — , and Mathias Thoenig, “Welfare and Trade without Pareto,” *American Economic Review*, 2014, 104 (5), 310–16.
- Krugman, Paul, “Scale economies, product differentiation, and the pattern of trade,” *The American Economic Review*, 1980, 70 (5), 950–959.
- Manova, Kalina and Zhiwei Zhang, “Export Prices Across Firms and Destinations,” *The Quarterly Journal of Economics*, 2012, 127 (1), 379–436.
- Mayer, Thierry and Soledad Zignago, Working Papers 2011-25, CEPII 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,” Working Paper 18113, National Bureau of Economic Research May 2012.
- Perla, Jesse, Christopher Tonetti, and Michael E. Waugh, “Equilibrium Technology Diffusion, Trade, and Growth,” 2015.

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	Krygyzstan	2006	2012	ZMB	Zambia	2003	2011

* indicates that 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]	[0.0049]
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]	[0.0022]
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 represents the estimated coefficients of the regression of log average exports on log total exports with year fixed effects (column 1), origin \times year fixed effects (column 2), origin \times year and destination \times year fixed effects (column 3). The data are aggregated at the year-origin-destination 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: IME regression, small firms excluded, core sample

	Coefficient from $\ln x_{ij}$ on $\ln X_{ij}$		
<i>Panel a: country pairs with $N_{ij} \geq 100$</i>			
IM elasticity	0.437***	0.459***	0.398***
Standard error	[0.0058]	[0.0042]	[0.0050]
R^2	0.54	0.74	0.85
Variation in $\ln X_{ij}$ explained by FE,%	0.01	0.19	0.59
Observations	7,698	7,684	7,234
<i>Panel b: all country pairs</i>			
IM elasticity	0.497***	0.525***	0.573***
Standard error	[0.0013]	[0.0013]	[0.0015]
R^2	0.77	0.81	0.84
Variation in $\ln X_{ij}$ explained by FE, %	0.00	0.19	0.50
Observations	46,925	46,925	46,832
Year FE	Yes		
Origin \times year FE		Yes	Yes
Destination \times year FE			Yes

Note: the table represents the estimated coefficients of the regression of log average exports on log total exports with year fixed effects (column 1), origin \times year fixed effects (column 2), origin \times year and destination \times year fixed effects (column 3). The data are aggregated at the year-origin-destination level for a set of origin-years listed in Table 1. Average and total exports per destination are calculated using the sales of firms with at least \$1000 to that destination. 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 4: IME regression, disaggregated within manufacturing, core sample

Coefficient from $\ln x_{ij}$ on $\ln X_{ij}$			
<i>Panel a: HS 2-digit</i>			
IM elasticity	0.569***	0.510***	0.467***
Standard error	[0.0022]	[0.0017]	[0.0049]
Observations	37,321	35,621	10,732
<i>Panel b: HS 4-digit</i>			
IM elasticity	0.651***	0.569***	0.515***
Standard error	[0.0019]	[0.0013]	[0.0069]
Observations	62,776	58,516	4,640
<i>Panel c: HS 6-digit</i>			
IM elasticity	0.664***	0.593***	0.508***
Standard error	[0.0020]	[0.0014]	[0.0094]
Observations	67,967	61,501	2,972
Year \times HS FE	Yes		
Origin \times Year \times HS FE		Yes	Yes
Destination \times Year \times HS FE			Yes

Note: the table represents the estimated coefficients of the regression of log average exports on log total exports with year fixed effects (column 1), origin \times year fixed effects (column 2), origin \times year and destination \times year fixed effects (column 3). The data are aggregated at the year-origin-destination-HS industry level for a set of origin-years listed in Table 1. Panels a), b), and c) represent the regressions for industries defined at the HS 2-digit, 4-digit, and 6-digit levels respectively. The sample is restricted to the origin-destination-product cells with at least 100 exporters. Robust standard errors are reported in brackets. *, **, and *** represent the 5%, 1%, and 0.1% significance levels respectively.

Table 5: Product-level IME regression, core sample

	Coefficient from $\ln x_{ij}^p$ on $\ln X_{ij}$		
IM elasticity	0.380***	0.397***	0.288***
Standard error	[0.0070]	[0.0054]	[0.0073]
R^2	0.35	0.62	0.78
Variation in $\ln X_{ij}$ explained by FE,%	0.01	0.20	0.59
Observations	7781	7,768	7,324
Year FE	Yes		
Origin \times year FE		Yes	Yes
Destination \times year FE			Yes

Note: the table represents the estimated coefficients of the regression of log average exports per product (average exports divided by the number of HS6 products exported by all firms from origin i to destination j in a given year) on log total exports with year fixed effects (column 1), origin \times year fixed effects (column 2), origin \times year and destination \times year fixed effects (column 3). The data are aggregated at the year-origin-destination level for a set of origin-years listed in Table 1. The sample is restricted to the origin-destination pairs with at least 100 exporters. Robust standard errors are reported in brackets. *, **, and *** represent the 5%, 1%, and 0.1% significance levels respectively.

Table 6: Margins of trade and distance

Elasticity with respect to distance		
x_{ij}	0.123***	-0.280***
Standard error	[0.0150]	[0.0130]
N_{ij}	-0.416***	-1.010***
Standard error	[0.0134]	[0.0128]
x_{ij}^p	0.288***	-0.071***
Standard error	[0.0158]	[0.0146]
m_{ij}	-0.165***	-0.209***
Standard error	[0.0059]	[0.0051]
Observations	7,725	7,320
Origin \times year FE	Yes	Yes
Destination \times year FE		Yes

Note: the table represents the estimated coefficients of the regression of log average exports, number of firms, average exports per product, and number of products on log distance between origins and destinations with origin \times year fixed effects (column 2), origin \times year and destination \times year fixed effects (column 3). The data are aggregated at the year-origin-destination 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. Robust standard errors are reported in brackets. *, **, and *** represent the 5%, 1%, and 0.1% significance levels respectively.

Table 7: Trade costs and distance

	$\ln \tilde{F}_{ij}$	$\ln \tilde{\tau}_{ij}$	$\ln \tilde{f}_{ij}$
$\ln dist_{ij}$	-0.280***	0.272***	-0.071***
Standard error	[0.0140]	[0.0046]	[0.0146]
Observations	7,320	7,320	7,320

Note: the table represents the estimated coefficients of the regression of the implied log fixed firm-level trade costs (column 1), log variable trade costs (column 2), and log fixed product-level trade costs (column 3) on log distance between origins and destinations. We calculate trade costs as per equations 12 and 13 with $\theta = 5$. The data are aggregated at the year-origin-destination 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. Robust standard errors are reported in brackets. *, **, and *** represent the 5%, 1%, and 0.1% significance levels respectively.

Table 8: Fixed trade costs distance elasticity and granularity

	Fixed trade costs elasticity	
	Firm level	Product level
ζ	-0.022***	-0.007**
Standard error	[0.0029]	[0.0026]
Observations	7,320	7,320

Note: the table represents the estimated coefficients of the regression of the implied log fixed firm-level trade costs and log fixed product-level trade costs (column 2) on log distance between origins and destinations using Poisson pseudo maximum likelihood procedure discussed in the Online Appendix B. 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. Robust standard errors are reported in brackets. *, **, and *** represent the 5%, 1%, and 0.1% significance levels respectively.

Table 9: IME under granularity

	$corr(\alpha_j\varphi, \alpha_k\varphi)$	
	0	1
$\tilde{\theta} = 2.4$	0.005	0.001
$\tilde{\theta} = 1.25$	0.133	0.036
$\tilde{\theta} = 1$	0.333	0.103

Note: the table represents the estimated coefficients of the regression of the implied log average exports on log total exports using numerical simulations as discussed in the Online Appendix B. The sample is restricted to the year 2007 and to the origin-destination pairs with at least 100 exporters (867 observations). The first column reports the results from the model with zero correlation between the product of demand and productivity shocks across destinations. The second column reports the results for the model with perfect correlation. Four different values of $\tilde{\theta}$ were used.

Table 10: Number of firms and population

	log number of firms	
log population	0.945***	0.944***
Standard error	[0.0136]	[0.0139]
Observations	468	468
Year FE		Yes

Note: the table represents the regression of log number of firms, taken from Bento and Restuccia (2015) on log population, taken from World Development Indicators. Robust standard errors are reported in brackets. *, **, and *** represent the 5%, 1%, and 0.1% significance levels respectively.

Table 11: Trade costs and distance, Melitz-lognormal model

	log fixed costs	log variable costs
$\ln dist$	0.156***	0.299***
Standard error	[0.0155]	[0.0051]
Observations	7738	7738

Note: the table represents 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. Robust standard errors are reported in brackets. *, **, and *** represent the 5%, 1%, and 0.1% significance levels respectively.

Table 12: 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 represents the estimates of the full lognormal model. The estimation procedure is discussed in the 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 13: Implied IME in full Melitz-lognormal model

	IME	95% CI
Data	0.67	[0.61, 0.73]
Full Melitz-lognormal model	0.63	[0.59, 0.67]

Note: the table represents the coefficient from the regression of log average exports on log total exports with origin and destination fixed effects implied by the simulated full Melitz-lognormal 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.

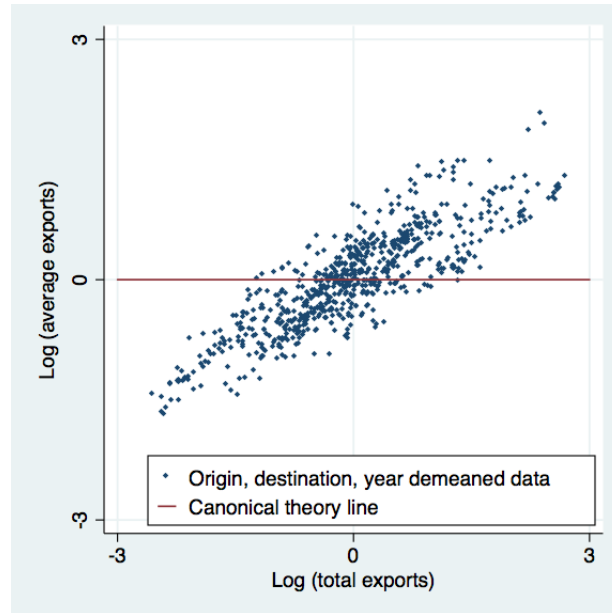
Table 14: Implied trade costs in full Melitz-lognormal model

	Estimate	95% CI
$corr(\tilde{F}_{ij}, \tilde{\tau}_{ij})$	-0.31	[-0.45, -0.1]
Distance elasticity		
Fixed costs	0.31	[0.18, 0.41]
Variable costs	0.34	[0.30, 0.37]

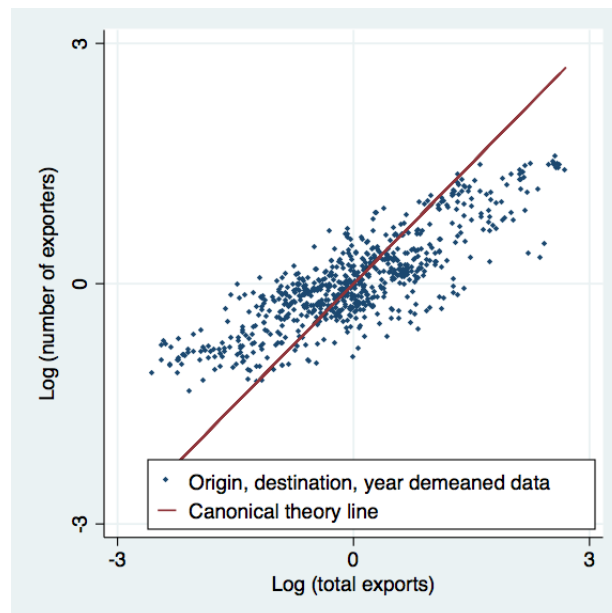
Note: the table represents the coefficient from the regression of log fixed and variable trade costs on distance, origin, and destination fixed effects implied by the simulated full lognormal model. The sample includes 37 origins and 4 main destinations (USA, Germany, France, and Japan) in 2007. The 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 (intensive margin) and total exports



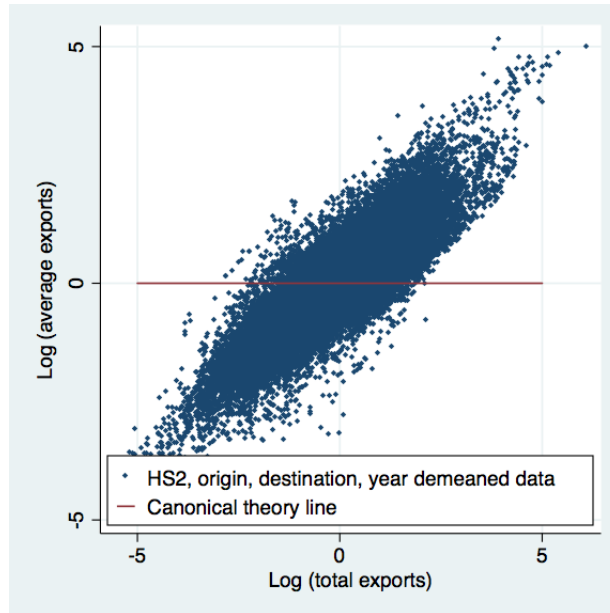
Panel B: Number of exporters (extensive margin) and total exports



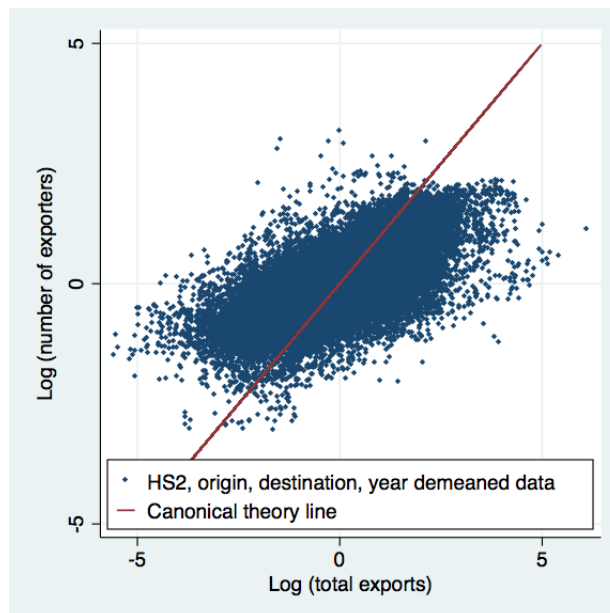
Note: the source are the statistics in the Exporter Dynamics Database for the extended sample. The x-axis represents log total exports at the origin country-destination country-year level demeaned by origin-year, and destination-year fixed effects. Only origin-destination pairs with more than 100 exporting firms considered. The dots represent the raw measures. The line is the slope predicted by the Melitz-Pareto model.

Figure 2: Intensive and Extensive margins of exporting, by industry

Panel a: Average size of exporters (intensive margin) and total exports

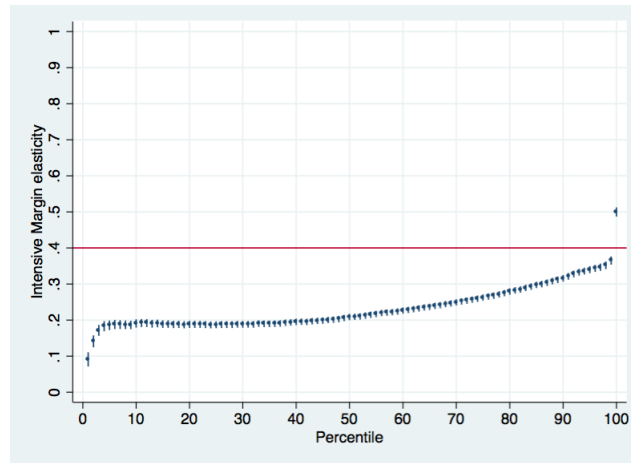


Panel b: Number of exporters (extensive margin) and total exports



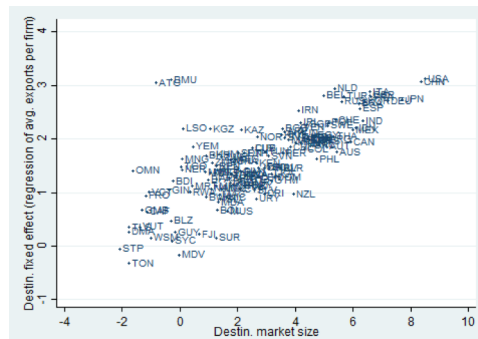
Note: the source are the statistics in the Exporter Dynamics Database for the extended sample. The x-axis represents log total exports at the origin country-HS 2-digit product-destination country-year level demeaned by origin-HS 2-digit-year and destination-HS 2-digit-year fixed effects. Only origin-HS 2-digit-destination triplets with more than 100 exporting firms are considered. The line is the slope predicted by the Melitz-Pareto model.

Figure 3: IME for each percentile, data



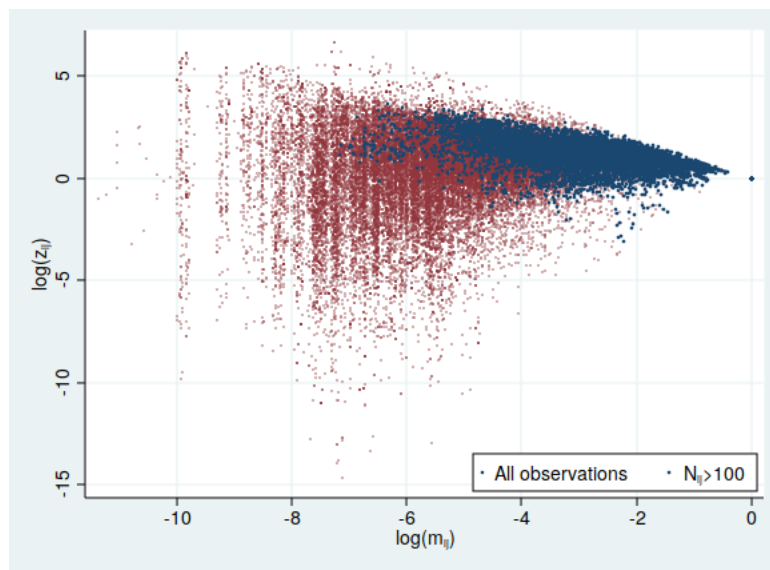
Note: the source is the exporter-level data used for the Exporter Dynamics Database for the core sample. The x-axis represents percentiles of the average exporter size distribution. Each dot represents the coefficient from the regression of log average exports per firm in an exporter size percentile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

Figure 4: Manufacturing absorption and averaged exports per firm (destination fixed effects)



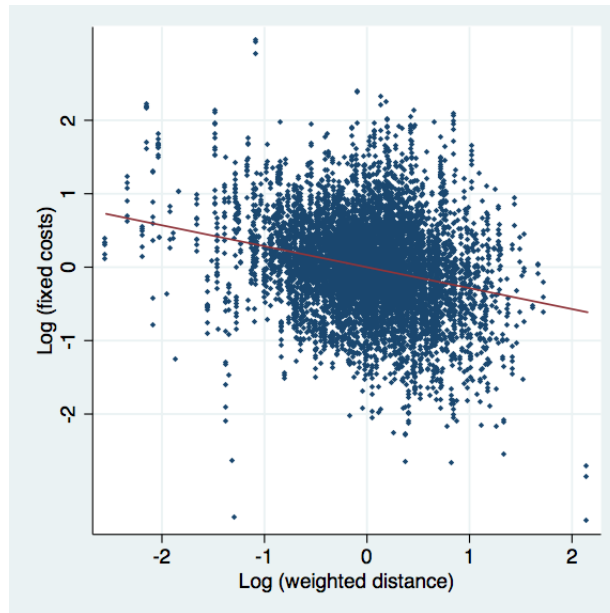
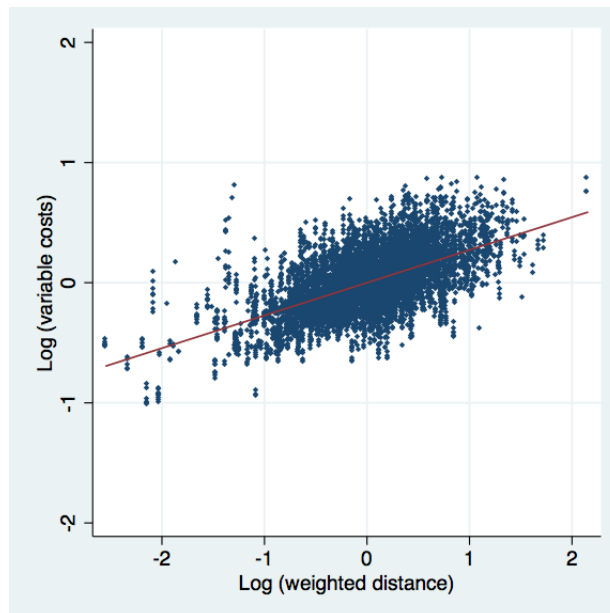
Note: the source is the exporter-level data used for the Exporter Dynamics Database for the core sample. The x-axis represents the log of manufacturing absorption in each destination country measured by manufacturing gross production plus manufacturing imports minus manufacturing exports (measured in billions of USD). The y-axis represents the estimated destination fixed effects obtained from a regression of log average exports per firm on origin, destination, and year fixed effects based on the core sample considering origin-destination pairs with more than 100 exporting firms. Manufacturing gross production is calculated as manufacturing value-added from the World Development Indicators divided by 0.418 (the factor used by EKK). Manufacturing imports and exports are obtained from COMTRADE/WITS.

Figure 5: Exports to largest destination and market entry



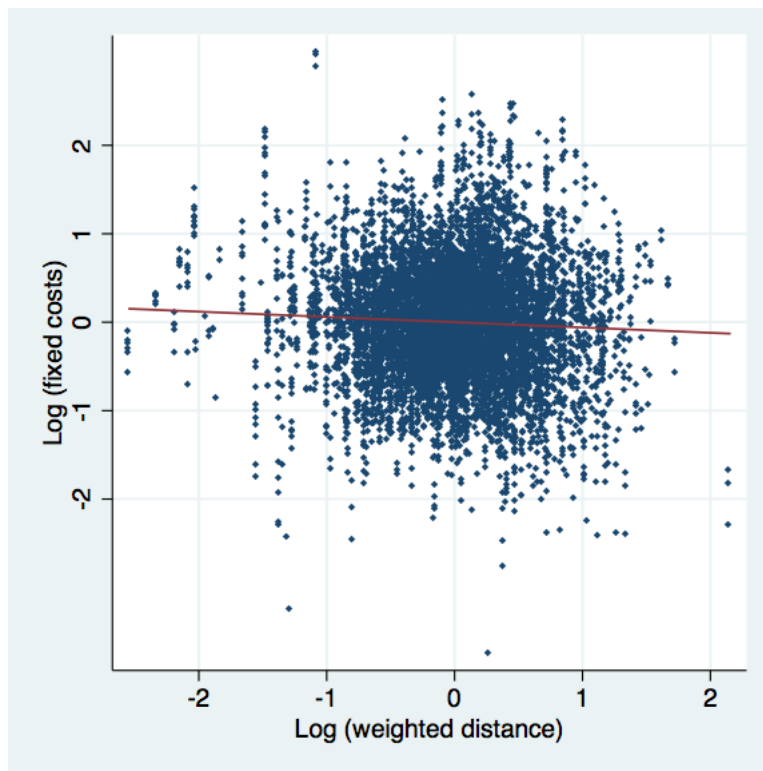
Note: the source is the exporter-level data used for the Exporter Dynamics Database for the core sample. The x-axis represents for each country i the log of the ratio of average exports per exporter to destination j to average exports per exporter to i 's most popular destination market, $\log m_{ij}$. The y-axis represents for each country i the log of the ratio of the number of exporters to destination j to the number of exporters to i 's most popular destination market, $\log z_{ij}$. A more formal definition of the variables can be found in the Online Appendix A. For the calculation of both average exports per exporter and number of exporters we focus only on firms from i that sell both in j and in the most popular destination.

Figure 6: Model-implied fixed and variable trade costs and distance

Panel a: fixed trade costs and distance*Panel b: variable trade costs and distance*

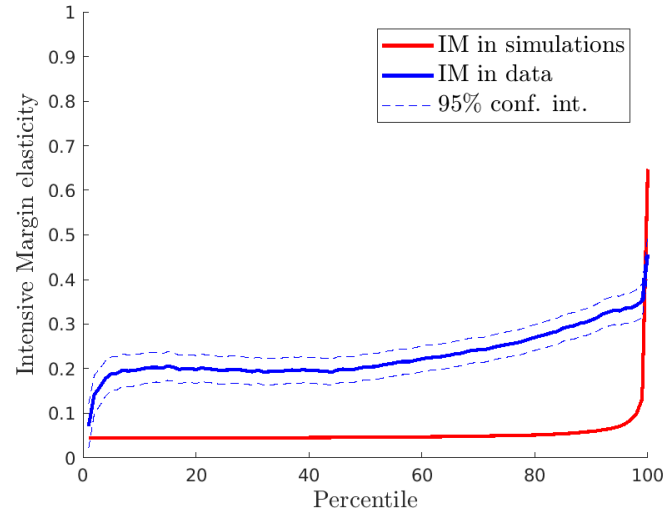
Note: the source is the exporter-level data used for the Exporter Dynamics Database. The x-axis represents log distance demeaned by origin and destination fixed effects taken from Mayer and Zignago (2011). The y-axis represents the fixed or variable trade costs implied by the basic Melitz-Pareto model demeaned by origin-year and destination-year fixed effects. To calculate the model-implied fixed and variable trade costs we use $\theta = 5$ from Head and Mayer (2014) and $\sigma = 5$ from Bas et al. (2015)

Figure 7: Fixed product-level trade costs and distance



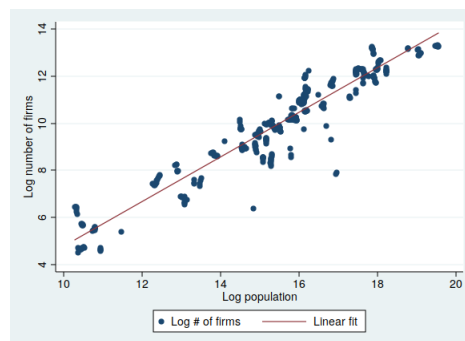
Note: the source is the exporter-level data used for the Exporter Dynamics Database. The x-axis represents log distance taken from Mayer and Zignago (2011). The y-axis represents fixed product-level trade costs by the the multi-product extension of the Melitz-Pareto model demeaned by origin-year and destination-year fixed effects.

Figure 8: IME for each percentile, Pareto and granularity



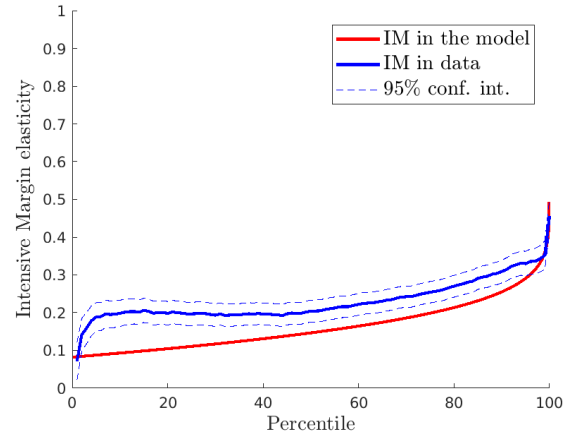
Note: the source is the exporter-level data used for the Exporter Dynamics Database. The darker solid line corresponds to IME for each percentile estimated using EDD and four main destinations: France, Germany, Japan and the U.S.. Dashed lines indicate 95% confidence intervals. The lighter solid line is IME for each percentile implied by the model with Pareto distribution of productivity and granularity, $\tilde{\theta} = 1$. The level of bilateral fixed trade costs was chosen to match overall IME in the data. The number of draws for each origin-destination pair is equal to the number of exporters from origin to destination in EDD as of 2007.

Figure 9: Number of firms and population



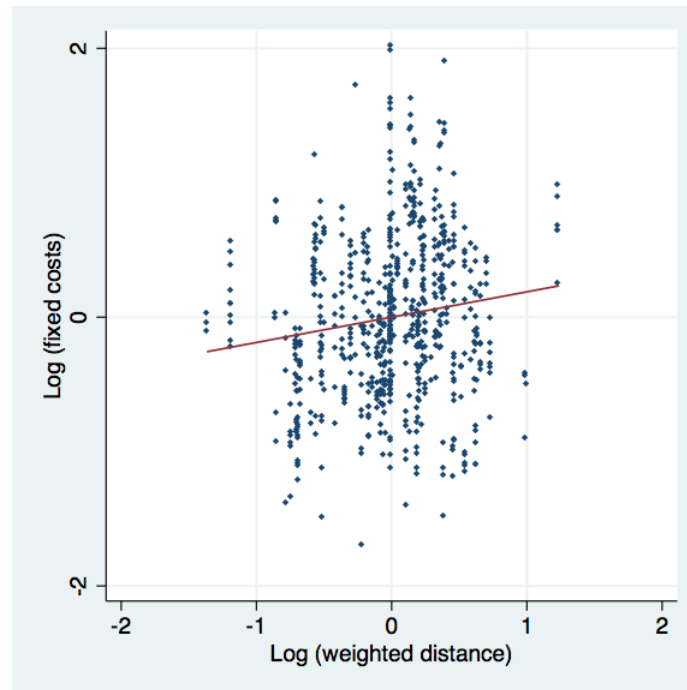
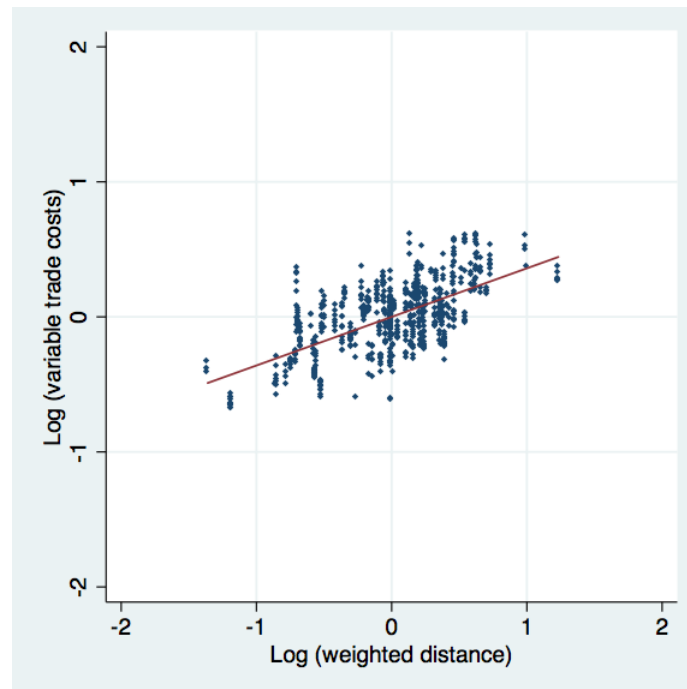
Note: the x-axis represents log of population taken from the World Development Indicators. The y-axis represents the number of firms as computed by Bento and Restuccia (2015). The sample includes all country-years for which EDD and data from Bento and Restuccia (2015) overlap.

Figure 10: IME for each percentile, lognormal



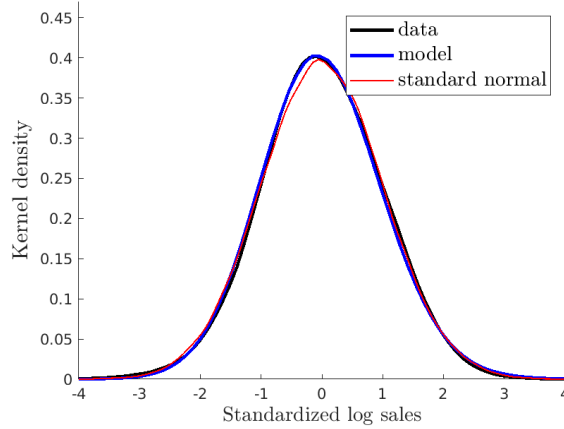
Note: the source is the exporter-level data used for the Exporter Dynamics Database. The darker solid line corresponds to IME for each percentile estimated using EDD and four main destinations: France, Germany, Japan and the U.S.. Dashed lines indicate 95% confidence intervals. The lighter solid line 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. (2015). 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 (2015).

Figure 11: Fixed and variable trade costs and distance, lognormal

Panel a: fixed trade costs and distance*Panel b: variable trade costs and distance*

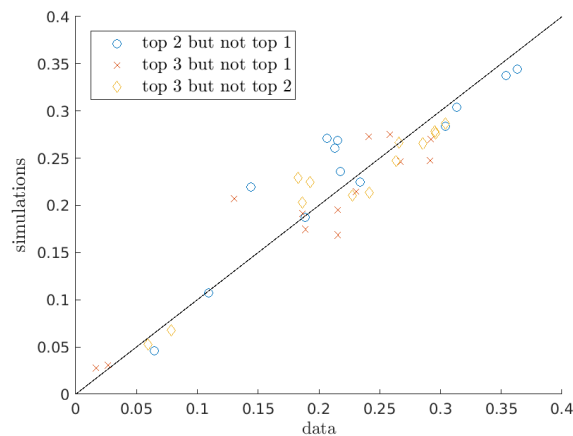
Note: source is the exporter-level data used for the Exporter Dynamics Database. The x-axis represents log distance taken from Mayer and Zignago (2011). Only four destination countries are considered: France, Germany, Japan, and the U.S.. To calculate the model-implied fixed and variable trade costs we use our estimate of $\sigma_\varphi = 4.02$ and $\sigma = 5$ from Bas et al. (2015), and the implied number of firms from Bento and Restuccia (2015).

Figure 12: Full Melitz-lognormal model, pdf of log sales



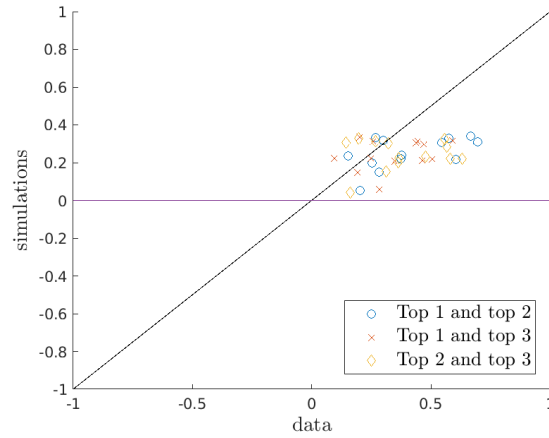
Note: the source is the exporter-level data used for the Exporter Dynamics Database and authors calculations. The black line corresponds to the standardized log sales (demeaned, divided by standard deviation) pooled across different origin-destination cells. The blue line corresponds to the standardized log sales pooled across different cells in the model. We used 1MM simulated draws for each origin-destination pair to calculate standardized log sales in the full Melitz-lognormal model. The red line corresponds to the standard normal distribution.

Figure 13: Share of firms selling to destination X but not to destination Y



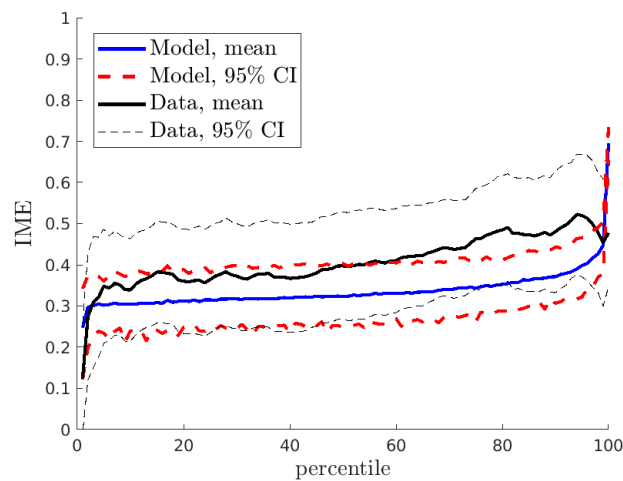
Note: the source is the exporter-level data used for the Exporter Dynamics Database and authors' calculations. Each point corresponds to the share of firms exporting only to less popular markets in the data (horizontal axis) and according to the estimated model (vertical axis) for each origin. The figure also includes a 45° line.

Figure 14: Correlation between log exports to top destinations



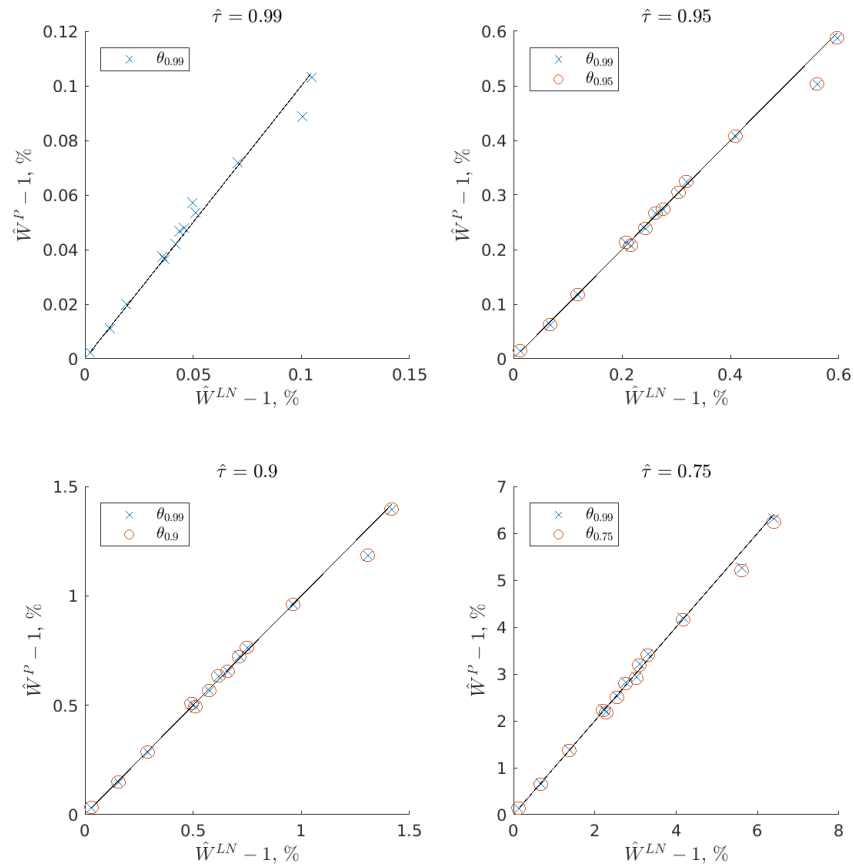
Note: the source is the exporter-level data used for the Exporter Dynamics Database and authors' calculations. Each point corresponds for each origin and any two destinations among the three most popular ones, the correlation in export value across all firms that sell in those two destinations in the data (horizontal axis) and according to the estimated model (vertical axis).

Figure 15: IME for each percentile, data and full Melitz-lognormal model



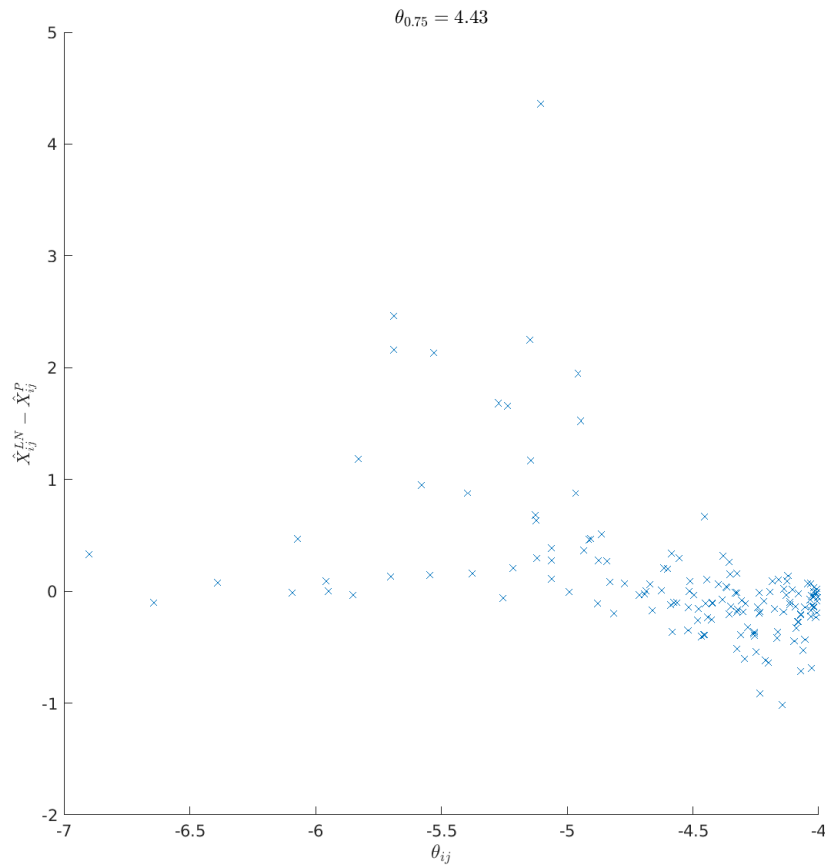
Note: the source is the exporter-level data used for the Exporter Dynamics Database and authors' calculations. The x-axis represent percentiles. The blue solid line represents coefficient from the regression of log average exports in each percentile on log total exports in the model; the dashed red lines 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 95% confidence interval.

Figure 16: Gains from trade liberalization



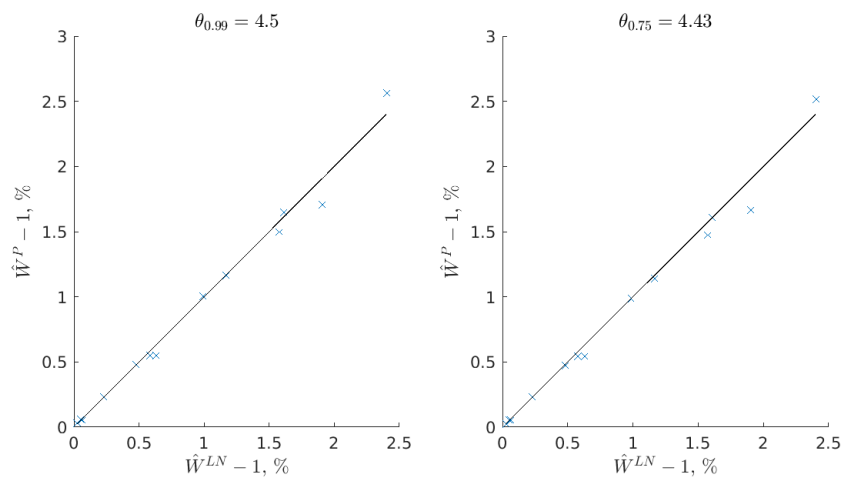
Note: the figure represents the change in welfare in response to a variable trade costs shock in the full Melitz-lognormal model and the Melitz-Pareto model. To calculate welfare gains in the full Melitz-lognormal model we used the parameter estimates from the Monte-Carlo Markov chain. Section 5.1. describes 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 for different value of trade elasticity used to calculate the gains. Each of the four panels reports the results for a different change in trade costs (1%, 5%, 10%, 25%). Blue crosses correspond to welfare change in the Melitz-Pareto model when we use the trade elasticity estimated from equation 49 using changes in trade flows implied by a full Melitz-lognormal model and a 1% reduction in trade costs. Orange circles correspond to welfare change in the Melitz-Pareto model when we use the trade elasticity estimated from equation 49 using changes in trade flows implied by a full Melitz-lognormal model and a reduction in trade costs that is indicated in the title of each panel.

Figure 17: Counterfactual changes in trade flows



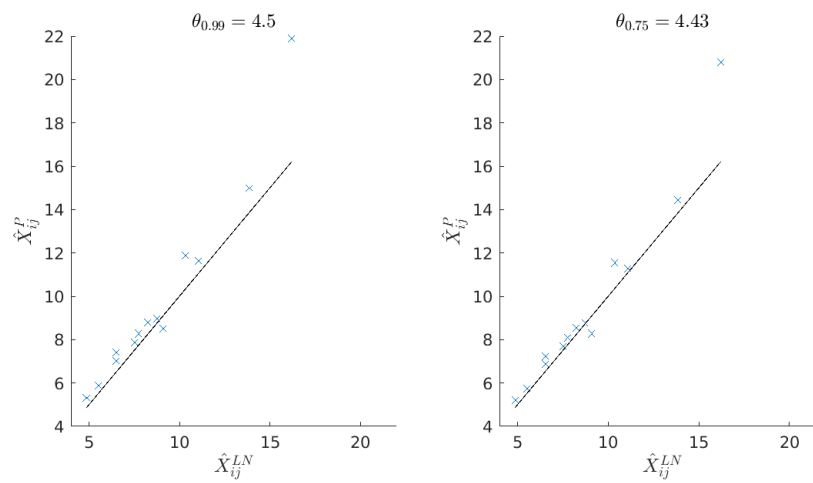
Note: the figure represents the difference between changes in trade flows in the full Melitz-lognormal model and the Melitz-Pareto model in response to a reduction of variable trade costs on the vertical axis, and the trade elasticity implied by the full Melitz-lognormal model on the horizontal axis. 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. We used equation 49 and changes in trade flows implied by a 25% reduction in trade costs in the full Melitz-lognormal model to estimate the trade elasticity that was used to calculate the changes in trade flows in the Melitz-Pareto model.

Figure 18: 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. To calculate gains in the Melitz-Pareto model the left panel (right panel) uses trade elasticity estimated from equation 49 using changes in trade flows implied by a 1% (25%) reduction in trade costs in the full Melitz-lognormal model.

Figure 19: 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. To calculate changes in trade flows in the Melitz-Pareto model the left panel (right panel) uses the trade elasticity estimated from equation 49 using changes in trade flows implied by a 1% (25%) reduction in trade costs in the full Melitz-lognormal model.