From Micro to Macro: Demand, Supply, and Heterogeneity in the Trade Elasticity*

Maria Bas[†] Thierry Mayer[‡] Mathias Thoenig [§]

September 16, 2015

Abstract

Models of heterogeneous firms with selection into export market participation generically exhibit aggregate trade elasticities that vary across country-pairs. Only when heterogeneity is assumed Pareto-distributed do all elasticities collapse into an unique elasticity, estimable with a gravity equation. This paper provides a theory-based method for quantifying country-pair specific elasticities when moving away from Pareto, i.e. when gravity does not hold. Combining two firm-level customs datasets for which we observe French and Chinese individual sales on the same destination market over the 2000-2006 period, we are able to estimate all the components of the dyadic elasticity: i) the demand-side parameter that governs the intensive margin and ii) the supply side parameters that drive the extensive margin. These components are then assembled under theoretical guidance to calculate bilateral aggregate elasticities over the whole set of destinations, and their decomposition into different margins. Our predictions fit well with econometric estimates, supporting our view that micro-data is a key element in the quantification of non-constant macro trade elasticities.

Keywords: trade elasticity, firm-level data, heterogeneity, gravity, Pareto, log-normal.

JEL Classification: F1

*This research has received funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) Grant Agreement No. 313522. We thank David Atkin, Dave Donaldson, Swati Dhingra, Ben Faber, Pablo Fagelbaum, Jean Imbs, Oleg Itskhoki, Jan de Loecker, Peter Morrow, Steve Redding, Andres Rodriguez Clare, Esteban Rossi-Hansberg, Katheryn Russ, Nico Voigtlander, and Yoto Yotov for useful comments on an early version, and participants at seminars in UC Berkeley, UCLA, Banque de France, WTO, CEPII, ISGEP in Stockholm, University of Nottingham, Banca d'Italia, Princeton University and CEMFI.

[†]University of Paris 1.

[‡]Sciences Po, Banque de France, CEPII and CEPR. Email: thierry.mayer@sciencespo.fr. Postal address: 28, rue des Saints-Peres, 75007 Paris, France.

[§]Faculty of Business and Economics, University of Lausanne and CEPR.

1 Introduction

The response of trade flows to a change in trade costs, the aggregate trade elasticity, is a central element in any evaluation of the welfare impacts of trade liberalization. Arkolakis et al. (2012) recently showed that this parameter, denoted ε for the rest of the paper, is actually one of the (only) two sufficient statistics needed to calculate Gains From Trade (GFT) under a surprisingly large set of alternative modeling assumptions—the ones most commonly used by recent research in the field. Measuring those elasticities has therefore been the topic of a long-standing literature in international economics. The most common usage (and the one recommended by Arkolakis et al., 2012) is to estimate this elasticity in a *macro-level* bilateral trade equation referred to as structural gravity in the literature following the initial impulse by Anderson and van Wincoop (2003). In order for this estimate of ε to be relevant for a particular experiment of trade liberalization, it is crucial for this bilateral trade equation to be correctly specified as a structural gravity model with, in particular, a *unique* elasticity to be estimated across dyads.

Our starting point is that the model of heterogeneous firms with selection into export market participation (Melitz, 2003) will in general exhibit a *dyad-specific* elasticity, i.e. an ε_{ni}, which applies to each country pair. Only when heterogeneity is assumed Pareto-distributed do all ε_{ni} collapse to a single ε. Under any other distributional assumption, obtaining an estimate of the aggregate trade elasticity from a macro-level bilateral trade equation becomes problematic, since there is a now a whole set of ε_{ni} to be estimated, and structural gravity does not hold anymore. We argue that in this case quantifying trade elasticities at the aggregate level makes it necessary to use micro-level information. To this purpose we exploit a rich panel that combines sales of French and Chinese exporters over 2000-2006 on many destination-product combinations for which we also observe the applied tariff. We propose a theory-based method using this firm-level export data for estimating all the components of the dyad-specific trade elasticity: i) the demand-side parameter that governs the intensive margin and ii) the supply side parameters that drive the extensive margin. These components are then assembled under theoretical guidance to calculate the dyadic aggregate elasticities over the whole set of destination-product.

Taking into account cross-dyadic heterogeneity in trade elasticities is crucial for quantifying the expected impact of various trade policy experiments.[1] Consider the example of the current negotiations over a transatlantic trade agreement between the USA and the EU (TTIP). Under the simplifying assumption of a unique elasticity, whether the trade liberalization takes place with a proximate vs distant, large vs small economy, etc. is irrelevant in terms of trade-promoting effect or welfare gains calculations. By contrast, our results suggest that the relevant ε_{ni} should be smaller (in absolute value) than if the United States were considering a comparable agreement with countries where the expected volume of trade is smaller. Regarding welfare, Melitz and Redding (2015) and Head et al. (2014) have shown theoretically that the GFT can be quite substantially mis-estimated if one assumes a constant trade elasticity when the "true" elasticity is variable (the margin of error can exceed 100 percent in both papers). The expected changes in trade patterns and welfare effects of agreements such as TTIP will therefore be different compared to the unique elasticity case. One of the main objectives of our paper is to quantify how wrong can one be when making predictions based on a constant trade elasticity assumption.

Our approach maintains the traditional CES (σ) demand system combined with monopolistic competition. It features several steps that are structured around the following decomposition of aggregate trade elasticity into the sum of the intensive margin and the (weighted) extensive margin:

$$\varepsilon_{ni} = \underbrace{1-\sigma}_{\text{intensive margin}} + \underbrace{\frac{1}{\bar{x}_{ni}/x_{ni}^{\texttt{MIN}}}}_{\text{min-to-mean}} \times \underbrace{\frac{d\ln N_{ni}}{d\ln \tau_{ni}}}_{\text{extensive margin}}, \tag{1}$$

[1] Imbs and Méjean (2015) and Ossa (2015) recently argued that another source of heterogeneity, the cross-sectoral one, raises important aggregation issues that matter for aggregate outcomes of trade liberalization. We abstract from this issue (which would reinforce the importance of heterogeneity for aggregate outcomes) in our paper, and mostly omit cross-sectional variation in ε, apart from section 5.5 where we use industry-level estimates to show that both demand and supply side determinants enter aggregate elasticities.

The weight is the *mean-to-min ratio*, our observable measuring the dyadic dispersion of firm-level performance, that is defined as the ratio of average to minimum sales across markets. Intuitively, the weight of the extensive margin should be decreasing in easy markets where the increasing presence of weaker firms augments productivity dispersion. When assuming Pareto with shape parameter θ, the last part of the elasticity reduces to $\sigma - 1 - \theta$, and the overall elasticity becomes constant and reflects only the supply side homogeneity in the distribution of productivity: $\varepsilon^P_{ni} = \varepsilon^P = -\theta$ (Chaney, 2008).

Our first step aims to estimate the demand side parameter σ using firm-level exports. Since protection is imposed on all firms from a given origin, higher demand and lower protection are not separately identifiable when using only one country of exports. With CES, firms are all confronted to the same aggregate demand conditions. Thus, considering a second country of origin enables to isolate the effects of trade policy, if the latter is discriminatory. We therefore combine shipments by French and Chinese exporters to destinations that confront those firms with different levels of tariffs. Our setup yields a *firm-level* gravity equation specified as a ratio-type estimation so as to eliminate unobserved characteristics of both the exporting firm and the importer country, while keeping tariffs in the regression.[2] We explore different sources of variance in the data with comparable estimates of the intensive margin trade elasticity that imply an average value of σ around 5.

Our second and main step applies equation (1) and assembles the estimates of the intensive margin ($\hat{\sigma}$) with the central supply side parameter —reflecting dispersion in the distribution of productivity— estimated on the same datasets, to obtain predicted aggregate elasticities of total export, number of exporters and average exports to each destination. Those dyadic predictions (one elasticity for each exporter-importer combination) require knowledge of the bilateral export productivity cutoff under which firms find exports to be unprofitable. We also make use of the mean-to-min ratio to reveal those cutoffs. A key element of our procedure is the calibration of the productivity distribution. As an alternative to Pareto we consider the log-normal distribution that fits the micro-data on firm-level sales very well. We show that under log-normal the ε_{ni} are larger (in absolute value) for pairs with low volumes of trade. Hence the trade-promoting impact of liberalization is expected to be larger for this kind of trade partners.

A side result of our paper is to discriminate between Pareto and log-normal as potential distributions for the underlying firm-level heterogeneity, suggesting that log-normal does a better job at matching the non-unique response of exports to changes in trade costs. Two pieces of evidence in that direction are provided.[3] The first provides direct evidence that aggregate elasticities are non-constant across dyads. The second is a positive and statistically significant correlation across industries between firm-level and aggregate elasticities–at odds with the prediction of a null correlation under Pareto. We also find that the heterogeneity in trade elasticities is quantitatively important: Although the cross-dyadic average of bilateral elasticities is quite well approximated by a standard gravity model constraining the estimated parameter to be constant, deviations from this average level can be large. For Chinese exports, assuming a unique elasticity would yield to underestimate the trade impact of a tariff liberalization by about 25% for countries with initially very small trade flows (Somalia, Chad or Azerbaijan for instance). By contrast, the error would be to overestimate by around 20% the exports created when the United States or Japan reduce their trade costs.

In terms of literature, our paper relates to several recent papers studying patterns and consequences of heterogeneity in trade elasticities. Berman et al. (2012) and Gopinath and Neiman (2014) find that

[2] Other work in the literature also relies on the ratio-type estimation. Romalis (2007) uses a similar method to estimate the effect of tariffs on trade flows at the product-country level. He estimates the effects of applied tariff changes within NAFTA countries (Canada and Mexico) on US imports at the product level. Hallak (2006) estimates a fixed effects gravity model and then uses a ratio of ratios method in a quantification exercise. Caliendo and Parro (2015) also use ratios of ratios and rely on asymmetries in tariffs to identify industry-level elasticities.

[3] Head et al. (2014) provide evidence and references for several micro-level datasets that individual sales are much better approximated by a log-normal distribution when the entire distribution is considered (without left-tail truncation). Freund and Pierola (2015) is a recent example showing very large deviations from the Pareto distribution if the data is not vastly truncated for all of the 32 countries used. Our findings complement those papers by providing industry- and aggregate-level evidence on trade elasticities.

in order to predict correctly the aggregate patterns of trade adjustments to price shocks, one has to take into account firm-level heterogeneity with use of micro data. In their case, heterogeneity matters because firms have different individual responses in export and/or import behavior. In particular, both papers find that the firm-level elasticity depends negatively on the size of the firm. Our paper also finds that measuring aggregate trade responses requires usage of firm-level data. It is however for a different reason: In our case, heterogeneity in aggregate trade elasticities simply originates in a departure from the common assumption that productive efficiency is Pareto-distributed. While we do recognize that trade elasticities might differ across firms, our paper shows that this is not required to ensure that heterogeneity matters for the aggregate economy and investigates a different, complementary, channel.

In the empirical literature estimating trade elasticities, different approaches and proxies for trade costs have been used, with an almost exclusive focus on aggregate country or industry-level data. The gravity approach to estimating those elasticities mostly uses tariff data to estimate bilateral responses to variation in applied tariff levels. Most of the time, identification is in the cross-section of country pairs, with origin and destination determinants being controlled through fixed effects (Baier and Bergstrand (2001), Head and Ries (2001), Caliendo and Parro (2015), Hummels (1999), Romalis (2007) are examples). A related approach is to use the fact that most foundations of gravity have the same coefficient on trade costs and domestic cost shifters to estimate that elasticity from the effect on bilateral trade of exporter-specific changes in productivity, export prices or exchange rates (Costinot et al. (2012) is a recent example).[4] Baier and Bergstrand (2001) find a demand side elasticity ranging from -4 to -2 using aggregate bilateral trade flows from 1958 to 1988. Using product-level information on trade flows and tariffs, this elasticity is estimated by Head and Ries (2001), Romalis (2007) and Caliendo and Parro (2015) with benchmark average elasticities of -6.88, -8.5 and -4.45 respectively. Costinot et al. (2012) also use industry-level data for OECD countries, and obtains a preferred elasticity of -6.53 using productivity based on producer prices of the exporter as the identifying variable. Our paper also has consequences for how to interpret those numbers in terms of underlying structural parameters. With a homogeneous firms model of the Krugman (1980) type in mind, the estimated elasticity turns out to reveal a demand-side parameter only (this is also the case with Armington differentiation and perfect competition as in Anderson and van Wincoop (2003)). When instead considering heterogeneous firms à la Melitz (2003), the literature has proposed that the macro-level trade elasticity is driven solely by a supply-side parameter describing the dispersion of the underlying heterogeneity distribution of firms. This result has been shown with several demand systems (CES by Chaney (2008), linear by Melitz and Ottaviano (2008), translog by Arkolakis et al. (2010) for instance), but again relies critically on the assumption of a Pareto distribution. The trade elasticity then provides an estimate of the dispersion parameter of the Pareto.[5] We show here that both existing interpretations of the estimated elasticities are too extreme: When the Pareto assumption is relaxed, the aggregate trade elasticity is a mix of demand and supply parameters.

There is a small set of papers that estimate the intensive margin elasticity at the exporter level. Berman et al. (2012) presents estimates of the trade elasticity with respect to real exchange rate variations across countries and over time using firm-level data from France. Fitzgerald and Haller (2014) use firm-level data from Ireland, real exchange rate and weighted average firm-level applied tariffs as price shifters to estimate the trade elasticity to trade costs. The results for the impact of real exchange rate on firms' export sales are of a similar magnitude, around 0.8 to 1. Applied tariffs vary at the product-destination-year level. Fitzgerald and Haller (2014) create a firm-level destination tariff as the weighted average over

[4]Other methodologies (also used for aggregate elasticities) use identification via heteroskedasticity in bilateral flows, and have been developed by Feenstra (1994) and applied widely by Broda and Weinstein (2006) and **?**. Yet another alternative is to proxy trade costs using retail price gaps and their impact on trade volumes, as proposed by Eaton and Kortum (2002) and extended by Simonovska and Waugh (2011).

[5]This result of a constant trade elasticity reflecting the Pareto shape holds when maintaining the CES demand system but making other improvements to the model such as heterogeneous marketing and/or fixed export costs (Arkolakis, 2010; Eaton et al., 2011). In the Ricardian setup of Eaton and Kortum (2002) , the trade elasticity is also a (constant) supply side parameter reflecting heterogeneity, but this heterogeneity takes place at the national level, and reflects the scope for comparative advantage.

all hs6 products exported by a firm to a destination in a year using export sales as weights. Relying on this construction, they find a tariff elasticity of around -2.5 at the micro level. This is also the preferred estimates of Berthou and Fontagné (2015), who use the response of the largest French exporters in the United States to the levels of applied tariffs. We depart from those papers by using an alternative methodology to identify the trade elasticity with respect to applied tariffs; i.e. the differential treatment of exporters from two distinct countries (France and China) in a set of product-destination markets.

Our paper also contributes to the literature studying the importance of the distribution assumption of heterogeneity for trade patterns, trade elasticities and welfare. Head et al. (2014), Yang (2014), Melitz and Redding (2015) and Feenstra (2013) have recently argued that the simple gains from trade formula proposed by Arkolakis et al. (2012) relies crucially on the Pareto assumption, which mutes important channels of gains in the heterogenous firms case. Barba Navaretti et al. (2015) present gravity-based evidence that the exporting country fixed effects depends on characteristics of firms' distribution that go beyond the simple mean productivity, a feature incompatible with the usually specified Pareto heterogeneity. Fernandes et al. (2015) use customs data for numerous developing countries to show that a decomposition of total bilateral exports into intensive and an extensive margins exhibits an important role for the latter, with patterns consistent with log-normally distributed heterogeneity and incompatible with untruncated Pareto. The alternatives to Pareto considered to date in welfare gains quantification exercises are i) the truncated Pareto by Helpman et al. (2008), Melitz and Redding (2015) and Feenstra (2013), and ii) the log-normal by Head et al. (2014) and Yang (2014). A key simplifying feature of Pareto is to yield a constant trade elasticity, which is not the case for alternative distributions. Helpman et al. (2008) and Novy (2013) have produced gravity-based evidence showing substantial variation in the trade cost elasticity across country pairs. Our contribution to that literature is to use the estimated demand and supply-side parameters to construct predicted bilateral elasticities for aggregate flows under the log-normal assumption, and compare their first moments to gravity-based estimates. It is possible to generate bilateral trade elasticities changing another feature of the standard model. The most obvious is to depart from the simple CES demand system. Novy (2013) builds on Feenstra (2003), using the translog demand system with homogeneous firms to obtain variable trade elasticities. Atkeson and Burstein (2008) is another example maintaining CES demand, and generating heterogeneity in elasticities trough monopolistic competition. We choose here to keep the change with respect to the benchmark Melitz/Chaney framework to a minimal extent, keeping CES and monopolistic competition, while changing only the distributional assumption.

The next section of the paper describes our model and empirical strategy. The third section presents the different firm-level data and the product-country level tariff data used in the empirical analysis. The fourth section reports the estimates of the intensive margin elasticity. Section 5 computes predicted macro-level trade elasticities and compares them with estimates from the Chinese and French aggregate export data. It also provides two additional pieces of evidence in favor of non-constant trade elasticities. The final section concludes.

2 Empirical strategy for estimating the demand side parameter

2.1 A firm-level export equation

Consider a set of potential exporting firms, all located in the same origin country i and producing product p (omitting those indexes for the start of exposition). We use the Melitz (2003) theoretical framework of heterogeneous firms facing constant price elasticity demand (CES utility combined with iceberg costs) and contemplating exports to several destinations. In this setup, firm-level exports to country n depend upon the firm-specific unit input requirement (α), wages (w), and "real" expenditure in n, $X_n P_n^{\sigma-1}$, with P_n the ideal CES price index relevant for sales in n. There are trade costs associated with reaching market n, consisting of an observable iceberg-type part (τ_n), and a shock that affects firms differently on each market, $b_n(\alpha)$:[6]

$$x_n(\alpha) = \left(\frac{\sigma}{\sigma - 1}\right)^{1-\sigma} [\alpha w \tau_n b_n(\alpha)]^{1-\sigma} \frac{X_n}{P_n^{1-\sigma}} \tag{2}$$

Taking logs of the demand equation (2), and noting with $\epsilon_n(\alpha) \equiv b_n^{1-\sigma}$ our unobservable firm-destination error term, and with $A_n \equiv X_n P_n^{\sigma-1}$ the attractiveness of country n (expenditure discounted by the degree of competition on this market), a "firm-level gravity" equation can be derived:

$$\ln x_n(\alpha) = (1-\sigma)\ln\left(\frac{\sigma}{\sigma-1}\right) + (1-\sigma)\ln(\alpha w) + (1-\sigma)\ln \tau_n + \ln A_n + \ln \epsilon_n(\alpha) \tag{3}$$

Our objective is to estimate the trade elasticity, $1-\sigma$ identified on cross-country differences in applied tariffs (that are part of τ_n). This involves controlling for a number of other determinants ("nuisance" terms) in equation (3). First, it is problematic to proxy for A_n, since it includes the ideal CES price index P_n, which is a complex non-linear construction that itself requires knowledge of σ. A well-known solution used in the gravity literature is to capture (A_n) with destination country fixed effects (which also solves any issue arising from omitted unobservable n-specific determinants). This is however not applicable here since A_n and τ_n vary across the same dimension. To separate those two determinants, we use a second set of exporters, based in a country that faces different levels of applied tariffs, such that we recover a bilateral dimension on τ. The firm-level sales become

$$\ln x_{ni}(\alpha) = (1-\sigma)\ln\left(\frac{\sigma}{\sigma-1}\right) + (1-\sigma)\ln(\alpha w_i) + (1-\sigma)\ln \tau_{ni} + \ln A_n + \ln \epsilon_{ni}(\alpha), \tag{4}$$

where each firm can now be based in one of the two origin countries for which we have customs data, France and China, $i = [\text{FR}, \text{CN}]$. A second issue is that we need to control for firm-level marginal costs (αw_i). Again measures of firm-level productivity and wages are hard to obtain for two different source countries on an exhaustive basis. In addition, there might be a myriad of other firm-level determinants of export performance, such as quality of products exported, managerial capabilities... which will remain unobservable. Capturing those determinants through firm-level and country fixed effects is an option which proves computationally intensive in a setup with endogenous selection into export markets. We adopt an alternative approach, a ratio-type estimation inspired by Hallak (2006), Romalis (2007), Head et al. (2010), and Caliendo and Parro (2015) that removes observable and unobservable determinants for both firm-level and destination factors.[7] This method uses four individual export flows to calculate ratios of ratios: an approach referred to as *tetrads* from now on. We now turn to a presentation of this method.

2.2 A ratio-type estimating equation

Estimating micro-level tetrads implies dividing product-level exports of a firm located in France to country n by the exports of the same product by that *same* firm to a reference country, denoted k. Then, calculate a similar ratio for a Chinese exporter (same product and countries). Finally the ratio of those two ratios uses the multiplicative nature of the CES demand system to get rid of all the "nuisance" terms mentioned above. Because there is quite a large number of exporters, taking all possible firm-destination-product combinations is not feasible. We therefore concentrate our identification on the largest exporters for each product.[8] We rank firms based on export value for each hs6 product and reference importer country k. For a given product, taking the ratio of exports of a French firm with rank j exporting to country n, over

[6] An example of such unobservable term would be the presence of workers from country n in firm α, that would increase the internal knowledge on how to reach consumers in n, and therefore reduce trade costs for that specific company in that particular market (b being a mnemonic for barrier to trade). Note that this type of random shock is isomorphic to assuming a firm-destination demand shock in this CES-monopolistic competition model.

[7] The appendix provides a robustness check of our baseline table of results based on a fixed effect approach.

[8] The appendix presents an alternative strategy that keeps all exporters and explicitly takes into account selection issues.

the flow to the reference importer country k, removes the need to proxy for firm-level characteristics in equation (4):

$$\frac{x_n(\alpha_{j,\mathrm{FR}})}{x_k(\alpha_{j,\mathrm{FR}})} = \left(\frac{\tau_{n\mathrm{FR}}}{\tau_{k\mathrm{FR}}}\right)^{1-\sigma} \times \frac{A_n}{A_k} \times \frac{\epsilon_n(\alpha_{j,\mathrm{FR}})}{\epsilon_k(\alpha_{j,\mathrm{FR}})} \tag{5}$$

To eliminate the aggregate attributes of importing countries n and k, we need the two sources of firm-level exports to have information on sales by destination country. This allows to take the ratio of equation (5) over the same ratio for a firm with rank j located in China:

$$\frac{x_n(\alpha_{j,\mathrm{FR}})/x_k(\alpha_{j,\mathrm{FR}})}{x_n(\alpha_{j,\mathrm{CN}})/x_k(\alpha_{j,\mathrm{CN}})} = \left(\frac{\tau_{n\mathrm{FR}}/\tau_{k\mathrm{FR}}}{\tau_{n\mathrm{CN}}/\tau_{k\mathrm{CN}}}\right)^{1-\sigma} \times \frac{\epsilon_n(\alpha_{j,\mathrm{FR}})/\epsilon_k(\alpha_{j,\mathrm{FR}})}{\epsilon_n(\alpha_{j,\mathrm{CN}})/\epsilon_k(\alpha_{j,\mathrm{CN}})}. \tag{6}$$

Denoting tetradic terms with a $\widetilde{}$ symbol, one can re-write equation (6) as

$$\widetilde{x}_{\{j,n,k\}} = \widetilde{\tau}^{1-\sigma}_{\{n,k\}} \times \widetilde{\epsilon}_{\{j,n,k\}}, \tag{7}$$

which will be our main foundation for estimation.

Restoring the product subscript (p), and using i = FR or CN as the origin country index, we specify bilateral trade costs as a function of applied tariffs, with ad valorem rate t^p_{ni} and of a collection of other barriers, denoted with D_{ni}. Those include the classical gravity covariates such as distance, common language, colonial link and common border. Taking the example of a continuous variable such as distance for D_{ni}, $\tau^p_{ni} = (1 + t^p_{ni})D^{\delta}_{ni}$, which, once introduced in the logged version of (7) leads to our estimable equation

$$\ln \widetilde{x}^p_{\{j,n,k\}} = (1-\sigma)\ln\left(\widetilde{1 + t^p_{\{n,k\}}}\right) + (1-\sigma)\delta \ln \widetilde{D_{\{n,k\}}} + \ln \widetilde{\epsilon}^p_{\{j,n,k\}}. \tag{8}$$

The dependent variable corresponds to the ratio of ratios of exports for a certain rank j. In order to obtain a valid observation for each product, we iterate from $j = 1$ to 10, that is firms ranking from the top to the 10th exporter for a given product. Our precise procedure is the following: Firms are ranked according to their export value for each product and reference importer country k. We then take the tetrad of exports of the top French firm over the top Chinese firm exporting the same product to the same destination. This gives us a first sample, labeled "Top 1". We then fill in the missing values with lower ranked export tetrads: We start with the top Chinese exporter ($j = 1$) flow which divides French exporter's flow, iterating the French firm over $j = 2$ to 10, until a non-missing tetrad is generated. If the tetrad is still missing, the procedure then goes to the Chinese exporter ranked $j = 2$ and restarts iterating until Chinese exporter ranked $j = 10$ is reached. The resulting is an extended sample (around three-fold increase), with one observation for each product×destination×reference, labeled "Top 1 to 10". Comparing results from the restricted and the extended sample also helps to alleviate potential concerns regarding the influence of relying only on the top exporter. Bias might arise for instance if the largest firm has a different pass-through of tariffs to its price. Adding lower rank firms and checking similarity of results is a robustness check that selecting the top firm is not critically influential.

3 Data

We combine French and Chinese firm-level datasets from the corresponding customs administrations which report export value by firm at the hs6 level for all destinations in 2000. The firm-level customs datasets are matched with data on tariffs effectively applied to each exporting country (China and France) at the same level of product disaggregation for each destination. Focusing on 2000 allows us to exploit variation in tariffs applied to each exporter country (France/China) at the product level by the importer countries since it precedes the entry of China into WTO at the end of 2001. We also exploit the variation over time of trade and tariffs from 2000 to 2006 in a set of robustness checks.

Trade: The French trade data comes from the French Customs, which provide annual export data at the product level for French firms.[9] The customs data are available at the 8-digit product level Combined Nomenclature (CN) and specify the country of destination of exports. The free on board (f.o.b) value of exports is reported in euros and we converted those to US dollars using the real exchange rate from Penn World Tables for 2000. The Chinese transaction data comes from the Chinese Customs Trade Statistics (CCTS) database which is compiled by the General Administration of Customs of China. This database includes monthly firm-level exports at the 8-digit HS product-level (also reported f.o.b) in US dollars. The data is collapsed to yearly frequency. The database also records the country of destination of exports. In both cases, export values are aggregated at the firm-product(hs6)-destination level in order to match with applied tariffs information that are available at the hs6-destination level.[10]

Tariffs: Tariffs come from the WITS (World Bank) database.[11] We rely on the ad valorem rate effectively applied at the hs6 level by each importer country to France and China. In our cross-section analysis performed for the year 2000 before the entry of China into the World Trade Organization (WTO), we exploit different sources of variation within hs6 products across importing countries on the tariff applied to France and China. The first variation naturally comes from the European Union (EU) importing countries that apply zero tariffs to trade with EU partners (like France) and a common external tariff to extra-EU countries (like China). The second source of variation in the year 2000 is that several non-EU countries applied the Most Favored Nation tariff (MFN) to France, while the effective tariff applied to Chinese products was different (since China was not yet a member of WTO). One might be worried by the presence of unobserved bilateral trade costs correlated with our measure of applied tariffs. Even though it is not clear that the correlation with those omitted trade costs should be systematically positive, we use, as a robustness check, a more inclusive measure of applied trade costs, the Ad Valorem Equivalent (AVE) tariffs from WITS and MAcMAp databases.

Gravity controls: In all estimations, we include additional trade barriers variables that determine bilateral trade costs, such as distance, common (official) language, colony and common border (contiguity). The data come from the CEPII distance database.[12] We use the population-weighted great circle distance between the set of largest cities in the two countries.

The use of a reference country, k in equation (5), is crucial for a consistent identification of the trade elasticity. We choose reference importer countries with two criteria in mind. First, these countries should be those that are the main trade partners of France and China in the year 2000, since we want to minimize the number of zero trade flows in the denominator of the tetrad. The second criteria relies on the variation in the tariffs effectively applied by the importing country to France and China. Hence, among the main trade partners, we retain those countries for which the average difference between the effectively applied ad valorem tariffs to France and China is greater. These two criteria lead us to select the following set of 8 reference countries: Australia, Canada, Germany, Italy, Japan, New Zealand, Poland and the UK. Tables A.3 and A.4 in the appendix present, for each destination country, the count of products for which the difference in tariffs applied to France and China is positive, negative or zero, together with the average tariff gap. The difference in the effectively applied tariffs to France and China is illustrated at the industry level for two major reference importer countries (Germany and Japan) in figure 1 (with precise numbers provided in the appendix, Table A.1). As can be noticed, there is a significant variation across 2-digit industries in the average percentage point difference in applied tariffs to both exporting countries

[9] This database is quite exhaustive. Although reporting of firms by trade values below 250,000 euros (within the EU) or 1,000 euros (rest of the world) is not mandatory, there are in practice many observations below these thresholds.

[10] The hs6 classification changes over time. During our period of analysis it has only changed once in 2002. To take into account this change in the classification of products, we have converted the HS-2002 into HS-1996 classification using WITS conversion tables.

[11] Information on tariffs is available at `http://wits.worldbank.org/wits/`

[12] This dataset is available at `http://www.cepii.fr/anglaisgraph/bdd/distances.htm`

in the year 2000. This variation is even more pronounced at the hs6 product level. Our empirical strategy will exploit this variation within hs6 products and across destination countries.

Figure 1: Average percentage point difference between the applied tariff to France and China across industries by Germany and Japan (2000)

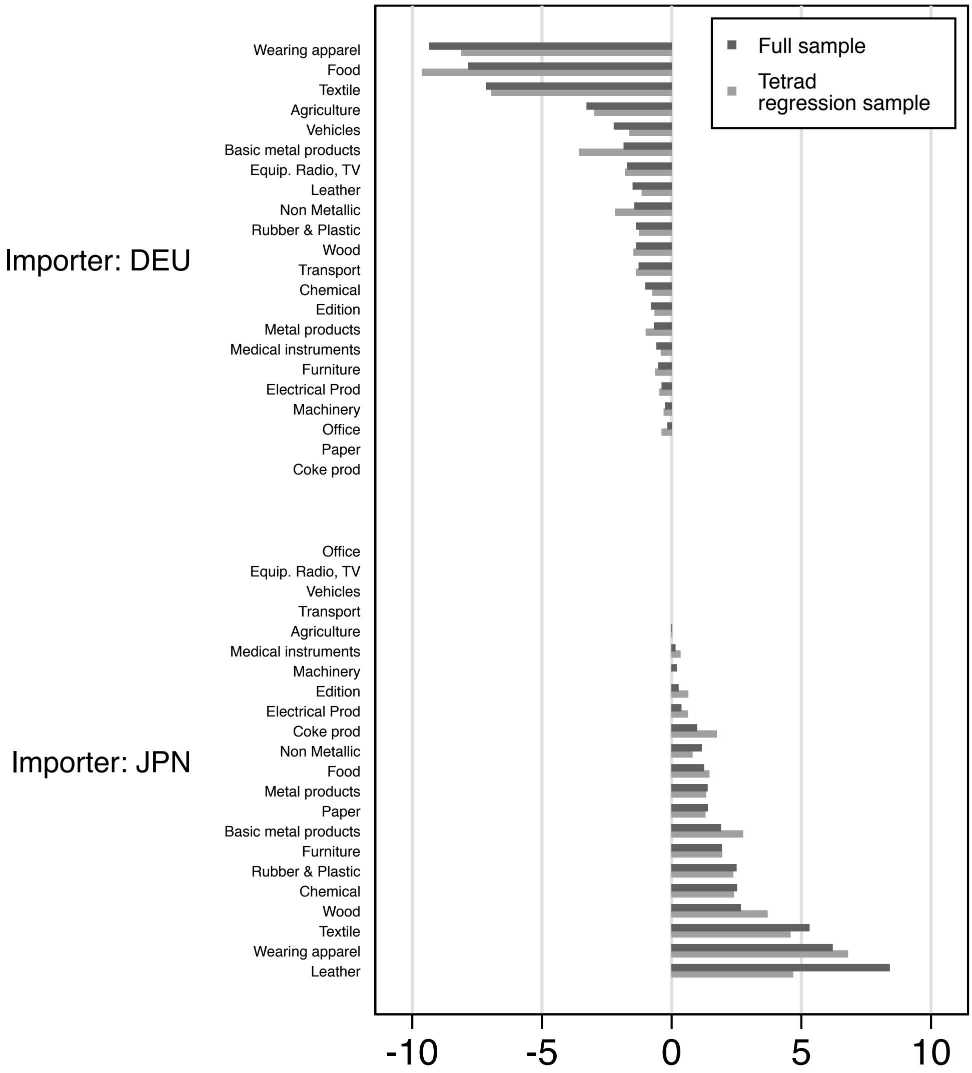

Source: Authors' calculation based on Tariff data from WITS (World Bank).

The final size of the "Top 1 to 10" estimating sample is 99,645 product-destination-reference country observations in the year 2000 (37,396 for the "Top 1" sample). The number of hs6 products and destination countries is lower than the ones available in the original French and Chinese customs datasets since we need that the top 1 (to top 10) French exporting firm exports the same hs6 product that the top 1 (to top 10) Chinese exporting firm to at least the reference country as well as the destination country. The total number of hs6 products in the estimating sample is 2649. The same restriction applies to destination countries. We manage to keep 74 such destination countries. Table A.2 in the appendix presents descriptive statistics of the main variables for the countries present in the estimating sample.

4 Estimates of the demand side parameter

4.1 Graphical illustration

Before estimation, we turn to describing graphically the relationship between export flows and applied tariffs tetrads for different destination countries across products. In the interest of parsimony we focus on two major reference importer countries (k is Germany or Japan) and a restricted set of destination countries (n is USA and Canada). We calculate for each product p the tetradic terms for exports ranked $j = 1$ to 10th, $\ln \widetilde{x}^p_{\{j,n,k\}}$, and the relevant tetradic term for applied tariffs $\ln \widetilde{(1 + t^p_{\{n,k\}})}$.

Figure 2 report these tetrad terms to document the raw (and unconditional) evidence of the effect of tariffs on exported values by individual firms. The graphs also display the regression line and estimated coefficients of this simple regression of the logged export tetrad on the log of tariff tetrad for each of the six destination countries. Each point corresponds to a given hs6 product, and we highlight the cases where the export tetrad is calculated out of the largest ($j = 1$) French and Chinese exporters with a circle. The observations corresponding to Germany as a reference importer country are marked by a triangle, when the symbol is a square for Japan.

These estimations exploit the variation across products on tariffs applied by the destination country n and reference importer country k to China and France. In all cases, the estimated coefficient on tariff is negative and highly significant as shown by the slope of the line reported in each of each graphs. Those coefficients are quite large in absolute value, denoting a very steep response of consumers to differences in applied tariffs.

Figure 2: Unconditional tetrad evidence: individual importers

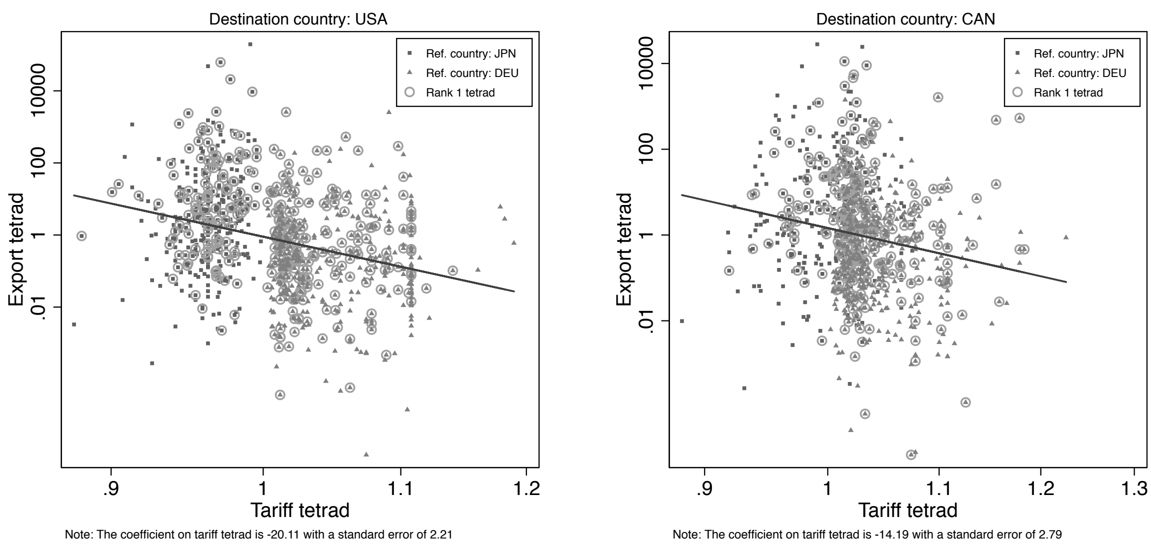

Figure 3 exposes a different dimension of identification, by looking at the impact of tariffs for specific products. We graph, following the logic of Figure 2 the tetrad of export value against the tetrad of tariffs for two individual products, which are the ones for which we maximize the number of observations in the dataset. Again those sectors exhibit strong reaction to tariff differences across importing countries. A synthesis of this evidence for individual sectors can be found by averaging tetrads over a larger set of products. Appendix figure A.1 provides this synthesis for the 184 products that have at least 30 destinations in common in our sample for French and Chinese exporters. The coefficient is again very large in absolute value and highly significant. The next section presents regression results with the full sample, both dimensions of identification, and the appropriate set of gravity control variables which will confirm this descriptive evidence and, as expected reduce the steepness of the estimated response.

Figure 3: Unconditional tetrad evidence: individual products

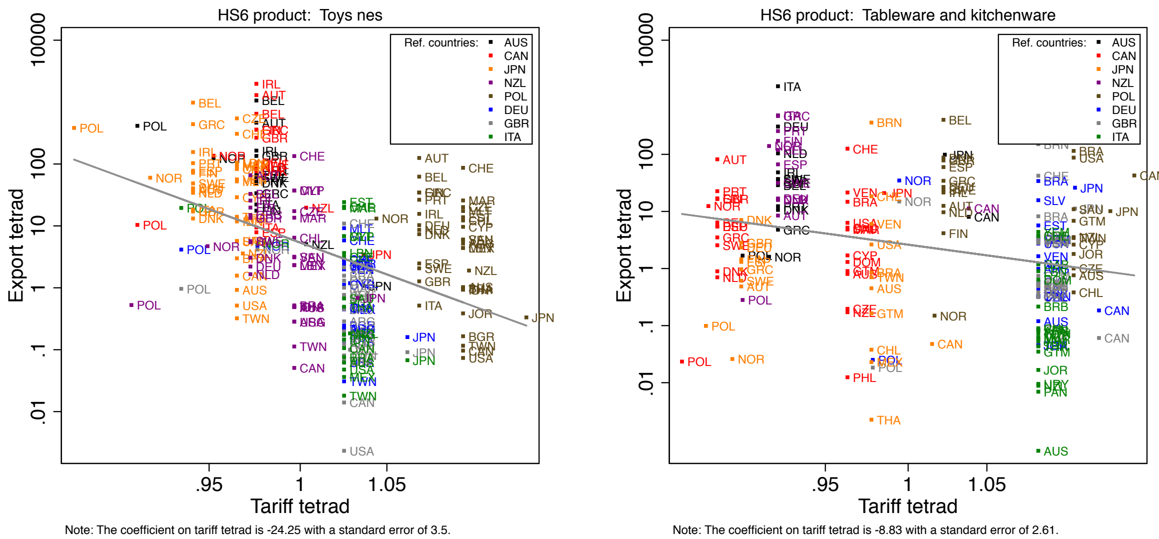

4.2 The intensive margin

This section presents estimates of equation (8) for all reference importer countries pooled in the same specification. In all specifications standard errors are clustered by destination×reference country.

Estimations in Table 1 exploit the variations in tariffs applied to France and China across both products and destination countries. Columns (1) to (3) show the results using as dependent variable the ratio of the top 1 exporting French and Chinese firm. Columns (2) presents estimations on the sample of positive tetraded tariffs and column (3) controls for the tetradic terms of Regional Trade Agreements (RTA). Columns (4) to (6) of Table 1 present the estimations using as dependent variable the ratio of firm-level exports of the top 1 to the top 10 French and Chinese firms. These estimations yield coefficients for the applied tariffs $(1-\sigma)$ that range between -5.74 and -2.66. Note that in both cases, the coefficients on applied tariffs are reduced when including the RTA as expected, but that the tariff variable retains statistical significance, showing that the effect of tariffs is not restricted to the binary impact of going from positive to zero tariffs.

In Table 2 we focus on the variations in tariffs within product across destination countries. Thus, all specifications in this table include (hs6-product × reference country) fixed effects. The coefficients for the applied tariffs $(1-\sigma)$ range from -6 and -3.2 for the pair of the Top 1 exporting French and Chinese firms (columns (1) to (3)). In the extended sample, applied tariffs vary from -4.1 to -1.65. While RTA has a positive and significant effect, it again does not capture the whole effect of tariff variations across destination countries on export flows. Note also that distance and contiguity have the usual (expected) signs and exhibit very high significance, while the presence of a colonial link and of a common language has a much more volatile influence.

In Appendix A.1.2. we present a number of alternative specifications of the intensive margin estimates. First we restrict the sample to destination countries applying non-MFN tariffs to France and China. (Australia, Canada, Japan, New Zealand and Poland) Table A.5 displays the results, where

Table 1: Intensive margin elasticities in 2000.

Dependent variable:	Top 1 firm-level exports			Top 1 to 10 firm-level exports		
	(1)	(2)	(3)	(4)	(5)	(6)
Applied Tariff	-5.74[a]	-4.83[a]	-3.83[a]	-4.54[a]	-4.65[a]	-2.66[a]
	(0.76)	(0.81)	(0.71)	(0.60)	(0.61)	(0.54)
Distance	-0.47[a]	-0.46[a]	-0.15[a]	-0.50[a]	-0.45[a]	-0.19[a]
	(0.03)	(0.03)	(0.04)	(0.02)	(0.02)	(0.03)
Contiguity	0.58[a]	0.75[a]	0.52[a]	0.60[a]	0.75[a]	0.54[a]
	(0.08)	(0.08)	(0.07)	(0.08)	(0.07)	(0.07)
Colony	0.27	0.63[c]	-0.24	-0.07	0.24	-0.61[a]
	(0.29)	(0.32)	(0.29)	(0.15)	(0.18)	(0.15)
Common language	0.10	-0.09	0.39[a]	0.08	-0.10	0.39[a]
	(0.09)	(0.09)	(0.09)	(0.08)	(0.07)	(0.07)
RTA			1.06[a]			1.07[a]
			(0.12)			(0.09)
Observations	37396	15477	37396	99645	41376	99645
R^2	0.137	0.189	0.143	0.146	0.181	0.153
rmse	2.99	3.01	2.98	3.08	3.12	3.06

Notes: Standard errors are clustered by destination×reference country. Applied tariff is the tetradic term of the logarithm of applied tariff plus one. Columns (2) and (5) present estimations on the sample of positive tetraded tariffs. [a], [b] and [c] denote statistical significance levels of one, five and ten percent respectively.

Table 2: Intensive margin elasticities in 2000. Within-product estimations.

Dependent variable:	Top 1 firm-level exports			Top 1 to 10 firm-level exports		
	(1)	(2)	(3)	(4)	(5)	(6)
Applied Tariff	-5.99[a]	-5.47[a]	-3.20[a]	-4.07[a]	-3.09[a]	-1.65[b]
	(0.79)	(1.07)	(0.79)	(0.72)	(0.75)	(0.68)
Distance	-0.54[a]	-0.49[a]	-0.21[a]	-0.59[a]	-0.55[a]	-0.29[a]
	(0.03)	(0.03)	(0.04)	(0.03)	(0.03)	(0.03)
Contiguity	0.93[a]	0.97[a]	0.84[a]	1.00[a]	0.94[a]	0.93[a]
	(0.08)	(0.09)	(0.07)	(0.07)	(0.09)	(0.07)
Colony	0.56[a]	0.48[c]	0.01	0.13	0.18	-0.34[a]
	(0.21)	(0.29)	(0.21)	(0.10)	(0.15)	(0.11)
Common language	-0.03	-0.00	0.25[a]	-0.07	-0.07	0.18[a]
	(0.07)	(0.08)	(0.06)	(0.06)	(0.07)	(0.06)
RTA			1.08[a]			0.94[a]
			(0.11)			(0.07)
Observations	37396	15477	37396	99645	41376	99645
R^2	0.145	0.128	0.153	0.140	0.115	0.146
rmse	2.14	1.99	2.13	2.42	2.26	2.41

Notes: Standard errors are clustered by destination×reference country. All estimations include (hs6-product×reference country) fixed effects. Applied tariff is the tetradic term of the logarithm of applied tariff plus one. Columns (2) and (5) present estimations on the sample of positive tetraded tariffs. [a], [b] and [c] denote statistical significance levels of one, five and ten percent respectively.

the coefficient on applied tariffs is statistically significant with a magnitude ranging from -5.47 to -3.24. Second, we exploit variations in applied tariffs within destination countries across hs6-products as an alternative dimension of identification (Table A.6). By contrast, our baseline estimations exploit variation of applied tariffs within hs6 products across destination countries and exporters (firms located in France and China). The results are robust to this new source of identification. Third, we complement the cross-sectional analysis of our baseline specifications—undertaken for the year 2000, i.e. before entry of China into WTO. We consider two additional cross-sectional samples, one after China entry into WTO (2001), the other for the final year of our sample (2006). Here again the results are qualitatively robust, although the coefficients on tariffs are lower since the difference of tariffs applied to France and China by destination countries is reduced after 2001. Fourth, we consider panel estimations over the 2000-2006 period. This analysis exploits the variations in tariffs within product-destination over time and across reference countries. The panel dimension allows for the inclusion of three sets of fixed effects: Product-destination, year and reference country. The coefficients of the intensive margin elasticity are close to the findings from the baseline cross-section estimations in 2000, and they range from -5.26 to -1.80.

Finally we address in the Appendix an econometric concern that is linked to endogenous selection into export markets. To understand the potential selection bias associated with estimating the trade elasticity it is useful to recall that selection is due to the presence of a fixed export cost that makes some firms unprofitable in some markets. Therefore higher tariff countries will be associated with firms having drawn a more favorable demand shock thus biasing downwards our estimate of the trade elasticity. Our approach of tetrads that focuses on highly ranked exporters for each hs6-market combination should however not be too sensitive to that issue, since those are firms that presumably have such a large productivity that their idiosyncratic destination shock is of second order.[13] In order to verify that intuition, we follow Eaton and Kortum (2001), applied to firm-level data by Crozet et al. (2012), yielding a generalized structural tobit. This method (EK tobit) keeps all individual exports to all possible destination markets (including zeroes). Strikingly, the EK tobit estimates are very comparable to our baseline tetrad estimates, giving us further confidence in an order of magnitude of the firm-level trade elasticity around located between -4 and -6.

5 Aggregate trade elasticities

The objective of this section is to provide a theory-consistent methodology for inferring, from firm-level data, the *aggregate* elasticity of trade with respect to trade costs. Given this objective, our methodology requires to account for the full distribution of firm-level productivity, i.e. we now need to add supply-side determinants of the trade elasticity to the demand-side aspects developed in previous sections (see equation 1). Following Head et al. (2014), we consider two alternative distributions—Pareto, as is standard in the literature, and log-normal—and we provide two sets of estimates, one for each considered distribution.[14] The Pareto assumption has this unique feature that the aggregate elasticity is constant, and depends only on the dispersion parameter of the Pareto, that is on supply only, a result first emphasized in Chaney (2008). Without Pareto, things are notably more complex, as the trade elasticity varies across country pairs. In addition, calculating this elasticity requires knowledge of the bilateral cost cutoff under which the considered country is unprofitable.

To calculate this bilateral cutoff, we combine our estimate of the demand side parameter $\hat{\sigma}$ with a dyadic micro-level observable, the *mean-to-min* ratio, that corresponds to the ratio of average over minimum sales of firms for a given country pair. In the model, this ratio measures the endogenous

[13] It might be the case that those top firms exhibit a different trade elasticity than the rest of the firms' population. The finding by Berman et al. (2012) that the reaction to exchange rate changes declines with productivity suggests that the estimates in this paper could be considered as a lower bound.

[14] Unless otherwise specified, Pareto is understood here as the *un-truncated version* used by most of the literature. See Helpman et al. (2008) and Melitz and Redding (2015) for results with the truncated version, where the trade elasticity recovers a bilateral dimension.

dispersion of cross-firm performance on a market, and more precisely the relative performance of entrants in this market following a change in our variable of interest: variable trade costs.

Under Pareto, the mean-to-min ratio, for a given origin, should be constant and independent of the size of the destination market. This pattern of scale-invariance is not observed in the data where we see that mean-to-min ratios increase massively in large markets—a feature consistent with a log-normal distribution of firm-level productivity. In the last step of the section we compare our micro-based predicted elasticities to those estimated with a gravity-like approach based on macro-data.

5.1 Quantifying aggregate trade elasticities from firm-level data: Theory

In order to obtain the theoretical predictions on aggregate trade elasticities, we start by summing, for each country pair, the sales equation (2) across all active firms:

$$X_{ni} = V_{ni} \times \left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma} (w_i \tau_{ni})^{1-\sigma} A_n M_i^e, \tag{9}$$

where M_i^e is the mass of entrant firms and V_{ni} denotes a cost-performance index of exporters located in country i and selling in n. This index is characterized by

$$V_{ni} \equiv \int_0^{a_{ni}^*} a^{1-\sigma} g(a) da, \tag{10}$$

where $a \equiv \alpha \times b(\alpha)$ corresponds to the unitary labor requirement *rescaled* by the firm-destination shock. In equation (10), $g(.)$ denotes the pdf of the rescaled unitary labor requirement and a_{ni}^* is the rescaled labor requirement of the cutoff firm. The solution for the cutoff is the cost satisfying the zero profit condition, i.e., $x_{ni}(a_{ni}^*) = \sigma w_i f_{ni}$. Using (2), this cutoff is characterized by

$$a_{ni}^* = \frac{1}{\tau_{ni} f_{ni}^{1/(\sigma-1)}} \left(\frac{1}{w_i}\right)^{\sigma/(\sigma-1)} \left(\frac{A_n}{\sigma}\right)^{1/(\sigma-1)}. \tag{11}$$

We are interested in the (partial) elasticity of aggregate trade value with-respect to variable trade costs, τ_{ni}. Partial means here holding constant origin-specific and destination-specific terms (income and price indices) as in Arkolakis et al. (2012) and Melitz and Redding (2015). In practical terms, the use of importer and exporter fixed effects in gravity regressions (the main source of estimates of the aggregate elasticity) holds w_i, M_i and A_n constant, so that, using (9), we have[15]

$$\varepsilon_{ni} \equiv \frac{d \ln X_{ni}}{d \ln \tau_{ni}} = 1 - \sigma - \gamma_{ni}, \tag{12}$$

which uses the fact that $d \ln a_{ni}^* / d \ln \tau_{ni} = -1$. In (12), γ_{ni} is a very useful term, studied by Arkolakis et al. (2012), describing how V_{ni} varies with an increase in the cutoff cost a_{ni}^*, that is an easier access of market n for firms in i:

$$\gamma_{ni} \equiv \frac{d \ln V_{ni}}{d \ln a_{ni}^*} = \frac{a_{ni}^{*2-\sigma} g(a_{ni}^*)}{V_{ni}}. \tag{13}$$

Equation (12) means that the aggregate trade elasticity may not be constant across country pairs because of the γ_{ni} term. In order to evaluate those bilateral trade elasticities, combining (13) with (10) reveals that we need to know the value of bilateral cutoffs a^*. In order to obtain those, we define the following function

$$\mathcal{H}(a^*) \equiv \frac{1}{a^{*1-\sigma}} \int_0^{a^*} a^{1-\sigma} \frac{g(a)}{G(a^*)} da, \tag{14}$$

[15] While this is literally true under Pareto because w_i, M_i and A_n enter a_{ni}^* multiplicatively, deviating from Pareto adds a potentially complex interaction term through a non-linear in logs effect of monadic terms on the dyadic cutoff. We expect this effect to be of second order.

a monotonic, invertible function which has a straightforward economic interpretation in this model. It is the ratio of average over minimum performance (measured as $a^{*1-\sigma}$) of firms located in i and exporting to n. Using equations (2) and (9), this ratio also corresponds to the observed mean-to-min ratio of sales:

$$\frac{\bar{x}_{ni}}{x_{ni}(a^*_{ni})} = \mathcal{H}(a^*_{ni}). \tag{15}$$

For our two origin countries (France and China), we observe the ratio of average to minimum trade flows for each destination country n. Using equation (15), one can calibrate $\hat{a}^*_{n\mathrm{FR}}$ and $\hat{a}^*_{n\mathrm{CN}}$, the estimated value of the export cutoff for French and Chinese firms exporting to n as a function of the mean-to-min ratio of French and Chinese sales on each destination market n

$$\hat{a}^*_{n\mathrm{FR}} = \mathcal{H}^{-1}\left(\frac{\bar{x}_{n\mathrm{FR}}}{x^{\mathtt{MIN}}_{n\mathrm{FR}}}\right), \qquad \text{and} \qquad \hat{a}^*_{n\mathrm{CN}} = \mathcal{H}^{-1}\left(\frac{\bar{x}_{n\mathrm{CN}}}{x^{\mathtt{MIN}}_{n\mathrm{CN}}}\right). \tag{16}$$

Equipped with the dyadic cutoffs we use equations (12) to (15) to obtain the aggregate trade elasticities

$$\varepsilon_{n\mathrm{FR}} = 1 - \hat{\sigma} - \frac{x^{\mathtt{MIN}}_{n\mathrm{FR}}}{\bar{x}_{n,\mathrm{FR}}} \times \frac{\hat{a}^*_{n\mathrm{FR}} g(\hat{a}^*_{n\mathrm{FR}})}{G(\hat{a}^*_{n\mathrm{FR}})}, \quad \text{and} \quad \varepsilon_{n\mathrm{CN}} = 1 - \hat{\sigma} - \frac{x^{\mathtt{MIN}}_{n\mathrm{CN}}}{\bar{x}_{n\mathrm{CN}}} \times \frac{\hat{a}^*_{n\mathrm{CN}} g(\hat{a}^*_{n\mathrm{CN}})}{G(\hat{a}^*_{n\mathrm{CN}})}, \tag{17}$$

where $\hat{\sigma}$ is our estimate of the intensive margin (the demand-side parameter) from previous sections. Our inference procedure is characterized by equations (16), and (17). We also calculate two other trade margins: the elasticity of the number of active exporters N_{ni} (the so-called extensive margin) and the elasticity of average shipments $\bar{x}_{ni}$. The number of active firms is closely related to the cutoff since $N_{ni} = M^e_i \times G(a^*_{ni})$, where M^e_i represents the mass of entrants (also absorbed by exporter fixed effects in gravity regressions). Differentiating the previous relationship and using (17) we can estimate the dyadic extensive margin of trade

$$\frac{d\ln N_{n\mathrm{FR}}}{d\ln \tau_{n\mathrm{FR}}} = -\frac{\hat{a}^*_{n\mathrm{FR}} g(\hat{a}^*_{n\mathrm{FR}})}{G(\hat{a}^*_{n\mathrm{FR}})}, \qquad \text{and} \qquad \frac{d\ln N_{n\mathrm{CN}}}{d\ln \tau_{n\mathrm{CN}}} = -\frac{\hat{a}^*_{n\mathrm{CN}} g(\hat{a}^*_{n\mathrm{CN}})}{G(\hat{a}^*_{n\mathrm{CN}})}, \tag{18}$$

From the accounting identity $X_{ni} \equiv N_{ni} \times \bar{x}_{ni}$, we obtain the (partial) elasticity of average shipments to trade simply as the difference between the estimated aggregate elasticities, (17) and the estimated extensive margins, (18).

$$\frac{d\ln \bar{x}_{n\mathrm{FR}}}{d\ln \tau_{n\mathrm{FR}}} = \varepsilon_{n\mathrm{FR}} - \frac{d\ln N_{n\mathrm{FR}}}{d\ln \tau_{n\mathrm{FR}}} \qquad \text{and} \qquad \frac{d\ln \bar{x}_{n\mathrm{CN}}}{d\ln \tau_{n\mathrm{CN}}} = \varepsilon_{n\mathrm{CN}} - \frac{d\ln N_{n\mathrm{CN}}}{d\ln \tau_{n\mathrm{CN}}}, \tag{19}$$

For the sake of interpreting the role of the mean-to-min, we combine (17) and (18) to obtain a relationship linking the aggregate elasticities to the (intensive and extensive) margins and to the mean-to-min ratio. Taking France as an origin country for instance, we obtain:

$$\varepsilon_{n\mathrm{FR}} = \underbrace{1-\hat{\sigma}}_{\text{intensive margin}} + \underbrace{\frac{1}{\bar{x}_{n\mathrm{FR}}/x^{\mathtt{MIN}}_{n\mathrm{FR}}}}_{\text{min-to-mean}} \times \underbrace{\frac{d\ln N_{n\mathrm{FR}}}{d\ln \tau_{n\mathrm{FR}}}}_{\text{extensive margin}}, \tag{20}$$

which is equation (1) presented in the introduction. This decomposition shows that the aggregate trade elasticity is the sum of the intensive margin and of the (weighted) extensive margin. The weight on the extensive margin depends only on the mean-to-min ratio, an observable measuring the dispersion of relative firm performance. Intuitively, the weight of the extensive margin should be decreasing when the market gets easier. Indeed easy markets have larger rates of entry, $G(a^*)$, and therefore increasing presence of weaker firms which augments dispersion measured as $\mathcal{H}(a^*_{ni})$. The marginal entrant in an easy market will therefore have less of an influence on aggregate exports, a smaller impact of the extensive margin. In the limit, the weight of the extensive margin becomes negligible and the whole of the aggregate elasticity is due to the intensive margin / demand parameter. In the Pareto case however this mechanism is not operational since $\mathcal{H}(a^*_{ni})$ and therefore the weight of the extensive margin is constant. We now turn to implementing our method with Pareto as opposed to an alternative distribution yielding non-constant dispersion of sales across destinations.

5.2 Quantifying aggregate trade elasticities from firm-level data: Results

A crucial step for our quantification procedure consists in specifying the distribution of rescaled labor requirement, $G(a)$, which is necessary to inverse the $\mathcal{H}$ function, reveal the bilateral cutoffs and obtain the bilateral trade elasticities. The literature has almost exclusively used the Pareto. Head et al. (2014) show that a credible alternative, which seems favored by firm-level export data, is the log-normal distribution. Pareto-distributed rescaled productivity $\varphi \equiv 1/a$ translates into a power law CDF for a, with shape parameter θ. A log-normal distribution of a retains the log-normality of productivity (with location parameter μ and dispersion parameter ν) but with a change in the log-mean parameter from μ to $-\mu$. The CDFs for a are therefore given by

$$G^{\mathrm{P}}(a) = \left(\frac{a}{\bar{a}}\right)^{\theta}, \qquad \text{and} \qquad G^{\mathrm{LN}}(a) = \Phi\left(\frac{\ln a + \mu}{\nu}\right), \tag{21}$$

where we use Φ to denote the CDF of the standard normal. Simple calculations using (21) in (14), and detailed in Appendix 2:, show that the resulting formulas for $\mathcal{H}$ are

$$\mathcal{H}^{\mathrm{P}}(a^*_{ni}) = \frac{\theta}{\theta - \sigma + 1}, \qquad \text{and} \qquad \mathcal{H}^{\mathrm{LN}}(a^*_{ni}) = \frac{h[(\ln a^*_{ni} + \mu)/\nu]}{h[(\ln a^*_{ni} + \mu)/\nu + (\sigma - 1)\nu]}, \tag{22}$$

where $h(x) \equiv \phi(x)/\Phi(x)$, the ratio of the PDF to the CDF of the standard normal.

Calculating $G^{\mathrm{P}}(.)$, $G^{\mathrm{LN}}(.)$, $\mathcal{H}^{\mathrm{P}}(.)$ and $\mathcal{H}^{\mathrm{LN}}(.)$ requires knowledge of underlying key supply-side distribution parameters θ and ν. For those, we use estimates from the Quantile-Quantile (QQ) regressions in Head et al. (2014). This method, based on a regression of empirical against theoretical quantiles of log sales, is applied on the same samples of exporters (Chinese and French) as here, and requires an estimate of the CES. We choose $\hat{\sigma} = 5$, which corresponds to a central value in our findings on the intensive margin above (where the average value of $1 - \hat{\sigma} \simeq -4$).

Panel (a) of Figure 4 depicts the theoretical relationship between the ratio of average to minimum sales, $\mathcal{H}(a^*_{ni})$, and the probability of serving the destination market , $G(a^*_{ni})$, spanning over values of the cutoff a^*_{ni}. Under Pareto heterogeneity, $\mathcal{H}$ is constant but this property of scale invariance is specific to the Pareto: Indeed it is increasing in G under log-normal. Panel (b) of figure 4 depicts the empirical counterpart of this relationship as observed for French and Chinese exporters in 2000 for all countries in the world. On the x-axis is the share of exporters serving each of those markets.[16] Immediately apparent is the non-constant nature of the mean-to-min ratio in the data, contradicting the Pareto prediction. This finding is very robust when considering alternatives to the minimum sales (which might be noisy if only because of statistical threshold effects) for the denominator of $\mathcal{H}$, that is different quantiles of the export distribution.[17]

[16] While this is not exactly the empirical counterpart of $G(a^*_{ni})$, the x-axis of panel (a), those two shares differ by a multiplicative constant, leaving the *shape* of the (logged) relationship unchanged.

[17] In a further effort to minimize noise in the calculation of the mean-to-min ratio, the figures are calculated for each of the 99 HS2 product categories and averaged. In the rest of the section, we will stick with this approach for the calculation of elasticities, done at the sector level before being averaged, which also simplifies exposition.

Figure 4: Theoretical and Empirical Mean-to-Min ratios

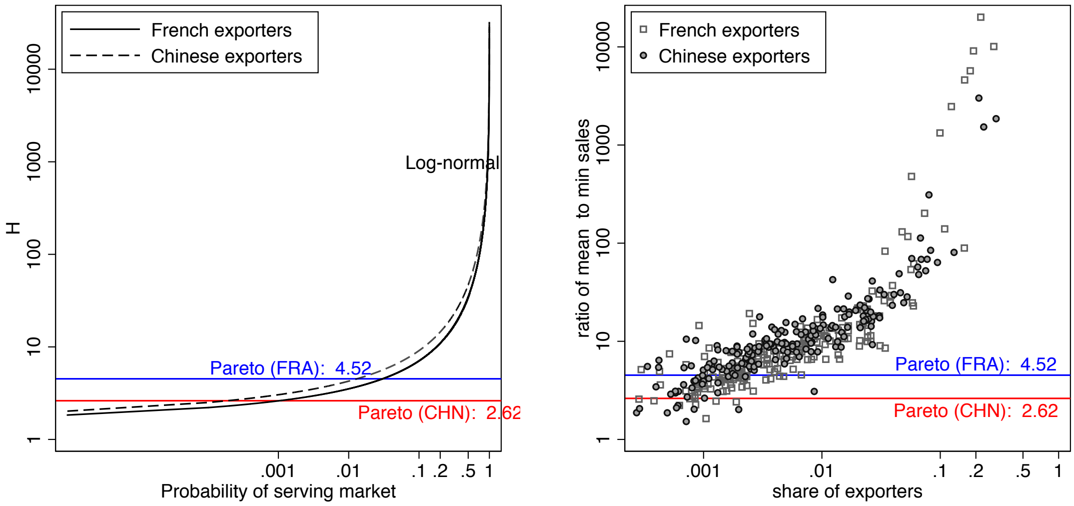

(a) theory (b) data

Figure 5 turns to the predicted trade elasticities under the two alternative distributions. Functional forms from (21) and (22) into (17), are used to deliver the two aggregate elasticities $\varepsilon^{\mathrm{P}}_{ni}$ and $\varepsilon^{\mathrm{LN}}_{ni}$:

$$\varepsilon^{\mathrm{P}}_{ni} = -\theta, \qquad \text{and} \qquad \varepsilon^{\mathrm{LN}}_{ni} = 1 - \sigma - \frac{1}{\nu} h \left(\frac{\ln a^*_{ni} + \mu}{\nu} + (\sigma - 1)\nu \right). \tag{23}$$

Parallel to figure 4, panel (a) of figure 5 shows the theoretical relationship between those elasticities and $G(a^*_{ni})$, while panel (b) plots the same elasticities evaluated for each individual destination country against the empirical counterpart of $G()$. Again, the Pareto case has a constant prediction (one for each exporter), while log normal predicts a trade elasticity that is declining (in absolute value) with easiness of the market. Panel (b) confirms the large variance of trade elasticities according to the share of exporters that are active in each of the markets. It also shows that the response of aggregate flows to trade costs is reduced (in absolute value) when the market becomes easier. The intuition is that for very difficult markets, the individual reaction of incumbent firms is supplemented with entry of exporters selected among the most efficient firms. The latter effect becomes negligible for the easiest markets, yielding ε to approach the intensive margin. This mechanism becomes very clear when looking at the patterns of the extensive margin and average export elasticities in figure 6.

The predicted elasticity on the extensive margin is also rising with market toughness as shown in panel (a) of figure 6. The inverse relationship is true for average exports (panel b). When a market is very easy and most exporters make it there, the extensive margin goes to zero, and the response of average exports goes to the value of the intensive margin (the firm-level response), $1 - \sigma$, as shown in figure 6 when the share of exporters increase. While this should intuitively be true in general, Pareto does not allow for this change in elasticities across markets, since the response of average exports should be uniformly 0, while the total response is entirely due to the (constant) extensive margin. In Table 3, we compute the average value and standard deviation of bilateral trade elasticities calculated using log-normal, and presented in figures 5 and 6. The first column presents the statistics for the French exporters' sample, the second one is the Chinese exporters' case, and the last column averages those. The mean elasticities obtained vary slightly between France and China, but the dominant feature is that the total elasticity is in neither case confined to the extensive margin. In both cases, average exports are predicted to react strongly to trade costs, a pattern we will confirm on actual data in the next subsection.

Figure 5: Predicted trade elasticities: $\varepsilon_{n\mathrm{FR}}$ and $\varepsilon_{n\mathrm{CN}}$

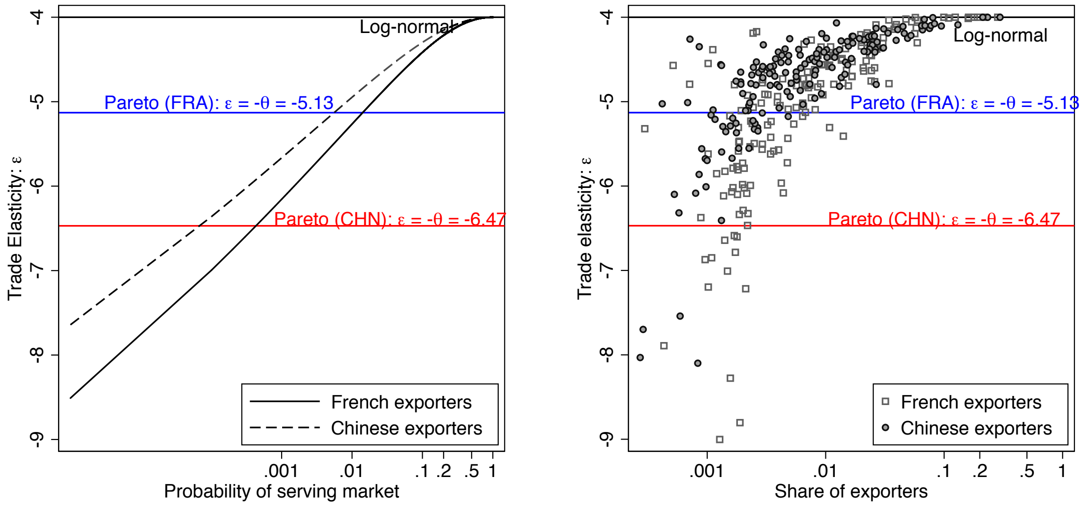

(a) theory

(b) data

Figure 6: Predicted elasticities: extensive and average exports

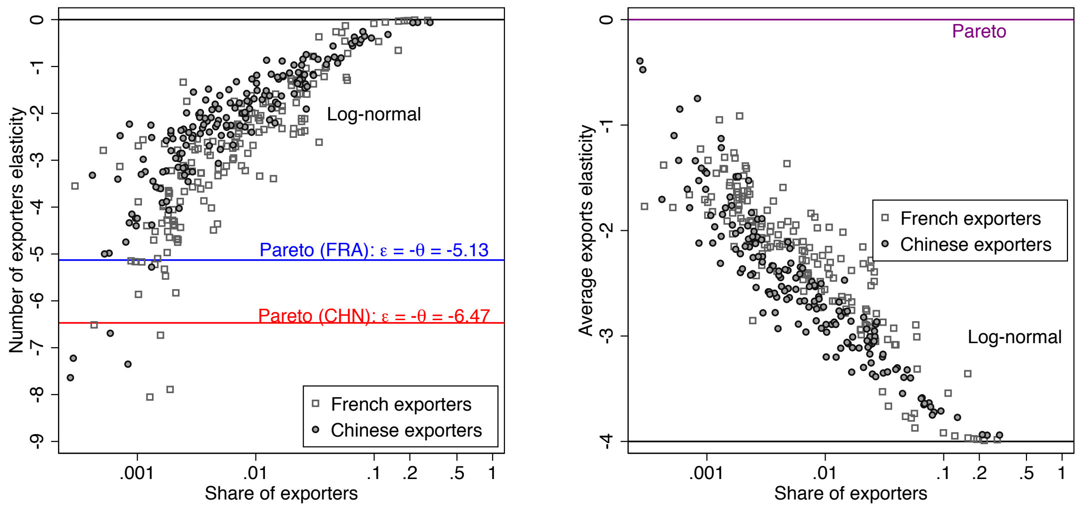

(a) extensive margin

(b) average exports

Table 3: Predicted bilateral trade elasticities (LN distribution)

LHS	France	China	Average
Total flows	-5.14	-4.792	-4.966
	(1.069)	(.788)	(.742)
Number of exporters	-2.866	-2.274	-2.57
	(1.657)	(1.472)	(1.335)
Average flows	-2.274	-2.517	-2.396
	(.687)	(.731)	(.64)

Notes: This table presents the predicted elasticities (mean and s.d.) on total exports, the number of exporting firms, and average export flows. Required parameters are σ, the CES, and ν, the dispersion parameter of the log normal distribution.

Figure 7 groups our bilateral trade elasticities (ε_{ni}) into ten bins of export shares for both France and China in a way similar to empirical evidence by Novy (2013), which reports that the aggregate trade cost elasticity decreases with bilateral trade intensity.[18] The qualitative pattern is very similar here, with the bilateral elasticity decreasing in absolute value with the share of exports going to a destination. One can use this variance in ε_{ni} to quantify the error that a practitioner would make when assuming a constant response of exports to a trade liberalization episode. Taking China as an example, decreasing trade costs by one percent would raise flows by around 6.5 percent for countries like Somalia, Chad or Azerbaijan (first bin of Chinese exports) and slightly more than 4 percent for the USA and Japan (top bin). Since the estimate that would be obtained when imposing a unique elasticity would be close to the average elasticity (4.79), this would entail about 25 percent underestimate of the trade growth for initially low traders (1.7/6.5) and an overestimate of around 20 percent (.8/4) for the top trade pairs.[19]

[18] Although Novy (2013) estimates variable distance elasticity, his section 3.4 assumes a constant trade costs to distance parameter to focus on the equivalent of our ε_{ni}.

[19] We thank Steve Redding for suggesting this quantification.

Figure 7: Variance in trade elasticities: $\varepsilon_{n\mathrm{FR}}$ and $\varepsilon_{n\mathrm{CN}}$

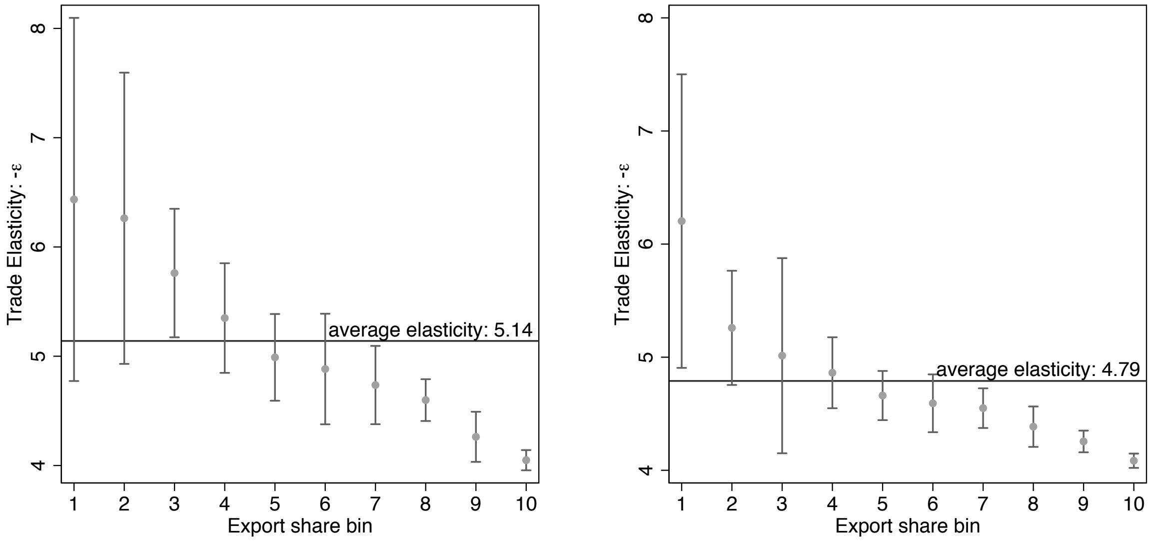

(a) French exports (b) Chinese exports

5.3 Comparison with macro-based estimates of trade elasticities

We now can turn to empirical estimates of aggregate elasticities to be compared with our predictions. Those are obtained using aggregate versions of our estimating tetrad equations presented above, which is very comparable to the method most often used in the literature: a gravity equation with country fixed effects and a set of bilateral trade costs covariates, on which a constant trade elasticity is assumed.[20] Column (1) of Table 4 uses the same sample of product-markets as in our benchmark firm-level estimations and runs the regression on the tetrad of aggregate rather than individual exports. Column (2) uses the same covariates but on the count of exporters, and column (3) completes the estimation by looking at the effects on average flows. An important finding is that the effect on average trade flow is estimated at -2.55, and is significant at the 1% level, contrary to the Pareto prediction (in which no variable trade cost should enter the equation for average flows).[21] This finding is robust to controlling for RTA (column 6) or constraining the sample to positive tariffs (column 9). The estimated median trade elasticity on total flows over all specifications at -4.79, is very close from the -5.03 found as the median estimate in the literature by Head and Mayer (2014).

Under Pareto, the aggregate elasticity should reflect fully the one on the number of exporters, and there should be no impact of tariffs on average exports. This prediction of the Pareto distribution is therefore strongly contradicted by our results. As a first pass at assessing whether, the data support the log-normal predictions, we compare the (unique) macro-based elasticity obtained in Table 4, with the corresponding average of bilateral elasticities shown in Table 3 of the preceding sub-section. The numbers obtained are quite comparable when the effects of RTAs are taken into account (columns (4) to (6)) or with

[20]Note that the gravity prediction on aggregate flows where origin, destination, and bilateral variables are multiplicatively separable and where there is a unique trade elasticity is only valid under Pareto. The heterogeneous elasticities generated by deviating from Pareto invalidate the usual gravity specification. Our intuition however is that the elasticity estimated using gravity/tetrads should be a reasonable approximation of the *average* bilateral elasticities. In order to verify this intuition, we run Monte Carlo simulations of the model with log-normal heterogeneity and find that indeed the average of micro-based heterogeneous elasticities is very close to the unique macro-based estimate in a gravity/tetrads equation on aggregate flows. Description of those simulations are in Appendix 3:.

[21]Note that the three dependent variables are computed for each hs6 product-destination, and therefore that the average exports do not contain an extensive margin where number of products would vary across destinations.

Table 4: Elasticites of total flows, count of exporters and average trade flows.

	Tot. (1)	# exp. (2)	Avg. (3)	Tot. (4)	# exp. (5)	Avg. (6)	Tot. (7)	# exp. (8)	Avg. (9)
Applied Tariff	-6.84^a	-4.29^a	-2.55^a	-4.00^a	-1.60^b	-2.41^a	-4.79^a	-2.13^a	-2.66^a
	(0.82)	(0.66)	(0.54)	(0.73)	(0.63)	(0.50)	(0.84)	(0.50)	(0.54)
Distance	-0.85^a	-0.61^a	-0.24^a	-0.51^a	-0.28^a	-0.23^a	-0.85^a	-0.60^a	-0.25^a
	(0.04)	(0.03)	(0.02)	(0.05)	(0.03)	(0.03)	(0.04)	(0.03)	(0.02)
Contiguity	0.62^a	0.30^a	0.32^a	0.53^a	0.21^a	0.32^a	0.64^a	0.35^a	0.29^a
	(0.12)	(0.09)	(0.06)	(0.11)	(0.07)	(0.05)	(0.12)	(0.09)	(0.06)
Colony	0.93^a	0.72^a	0.20^a	0.38^a	0.20^b	0.17^b	1.12^a	0.94^a	0.17^b
	(0.11)	(0.10)	(0.06)	(0.13)	(0.10)	(0.08)	(0.14)	(0.11)	(0.07)
Common language	0.09	0.16^c	-0.07	0.39^a	0.44^a	-0.05	-0.02	0.05	-0.07
	(0.09)	(0.09)	(0.06)	(0.09)	(0.08)	(0.06)	(0.10)	(0.08)	(0.06)
RTA				1.10^a	1.04^a	0.06			
				(0.11)	(0.06)	(0.07)			
Observations	99645	99645	99645	99645	99645	99645	41376	41376	41376
R^2	0.319	0.537	0.063	0.331	0.575	0.063	0.311	0.505	0.066
rmse	1.79	0.79	1.47	1.77	0.76	1.47	1.65	0.75	1.34

Notes: All estimations include fixed effects for each product-reference importer country combination. Standard errors are clustered at the destination-reference importer level. The dependent variable is the tetradic term of the logarithm of total exports at the hs6-destination-origin country level in columns (1), (4) and (7); of the number of exporting firms by hs6-destination and origin country in columns (2), (5) and (8) and of the average exports at the hs6-destination-origin country level in columns (3), (6) and (9). Applied tariff is the tetradic term of the logarithm of applied tariff plus one. Columns (7) to (9) present the estimations on the sample of positive tetraded tariffs and non-MFN tariffs. ^{a}, ^{b} and ^{c} denote statistical significance levels of one, five and ten percent respectively.

positive tetrad tariffs (columns (7) to (9)). Although this is not a definitive validation of the heterogenous firms model with log-normal distribution, our results clearly favor this distributional assumption over Pareto, and provides support for the empirical relevance of non-constant trade elasticities.

5.4 Direct evidence of non-constant trade elasticities

We can further use tetrads on aggregate trade flows in order to show *direct* empirical evidence of non-constant trade elasticities. Using aggregate bilateral flows from equation (9), and building tetrads with a procedure identical to the one used in Section 2.2, we obtain the (FR,CN, n, k)–tetrad of aggregate exports

$$\widetilde{X}_{\{n,k\}} \equiv \frac{X_{n\mathrm{FR}}/X_{k\mathrm{FR}}}{X_{n\mathrm{CN}}/X_{k\mathrm{CN}}} = \left(\frac{\tau_{n\mathrm{FR}}/\tau_{k\mathrm{FR}}}{\tau_{n\mathrm{CN}}/\tau_{k\mathrm{CN}}}\right)^{1-\sigma} \times \frac{V_{n\mathrm{FR}}/V_{k\mathrm{FR}}}{V_{n\mathrm{CN}}/V_{k\mathrm{CN}}} \tag{24}$$

Taking logs, differentiating with respect to tariffs and using the expression for the cutoff (11), we obtain

$$\begin{aligned} d\ln \widetilde{X}_{\{n,k\}} &= (1-\sigma-\gamma_{n\mathrm{FR}}) \times d\ln \tau_{n\mathrm{FR}} - (1-\sigma-\gamma_{k\mathrm{FR}}) \times d\ln \tau_{k\mathrm{FR}} \\ &\quad - (1-\sigma-\gamma_{n\mathrm{CN}}) \times d\ln \tau_{n\mathrm{CN}} + (1-\sigma-\gamma_{k\mathrm{CN}}) \times d\ln \tau_{k\mathrm{CN}}, \end{aligned} \tag{25}$$

where γ_{ni} is the elasticity of the cost performance index to a rise in the easiness of the market, defined in (13). For general distributions of heterogeneity, this elasticity is not constant across dyads as it depends on the dyad-specific cutoff a^*_{ni}. Hence, our interpretation of equation (25) is that the contribution to the (tetraded) total exports of a change in bilateral tariffs is larger for dyads that have a larger elasticity. Under Pareto, this elasticity is constant across dyads, $\gamma^{\mathrm{P}}_{ni} = 1-\sigma+\theta$. Combined with equation (25) this leads to

$$\widetilde{\varepsilon}^{\mathrm{P}}_{\{n,k\}} = \frac{d\ln \widetilde{X}_{\{n,k\}}}{d\ln \widetilde{\tau}_{\{n,k\}}} = -\theta, \tag{26}$$

where $\widetilde{\tau}_{\{n,k\}}$ is the vector of tetraded trade costs. This formula states that under Pareto, the elasticity of tetrad exports to tetrad tariffs is equal to the supply-side parameter θ. This transposes to the tetrad environment the well-known result of Chaney (2008) on gravity. Under non-Pareto heterogeneity, the four elasticities in (25) will remain different, a prediction we can put to a test. Results are shown in Table 5, where we pool observations for the years 2000 to 2006. Columns (1), (2) and (3) are the equivalent of the first three columns in Table 4, with the trade costs tetrads being split into its four components and the coefficients allowed to differ. The coefficients on tariffs to the destination country n show that the elasticity when considering France and China as an origin country differ significantly, consistent with the non-Pareto version of heterogeneity. Coefficients related to the reference importer k also differ significantly from each other, supporting further heterogeneity in the trade elasticities. A related approach is to confine identification on the destination country, neutralizing the change of reference country with a k fixed effect. Those results are shown in columns (4) to (6), where again most of the tariff elasticities differ across origin countries.[22]

[22]Table A.12 shows those same estimations for the two extreme years of our sample, 2000 and 2006, with significant evidence of non-constant elasticities in most cases.

Table 5: Non-constant trade elasticity

Dependent variable:	Tot.	# exp.	Avg.	Tot.	# exp.	Avg.
	(1)	(2)	(3)	(4)	(5)	(6)
Applied Tariff$_{n,\text{FR}}$	-4.25^a	-2.90^a	-1.34^a	-4.06^a	-2.76^a	-1.30^a
	(0.27)	(0.25)	(0.18)	(0.26)	(0.23)	(0.17)
Applied Tariff$_{n,\text{CN}}$	3.43^a	1.87^a	1.56^a	3.30^a	1.76^a	1.53^a
	(0.27)	(0.25)	(0.17)	(0.26)	(0.23)	(0.17)
Applied Tariff$_{k,\text{FR}}$	7.11^a	6.60^a	0.52^b			
	(0.36)	(0.20)	(0.24)			
Applied Tariff$_{k,\text{CN}}$	-3.79^a	-2.14^a	-1.66^a			
	(0.40)	(0.26)	(0.23)			
Observations	1077652	1077652	1077652	1085643	1085643	1085643
R^2	0.346	0.587	0.080	0.349	0.593	0.081
rmse	2.41	1.01	2.05	2.41	1.01	2.05

Notes: All estimations include a product and year fixed effects and the four components (n, FR; n, CN; k, FR; and k, CN) of each gravity control (distance, common language, contiguity and colony). In all estimations standard errors are clustered at the destination-reference country and year level.

5.5 Micro and Aggregate elasticities at the industry level

As a last exercise, we provide evidence that both demand and supply determinants enter the aggregate elasticity by looking at industry-level estimates. For each good, we can estimate a firm-level and an aggregate elasticity to tariffs (as in sections 4.2 and 5.3 respectively). Under the Pareto assumption, those two elasticities have no reason to be correlated, since the micro elasticity is a measure of product differentiation, while the macro one is capturing homogeneity in firms' productive efficiency. Under alternative distributions like the log-normal, the aggregate elasticity includes both determinants and therefore should be correlated with the micro one (equation 1).

We run our micro and macro-level tetrad estimations for each 2-digit ISIC industry separately including destination-reference country and year fixed effects. Table 6 presents the results. Columns (1) and (2) show the coefficients for the micro-level elasticity while columns (3) and (4) report the estimates of the aggregate elasticity using the tetrad term of total exports by product-destination-reference country and year. Columns (2) and (4) restrict the sample to EU destinations. Each cell reports the coefficient on the applied tariffs tetrad by industry with associated degree of statistical significance.

Estimates at the industry level yield coefficients of the intensive margin elasticity that average to -2.67 (column 1). The coefficients of the aggregate elasticity have a mean of -3.22 (column 3). More important for our main investigation, the intensive and aggregate elasticities are correlated (pairwise correlations are .68 for the full sample, and .74 for the EU one). Figure 8 shows graphical evidence of those correlations that exhibits overwhelming evidence in favor of the aggregate elasticity including demand side determinants.

Table 6: Micro and Aggregate elasticities by industry: 2000-2006

	Micro		Aggregate	
Dependent variable:	firm-level exports		total exports	
	(1)	(2)	(3)	(4)
	Full	EU	Full	EU
Agriculture	-4.85^{a}	-3.77	-4.66^{a}	-4.9^{c}
Food	-2.1^{a}	-3.09^{a}	-3^{a}	-3.57^{a}
Textile	-1.85^{a}	-5.74^{a}	-1.84^{a}	-2.97^{b}
Wearing	-3.2^{a}	-3.23^{a}	-3.4^{a}	-3.64^{a}
Leather	-4.13^{a}	-5.96^{a}	-6.84^{a}	-8.72^{a}
Wood	-9.4^{a}	-20.88^{a}	-9.59^{a}	-15.33^{a}
Paper	8.39^{a}	6.35	6.64^{b}	8.34
Edition	1.68	2.2	5.02^{b}	3.86
Coke prod	.02	1.39	.93	1.21
Chemical	-4.99^{a}	-7.29^{a}	-5.64^{a}	-7.66^{a}
Rubber	5.12^{a}	4.8^{a}	12.18^{a}	16.25^{a}
Basic metal	-4.55^{a}	-7.35^{a}	-3.17^{c}	-6.27^{c}
Metal products	-2.9^{a}	-6.33^{a}	-1.11^{c}	-3.32^{a}
Machinery	-3.18^{a}	-6.92^{a}	-2.89^{a}	-5.06^{a}
Office	-.06	-4.88	-21.14^{a}	-31.42^{b}
Electrical Prod	-2.49^{b}	-8.44^{a}	-.59	-1.34
Equip. Radio, TV	-8.38^{a}	-10.66^{a}	-7.08^{a}	-8.1^{a}
Medical instruments	-2.36^{a}	-3.9^{b}	-.38	-.29
Vehicles	-4.14^{c}	12.69^{c}	-5.38^{a}	15.52^{a}
Transport	-11.34^{a}	-18.28^{a}	-16.15^{a}	-17.37^{a}
Furniture	-1.35^{b}	-3.75^{a}	.26	0

Notes: All estimations are run by industry 2 digit. The cells report the coefficient on the applied tariffs tetrad by industry. All estimations include destination-reference country and year fixed effects. Standard errors are clustered by product-reference country and year. All estimations include a constant that is not reported. Applied tariff is the tetradic term of the logarithm of applied tariff plus one. ^{a}, ^{b} and ^{c} denote statistical significance levels of one, five and ten percent respectively.

Figure 8: Aggregate and intensive margin elasticities by industry 2 digit

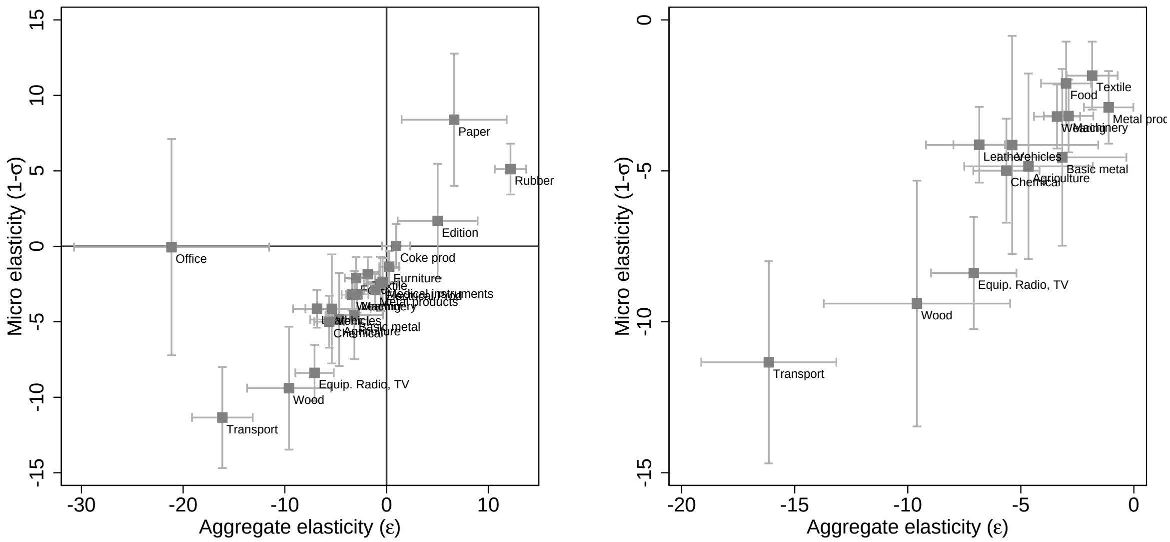

(a) Full set

(b) Coeff sig. < 0

6 Conclusion

We have argued in this paper that knowledge of the firm-level response to trade costs is a central element to our understanding of aggregate export reaction. In other words, we need micro-level data to understand the macro-level impacts of trade costs, a central element in any trade policy evaluation. This need for micro data is presumably true with the vast majority of possible heterogeneity distribution assumptions. There is one exception however where micro data is not needed to estimate the aggregate elasticity: the (untruncated) Pareto distribution. It is an exception the literature has been concentrating on for reasons of tractability that are perfectly legitimate, but the evidence presented in our paper points to systematic variation in bilateral aggregate trade elasticities that is both substantial and compatible with log-normal heterogeneity (in addition to be strongly preferred when looking at the micro-level distribution of export sales). We therefore call for a "micro approach" to estimating those elasticities as opposed to the "macro approach" using gravity specified so as to estimate a constant elasticity.

The micro- and macro- approaches differ substantially in several respects. On the one hand, gravity is a more direct and parsimonious route for estimating aggregate elasticities: (i) parametric assumptions are reduced to a minimum while our micro-based procedure depends on the calibration of the productivity distribution; (ii) gravity is less demanding in terms of data and makes possible the use of easily accessible dataset of bilateral aggregate trade flows. On the other hand, gravity provides, for each origin country, only a cross-destination average of elasticities while the micro-based approach provides the full cross-dyadic distribution of elasticities. Given this last limitation, we use our gravity estimates of averaged elasticities as a benchmark for discriminating between the two distributional assumptions made in our micro-based quantifications. We find that average value of bilateral trade elasticities obtained under a log-normal calibration is very close to the empirical gravity estimate which constrains the elasticity to be constant across country pairs. By contrast, the Pareto-based calibration leads to predictions that seem invalidated by the data. Namely, the invariance of average shipments to ad-valorem tariff variations, the lack of correlation between firm-level and aggregate level of elasticities estimated industry by industry, and the constant aggregate trade elasticities.

References

Anderson, J. E. and van Wincoop, E. (2003). Gravity with gravitas: A solution to the border puzzle. *American Economic Review*, 93 (1):170–192.

Arkolakis, C. (2010). Market penetration costs and the new consumers margin in international trade. *Journal of Political Economy*, 118(6):1151 – 1199.

Arkolakis, C., Costinot, A., and Rodríguez-Clare, A. (2010). Gains from trade under monopolistic competition: A simple example with translog expenditure functions and pareto distributions of firm-level productivity. manuscript.

Arkolakis, C., Costinot, A., and Rodriguez-Clare, A. (2012). New trade models, same old gains? *American Economic Review*, 102(1):94–130.

Atkeson, A. and Burstein, A. (2008). Pricing-to-market, trade costs, and international relative prices. *The American Economic Review*, 98(5):1998–2031.

Baier, S. and Bergstrand, J. H. (2001). The growth of world trade: tariffs, transport costs, and income similarity. *Journal of International Economics*, 53 (1):1–27.

Barba Navaretti, G., Bugamellli, M., Forlani, E., and Ottaviano, G. (2015). It takes (more than) a moment: Revisiting the link between firm productivity and aggregate exports. University of Milan Mimeo.

Berman, N., Martin, P., and Mayer, T. (2012). How do different exporters react to exchange rate changes? *The Quarterly Journal of Economics*, 127 (1):437–492.

Berthou, A. and Fontagné, L. (2015). Variable trade costs, composition effects, and the intensive margin of trade. Technical Report 543, Banque de France.

Broda, C. and Weinstein, D. (2006). Globalization and the gains from variety. *Quarterly Journal of Economics*, 121 (2):541–585.

Caliendo, L. and Parro, F. (2015). Estimates of the trade and welfare effects of nafta. *Review of Economic Studies*, 82(1):1–44.

Chaney, T. (2008). Distorted gravity: The intensive and extensive margins of international trade. *American Economic Review*, 98 (4):1707–1721.

Costinot, A., Donaldson, D., and Komunjer, I. (2012). What goods do countries trade? a quantitative exploration of ricardo's ideas. *Review of Economic Studies*, 79(2):581–608.

Crozet, M., Head, K., and Mayer, T. (2012). Quality sorting and trade: Firm-level evidence for french wine. *Review of Economic Studies*, 79(2):609 – 644.

Eaton, J. and Kortum, S. (2001). Trade in capital goods. *European Economic Review*, 45(7):1195–1235.

Eaton, J. and Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, 70 (5):1741–1779.

Eaton, J., Kortum, S., and Kramarz, F. (2011). An anatomy of international trade: Evidence from French firms. *Econometrica*, 79(5):1453–1498.

Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *American Economic Review*, 84:157–177.

Feenstra, R. C. (2003). A homothetic utility function for monopolistic competition models. *Economics Letters*, 78(2):79–86.

Feenstra, R. C. (2013). Restoring the product variety and pro-competitive gains from trade with heterogeneous firms and bounded productivity. UC Davis Mimeo.

Fernandes, A., Klenow, P., Meleshchuk, S., Pierola, M. D., and Rodriguez-Clare, U. A. (2015). The intensive margin puzzle in trade. mimeo.

Fitzgerald, D. and Haller, S. (2014). Exporters and shocks: Dissecting the international elasticity puzzle. manuscript.

Freund, C. and Pierola, M. D. (2015). Export superstars. *Review of Economics and Statistics*.

Gopinath, G. and Neiman, B. (2014). Trade adjustment and productivity in large crises. *American Economic Review*, 104(3):793–831.

Hallak, J. C. (2006). Product quality and the direction of trade. *Journal of International Economics*, 68 (1):238–265.

Head, K. and Mayer, T. (2014). Gravity equations: Workhorse, toolkit, and cookbook. In Helpman, E., Gopinath, G., and Rogoff, K., editors, *Handbook of International Economics*, volume 4. Elsevier.

Head, K., Mayer, T., and Ries, J. (2010). The erosion of colonial trade linkages after independence. *Journal of International Economics*, 81 (1):1–14.

Head, K., Mayer, T., and Thoenig, M. (2014). Welfare and trade without pareto. *American Economic Review*, 104(5):310–16.

Head, K. and Ries, J. (2001). Increasing returns versus national product differentiation as an explanation for the pattern of us-canada trade. *American Economic Review*, 91(4):858–876.

Helpman, E., Melitz, M., and Rubinstein, Y. (2008). Estimating trade flows: Trading partners and trading volumes. *Quarterly Journal of Economics*, 123(2):441–487.

Hummels, D. (1999). Towards a geography of trade costs. Technical Report 17, GTAP Working Paper.

Imbs, J. and Méjean, I. (2015). Elasticity optimism. *American Economic Journal: Macroeconomics*, 7 (3):43–83.

Krugman, P. (1980). Scale economies, product differentiation, and the pattern of trade. *American Economic Review*, 70:950–959.

Melitz, M. (2003). The impact of trade on intra-industry reallocations and agregate industry productivity. *Econometrica*, 71:1695–1725.

Melitz, M. and Ottaviano, G. (2008). Market size, trade, and productivity. *Review of Economic Studies*, 75(1):295–316.

Melitz, M. J. and Redding, S. J. (2015). New trade models, new welfare implications. *American Economic Review*.

Novy, D. (2013). International trade without CES: Estimating translog gravity. *Journal of International Economics*, 89(2):271–282.

Ossa, R. (2015). Why trade matters after all. *Journal of International Economics*.

Romalis, J. (2007). Nafta's and cusfta's impact on international trade. *Review of Economics and Statistics*, 89(3):416 − 435.

Simonovska, I. and Waugh, M. E. (2011). The elasticity of trade: Estimates and evidence. *Working Paper 16796, NBER.*

Yang, M.-J. (2014). Firm heterogeneity and trade in a distorted gravity world. University of Washington Mimeo.

Appendix 1: Empirical Appendix

A.1.1. Descriptive statistics

Table A.1: Average percentage point difference between the applied tariff to France and China across industries by Germany and Japan (2000)

Reference importer:	Germany		Japan	
	Full sample	Tetrad regression sample	Full sample	Tetrad regression sample
Agriculture	-3.27	-2.98	.01	.02
Food	-7.83	-9.63	1.24	1.45
Textile	-7.14	-6.95	5.31	4.58
Wearing apparel	-9.34	-8.11	6.2	6.79
Leather	-1.5	-1.15	8.4	4.68
Wood	-1.36	-1.47	2.66	3.69
Paper	0	0	1.39	1.3
Edition	-.79	-.65	.26	.64
Coke prod	0	0	.97	1.73
Chemical	-1.01	-.74	2.51	2.4
Rubber & Plastic	-1.37	-1.25	2.5	2.37
Non Metallic	-1.43	-2.18	1.16	.8
Basic metal products	-1.84	-3.56	1.89	2.75
Metal products	-.67	-.99	1.38	1.32
Machinery	-.25	-.3	.19	0
Office	-.16	-.38	0	0
Electrical Prod	-.38	-.47	.37	.62
Equip. Radio, TV	-1.72	-1.79	0	0
Medical instruments	-.58	-.41	.14	.34
Vehicles	-2.22	-1.63	0	0
Transport	-1.27	-1.37	0	0
Furniture	-.51	-.63	1.93	1.95

Table A.2 reports population and GDP for each destination country in 2000, as well as the ratios of total exports, average exports, total number of exporting firms, and distance between France and China. Only 12 countries in our estimating sample are closer to China than to France. In all of those, the number of Chinese exporters is larger than the number of French exporters, and the total value of Chinese exports largely exceeds the French one. On the other end of the spectrum, countries like Belgium and Switzerland witness much larger counts of exporters and total flows from France than from China, as expected.

Table A.2: Destination countries characteristics in 2000

	Population	GDP	Ratio France / China: Total exports	Average exports	Number exporters	Distance
CHE	7	246	23.03	1.27	18.1	.06
BEL	10	232	9.63	1.21	7.94	.06
NLD	16	387	2.04	1.01	2.02	.08
GBR	60	1443	4.89	2.37	2.06	.09
ESP	40	581	14.37	3.82	3.76	.1
DEU	82	1900	5.04	2.13	2.37	.1
ITA	57	1097	7.32	2.71	2.71	.11
AUT	8	194	10.68	1.83	5.83	.12
IRL	4	96	8.52	1.38	6.19	.12
CZE	10	57	5.56	2.48	2.24	.13
PRT	10	113	17.95	1.98	9.07	.13
DNK	5	160	3.15	1.07	2.94	.16
MAR	28	33	10.08	1.57	6.43	.16
POL	38	171	3.93	1.68	2.34	.18
MLT	0	4	8.29	5.18	1.6	.18
SWE	9	242	5.91	2.66	2.22	.22
NOR	4	167	2.47	1.66	1.49	.22
BGR	8	13	4.31	2.19	1.97	.24
GRC	11	115	4.41	1.75	2.52	.24
BLR	10	13	1.19	.22	5.29	.28
MDA	4	1	97.59	6.15	15.88	.28
EST	1	6	1.39	.45	3.07	.3
FIN	5	121	1.9	.87	2.2	.32
GHA	20	5	1.06	1.99	.53	.38
NGA	125	46	1.15	2.32	.5	.38
CYP	1	9	1.43	1.27	1.13	.39
LBN	4	17	3.19	2.06	1.55	.43
JOR	5	8	.92	2.31	.4	.45
GAB	1	5	78.18	2.17	36.09	.46
BRB	0	3	2.01	1.79	1.12	.47
SUR	0	1	1.38	5.22	.26	.47
ATG	0	1	5.22	1.08	4.82	.48
GUY	1	1	1.42	8.22	.17	.48
VCT	0	0	4.19	.98	4.27	.48
BRA	174	644	1.77	1.87	.94	.5
VEN	24	117	1.23	2.51	.49	.52
DOM	9	20	1.21	1.58	.77	.52
PRY	5	7	.31	1.18	.27	.54
BOL	8	8	2.89	2.36	1.22	.55
JAM	3	8	1.38	5.35	.26	.56
URY	3	21	.59	2.1	.28	.57
COL	42	84	1.66	2.36	.7	.57
ARG	37	284	1.72	2.94	.59	.57
CUB	11	.	1.04	1.41	.74	.58
PER	26	53	.82	1.97	.42	.6
PAN	3	12	.09	1.14	.08	.6
CHL	15	75	1.12	4.23	.27	.61
UGA	24	6	2.74	3.98	.69	.62
CRI	4	16	1.4	3.95	.35	.62
CAN	31	714	.73	1.51	.48	.62
SAU	21	188	1.21	2.52	.48	.62
NIC	5	4	.19	.74	.25	.63
HND	6	6	.39	1.56	.25	.64
SLV	6	13	3.86	11.98	.32	.65
GTM	11	19	.25	1	.25	.67
USA	282	9765	.54	1.13	.48	.67
KEN	31	13	1.09	2.14	.51	.68
YEM	18	9	.83	3.39	.25	.7
TZA	34	9	.28	1.26	.22	.71
IRN	64	101	1	1.56	.64	.72
MEX	98	581	.93	1.31	.71	.75
SYC	0	1	19.33	8.66	2.23	.98
LKA	19	16	1.2	7.96	.15	1.72
NZL	4	53	.54	1.96	.28	1.84
AUS	19	400	.38	1.73	.22	1.98
NPL	24	5	.08	.42	.2	2.24
IDN	206	165	.12	.94	.12	2.47
BGD	129	47	.14	1.91	.07	2.7
THA	61	123	.32	1.04	.31	3.17
BRN	0	4	.58	3.05	.19	3.22
LAO	5	2	.25	.69	.36	3.83
PHL	76	76	.4	2.49	.16	4.21
JPN	127	4650	.12	.6	.19	4.96
TWN	22	321	.35	1.32	.26	6.69

Notes: Population is expressed in millions and GDP in billions of US dollars.

Table A.3: Avg. difference between tariffs applied to France and China. Full sample

	France < China		France = China		France > China	
	Tariff	# HS6	Tariff	# HS6	Tariff	# HS6
ARG	.	0	0	5113	.	0
ATG	.	0	0	5097	.	0
AUS	.	0	0	4188	1.91	905
AUT	-5.7	2137	0	2825	.	0
BEL	-5.7	2137	0	2825	.	0
BGD	.	0	0	5106	.	0
BGR	.	0	0	5059	.	0
BLR	.	0	0	4559	.	0
BOL	.	0	0	5113	.	0
BRA	.	0	0	5113	.	0
BRB	.	0	0	2020	.	0
BRN	.	0	0	5079	.	0
CAN	-3.87	15	0	2877	3.07	2178
CHE	-10	1	0	4120	.	0
CHL	.	0	0	5113	.	0
COL	.	0	0	5113	.	0
CRI	.	0	0	5113	.	0
CUB	.	0	0	5112	.	0
CYP	.	0	0	4929	.	0
CZE	.	0	0	5113	.	0
DEU	-5.7	2137	0	2825	.	0
DNK	-5.7	2137	0	2825	.	0
DOM	.	0	0	5008	.	0
ESP	-5.7	2137	0	2825	.	0
EST	.	0	0	5113	.	0
FIN	-5.7	2137	0	2825	.	0
GAB	.	0	0	5108	.	0
GBR	-5.7	2137	0	2825	.	0
GHA	.	0	0	5019	.	0
GRC	-5.7	2137	0	2825	.	0
GTM	.	0	0	5113	.	0
GUY	.	0	0	2043	.	0
HND	.	0	0	5113	.	0
IDN	.	0	0	5110	.	0
IRL	-5.7	2137	0	2825	.	0
IRN	.	0	0	5113	.	0
ITA	-5.7	2137	0	2825	.	0
JAM	.	0	0	5113	.	0
JOR	.	0	0	5085	.	0
JPN	-.18	3	0	2771	4.06	2256
KEN	.	0	0	4554	.	0
LAO	.	0	0	4977	.	0
LBN	.	0	0	5067	.	0
LKA	.	0	0	5090	.	0
MAR	.	0	0	5113	.	0
MDA	.	0	0	5068	.	0
MEX	.	0	0	5084	.	0
MLT	.	0	0	5109	.	0
NGA	.	0	0	5113	.	0
NIC	.	0	0	5113	.	0
NLD	-5.7	2137	0	2825	.	0
NOR	-10.74	1210	0	3560	.	0
NPL	.	0	0	5096	.	0
NZL	.	0	0	3220	1.15	1876
PAN	.	0	0	5110	.	0
PER	.	0	0	5113	.	0
PHL	.	0	0	5112	.	0
POL	-9.51	4234	0	485	7.14	388
PRT	-5.7	2137	0	2825	.	0
PRY	.	0	0	5113	.	0
SAU	.	0	0	4799	.	0
SLV	.	0	0	5113	.	0
SUR	.	0	0	1170	.	0
SWE	-5.7	2137	0	2825	.	0
SYC	.	0	0	4849	.	0
THA	.	0	0	5056	.	0
TWN	.	0	0	5113	.	0
TZA	.	0	0	5113	.	0
UGA	.	0	0	5110	.	0
URY	.	0	0	4829	.	0
USA	.	0	0	4768	.	0
VCT	.	0	0	2040	.	0
VEN	.	0	0	5109	.	0
YEM	.	0	0	5111	.	0

Notes: The table reports the average difference across hs6 products of applied tariffs by destination country n to France and China and the corresponding number of hs6 products when the tariff applied to France is lower than to China (columns (1) and (2)), when the applied tariff to both origin countries is the equal (columns (3) and (4)) and when the tariff applied to France is higher than to China (columns (5) and (6)).

Table A.4: Avg. difference between tariffs applied to France and China. Tetrad sample.

	France < China		France = China		France > China	
	Tariff	# HS6	Tariff	# HS6	Tariff	# HS6
ARG	.	0	0	567	.	0
ATG	.	0	0	1	.	0
AUS	.	0	0	816	1.87	249
AUT	-6.06	357	0	377	.	0
BEL	-5.43	651	0	772	.	0
BGD	.	0	0	73	.	0
BGR	.	0	0	121	.	0
BLR	.	0	0	9	.	0
BOL	.	0	0	11	.	0
BRA	.	0	0	636	.	0
BRB	.	0	0	5	.	0
BRN	.	0	0	6	.	0
CAN	.	0	0	424	2.78	571
CHE	-10	1	0	658	.	0
CHL	.	0	0	419	.	0
COL	.	0	0	183	.	0
CRI	.	0	0	37	.	0
CUB	.	0	0	25	.	0
CYP	.	0	0	255	.	0
CZE	.	0	0	434	.	0
DEU	-5.58	831	0	982	.	0
DNK	-5.74	465	0	500	.	0
DOM	.	0	0	56	.	0
ESP	-5.41	720	0	859	.	0
EST	.	0	0	52	.	0
FIN	-5.47	367	0	412	.	0
GAB	.	0	0	24	.	0
GBR	-5.55	802	0	964	.	0
GHA	.	0	0	37	.	0
GRC	-5.37	446	0	556	.	0
GTM	.	0	0	51	.	0
GUY	.	0	0	2	.	0
HND	.	0	0	19	.	0
IDN	.	0	0	457	.	0
IRL	-5.32	257	0	291	.	0
IRN	.	0	0	162	.	0
ITA	-5.43	760	0	895	.	0
JAM	.	0	0	17	.	0
JOR	.	0	0	206	.	0
JPN	-.01	1	0	631	4.09	535
KEN	.	0	0	74	.	0
LAO	.	0	0	1	.	0
LBN	.	0	0	401	.	0
LKA	.	0	0	128	.	0
MAR	.	0	0	407	.	0
MDA	.	0	0	3	.	0
MEX	.	0	0	533	.	0
MLT	.	0	0	130	.	0
NGA	.	0	0	116	.	0
NIC	.	0	0	5	.	0
NLD	-5.64	748	0	920	.	0
NOR	-8.03	267	0	248	.	0
NPL	.	0	0	15	.	0
NZL	.	0	0	189	1.09	308
PAN	.	0	0	125	.	0
PER	.	0	0	119	.	0
PHL	.	0	0	420	.	0
POL	-9.6	564	0	41	8.1	9
PRT	-4.78	396	0	426	.	0
PRY	.	0	0	45	.	0
SAU	.	0	0	512	.	0
SLV	.	0	0	27	.	0
SUR	.	0	0	1	.	0
SWE	-5.94	489	0	511	.	0
SYC	.	0	0	1	.	0
THA	.	0	0	662	.	0
TWN	.	0	0	864	.	0
TZA	.	0	0	13	.	0
UGA	.	0	0	6	.	0
URY	.	0	0	180	.	0
USA	.	0	0	1809	.	0
VCT	.	0	0	1	.	0
VEN	.	0	0	258	.	0
YEM	.	0	0	59	.	0

Notes: The table reports the average difference across hs6 products of applied tariffs by destination country n to France and China and the corresponding number of hs6 products when the tariff applied to France is lower than to China (columns (1) and (2)), when the applied tariff to both origin countries is the equal (columns (3) and (4)) and when the tariff applied to France is higher than to China (columns (5) and (6)).

A.1.2. Alternative specifications of the intensive margin estimates

A.1.2.1. Graphical evidence

Figure A.1: Unconditional tetrad evidence: averaged over top products

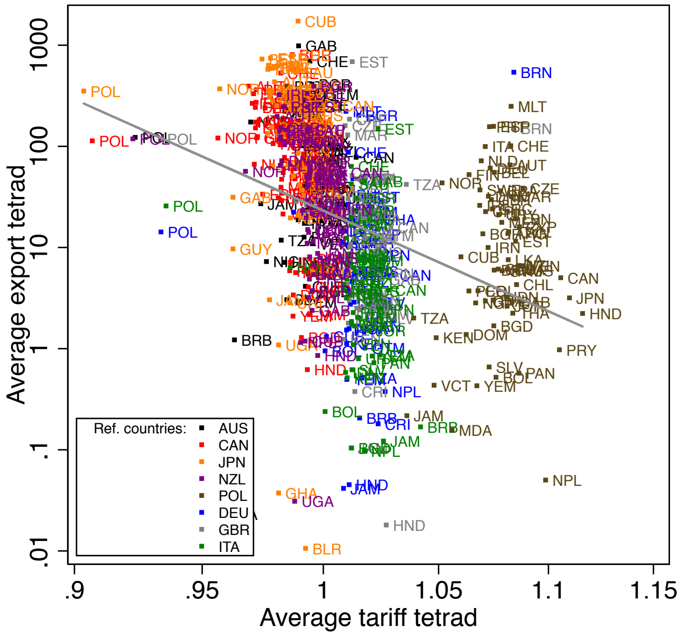

Note: Tetrads are averaged over the 184 products with at least 30 destinations in common.
The coefficient is -24.04 with a standard error of 2.89.

A.1.2.2. Non-MFN sample

As a more demanding specification, still identifying trade elasticity across destinations, we restrict the sample to destination countries applying non-MFN tariffs to France and China. The sample of such countries contains Australia, Canada, Japan, New Zealand and Poland.[23] Table A.5 displays the results. Common language, contiguity and colony are excluded from the estimation since there is no enough variance in the non-MFN sample. Our non-MFN sample also does not allow for including a RTA dummy. Estimations in columns (3) and (4) include fixed effect for each product×reference country. Columns (2) and (4) present estimations on the non-MFN sample of positive tetraded tariffs. In all cases, the coefficient of applied tariffs is negative and statistically significant with a magnitude from -5.47 to -3.24. Hence, in spite of the large reduction in sample size, the results are very comparable to those obtained on the full sample of tariffs.

A.1.2.3. Identification across products

Estimations on the intensive margin trade elasticity in the main text exploit variation of applied tariffs within hs6 products across destination countries and exporters (firms located in France and China). This section presents a set of estimations on alternative specifications that exploits the variation of applied tariffs within destination countries across hs6-products.

[23] To be on the conservative side, we exclude EU countries from the sample of non-MFN destinations since those share many other dimensions with France that might be correlated with the absence of tariffs (absence of Non-Tariff Barriers, free mobility of factors, etc.). Poland only enters the EU in 2004.

Table A.5: Intensive margin: non-MFN sample.

Dependent variable:	Top 1 to 10 firm-level exports			
	(1)	(2)	(3)	(4)
Applied Tariff	-3.87^a	-5.36^a	-3.24^a	-5.47^a
	(1.09)	(1.14)	(1.09)	(1.03)
Distance	-0.50^a	-0.41^a	-0.45^a	-0.36^a
	(0.03)	(0.03)	(0.05)	(0.05)
Observations	12992	9421	12992	9421
R^2	0.102	0.094	0.058	0.062
rmse	3.11	3.08	1.80	1.67

Notes: Standard errors are clustered by destination×reference country. Columns (3) and (4) include fixed effects at the (hs6 product×reference country) level. Applied tariff is the tetradic term of the logarithm of applied tariff plus one. Columns (2) and (4) present estimations on the sample of positive tetraded tariffs. ^{a}, ^{b} and ^{c} denote statistical significance levels of one, five and ten percent respectively.

Table A.6 reports the results from estimations including a destination-reference importer country fixed effect. In this case, standard errors are clustered by hs6-reference importer country. Including these fixed effects implies that the source of identification comes from variations within destination countries across hs6-products in applied tariffs to both origin countries, France and China, by the reference importer countries. Columns (1) and (3) present estimations on the full sample, while columns (2) and (4) report estimations on the sample of positive tetraded tariffs. The trade elasticity ranges from -3.57 to -5.07. Estimations in columns (5) and (6) restrict the destination countries to be the ones applying non-MFN duties and in column (7) to EU destination countries including countries with EU-trade agreements as Norway and Switzerland. The sample size drops radically, with the trade elasticities remaining of the expected sign and order of magnitude, but losing in statistical significance.

Table A.6: Intensive margin elasticities. Within-country estimations.

Dependent variable:	Top 1 firm-level exports		Top 1 to 10 firm-level exports				
Sample:	Full		Full		non-MFN		EU
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Applied Tariff	-3.53^a	-3.59^a	-3.10^a	-5.07^a	-2.93^a	-5.28^a	-2.48^a
	(1.07)	(1.26)	(0.70)	(0.83)	(0.83)	(0.94)	(0.61)
Observations	37396	15477	99645	41376	12992	9421	54198
R^2	0.001	0.002	0.001	0.003	0.001	0.004	0.000
rmse	2.92	2.93	3.02	3.05	3.09	3.06	2.95

Notes: All estimations include destination-reference importing country fixed effects Standard errors are clustered by hs6-reference importing country. Applied tariff is the tetradic term of the logarithm of applied tariff plus one. Columns (5) and (6) present the estimations for the non-MFN sample. Columns (2), (4) and (6) present estimations on the sample of positive tetraded tariffs. ^{a}, ^{b} and ^{c} denote statistical significance levels of one, five and ten percent respectively.

A.1.2.4. Cross-section estimations in 2001 and 2006

Results in the main text focus on cross-sectional analysis of the year 2000 (before entry of China into WTO). We now turn to additional cross-section evidence after China entry into WTO (2001) and for the final year of our sample (2006), in Table A.7. Estimations in columns (1) and (3) include a fixed effect identifying the product-reference country combination while columns (2) and (4) include a fixed effect identifying the destination-reference importing country. As expected the coefficients on tariffs are lower since the difference of tariffs applied to France and China by destination countries is reduced after 2001. The implied values of σ range from -3.6 to almost -2.

Table A.7: Intensive margin: cross-section 2001 and 2006

Dependent variable:	Top 1 to 10 firm-level exports			
	2001		2006	
	(1)	(2)	(3)	(4)
Applied Tariff	-1.83[a]	-2.59[a]	-0.95[b]	-1.42[b]
	(0.64)	(0.68)	(0.42)	(0.61)
Distance	-0.62[a]		-0.54[a]	
	(0.02)		(0.02)	
Contiguity	0.75[a]		0.98[a]	
	(0.08)		(0.06)	
Colony	-0.02		0.49[a]	
	(0.09)		(0.07)	
Common language	-0.13[b]		-0.03	
	(0.06)		(0.05)	
Observations	111039	111039	217732	217732
R^2	0.129	0.000	0.107	0.000
rmse	2.46	3.04	2.45	2.96

Notes: Estimations in columns (1) and (3) include a fixed effect identifying the hs6 product-reference importing country and standard errors are clustered by destination-reference importer country. Estimations in columns (2) and (4) include a fixed effect identifying the destination-reference importing country and standard errors are clustered by hs6 product-reference importer country. Applied tariff is the tetradic term of the logarithm of applied tariff plus one. [a], [b] and [c] denote statistical significance levels of one, five and ten percent respectively.

A.1.2.5. Panel estimations 2000-2006

Our dataset spans over the 2000-2006 period. This dimension allows to identify the variation of tariffs within product-destination over time and across reference countries. Table A.8 reports results. Columns (1) and (3) present the baseline estimation for the 2000-2006 period, columns (2) and (5) present estimations on the sample of non-MFN destinations and columns (3) and (6) on the sample of EU destination countries. All estimations include product-destination, year and importing reference country fixed effects. The coefficients of the intensive margin elasticity are close to the findings from the cross-section estimations of 2000, and they range from -5.26 to -1.80.

Table A.8: Intensive margin elasticities. Within-product-destination country estimations 2000-2006

Dependent variable:	Top 1 firm-level exports			Top 1 to 10 firm-level exports		
	(1)	(2)	(3)	(4)	(5)	(6)
Applied Tariff	-3.20^a	-2.47^a	-5.26^a	-1.99^a	-1.80^a	-3.66^a
	(0.40)	(0.62)	(0.57)	(0.25)	(0.39)	(0.36)
Distance	-0.38^a	-0.44^a	-0.28^a	-0.45^a	-0.51^a	-0.35^a
	(0.01)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)
Contiguity	0.28^a	0.23^a	0.31^a	0.53^a	0.52^a	0.31^a
	(0.03)	(0.04)	(0.03)	(0.02)	(0.03)	(0.02)
Common language	-0.31^a	-0.44^a	-0.24^a	-0.24^a	-0.35^a	-0.11^a
	(0.04)	(0.05)	(0.05)	(0.03)	(0.03)	(0.03)
Observations	379644	61882	198629	1077652	167758	558424
R^2	0.081	0.099	0.060	0.085	0.109	0.060
rmse	2.23	2.11	2.27	2.48	2.37	2.50

Notes: All estimations include hs6-destination country, reference importing country and year fixed effects. Standard errors are clustered by destination-reference importing country and year. Applied tariff is the tetradic term of the logarithm of applied tariff plus one. Columns (2) and (5) present estimations on the sample of non-MFN destinations and columns (3) and (6) on the sample of EU destination countries. ^{a}, ^{b} and ^{c} denote statistical significance levels of one, five and ten percent respectively.

A.1.3. High-dimensional fixed effects

Table A.9 re-runs our benchmark table of results (Table 1) with a method alternative to tetrads, using high-dimensional fixed effects (`reghdfe` in Stata) to control for the firm-level as well as destination-product determinants that enter the equation for individual export sales. Those two sets of fixed effects are denoted $\mathrm{FE}_i(\alpha)$ and FE_n^p. The estimated equation is therefore

$$\ln x_{ni}^p(\alpha) = \mathrm{FE}_i(\alpha) + \mathrm{FE}_n^p + (1-\sigma)\ln(1+t_{ni}^p) + (1-\sigma)\delta \ln D_n + \ln \epsilon_{ni}^p(\alpha), \tag{A.1}$$

which doubles the number of observations, since there is an independent observation for the flow of the French and of the Chinese exporter, when tetrads considers the ratio of those. Results point to reduced coefficients in absolute value, which still point to a substantial intensive margin, with statistically significant estimates of σ ranging between 2.6 and 4.5.

A.1.4. Selection bias

Not all firms export to all markets n, and the endogenous selection into different export destinations across firms is one of the core elements of the type of model we are using. To understand the potential selection bias associated with estimating the trade elasticity it is useful to recall the firm-level export equation (4):

$$\ln x_{ni}(\alpha) = (1-\sigma)\ln\left(\frac{\sigma}{\sigma-1}\right) + (1-\sigma)\ln(\alpha w_i) + (1-\sigma)\ln\tau_{ni} + \ln A_n + \ln \epsilon_{ni}(\alpha). \tag{A.2}$$

In this model, selection is due to the presence of a fixed export cost f_{ni} that makes some firms unprofitable in some markets. Assuming that fixed costs are paid using labor of the origin country, profits in this setup are given by $x_{ni}(\alpha)/\sigma - w_i f_{ni}$, which means that a firm is all the more likely to be present in market n that its $(1-\sigma)\ln(\alpha w_i) + (1-\sigma)\ln\tau_{ni} + \ln A_n + \ln\epsilon_{ni}(\alpha)$ is high. Therefore a firm with a low

Table A.9: Intensive margin elasticities in 2000, high-dimensional FEs.

Dependent variable:	Top 1 firm-level exports			Top 1 to 10 firm-level exports		
	(1)	(2)	(3)	(4)	(5)	(6)
Applied Tariff	-3.50^a	-3.27^a	-1.86^c	-1.56^a	-2.42^a	-0.56
	(0.97)	(1.27)	(0.99)	(0.60)	(0.73)	(0.61)
Distance	-0.38^a	-0.49^a	-0.16^a	-0.40^a	-0.48^a	-0.26^a
	(0.02)	(0.03)	(0.04)	(0.01)	(0.02)	(0.02)
Contiguity	0.72^a	0.91^a	0.66^a	0.69^a	0.80^a	0.65^a
	(0.05)	(0.08)	(0.05)	(0.04)	(0.05)	(0.04)
Colony	0.77^a	0.30	0.40^c	0.40^a	0.23	0.17
	(0.22)	(0.29)	(0.23)	(0.13)	(0.17)	(0.13)
Common language	0.08	-0.08	0.27^a	0.12^a	-0.04	0.25^a
	(0.06)	(0.08)	(0.06)	(0.04)	(0.05)	(0.04)
RTA			0.73^a			0.46^a
			(0.09)			(0.06)
Observations	74792	30954	74792	199290	82750	199290
R^2	0.772	0.808	0.773	0.713	0.753	0.714
RMSE	1.48	1.41	1.48	1.48	1.41	1.48

Notes: Standard errors in parentheses clustered by destination-product. Columns (2) and (5) present estimations on the sample of positive tetraded tariffs. ^{a}, ^{b} and ^{c} denote statistical significance levels of one, five and ten percent respectively.

cost (αw_i) can afford having a low draw on $\epsilon_{ni}(\alpha)$, creating a systematic bias on the cost variable. The same logic applies in attractive markets, (high A_n), which will be associated with lower average draws on the error term. Fortunately , our tetrad estimation technique removes the need to estimate αw_i and A_n, and therefore solves this issue.

However a similar problem arises with the trade cost variable, τ_{ni}, which is used to estimate the trade elasticity. Higher tariff countries will be associated with firms having drawn higher $\epsilon_{ni}(\alpha)$, thus biasing downwards our estimate of the trade elasticity. Our approach of tetrads that focuses on highly ranked exporters for each hs6-market combination should however not be too sensitive to that issue, since those are firms that presumably have such a large productivity that their idiosyncratic destination shock is of second order. In order to verify that intuition, we follow Eaton and Kortum (2001), applied to firm-level data by Crozet et al. (2012), who assume a normally distributed $\ln \epsilon_{ni}(\alpha)$, yielding a generalized structural tobit. This procedure uses the theoretical equation for minimum sales, $x_{ni}^{\mathtt{MIN}}(\alpha) = \sigma w_i f_{ni}$, which provides a natural estimate for the truncation point for each desination market. This method (EK tobit) keeps all individual exports to all possible destination markets (including zeroes).[24] When estimating equation (A.2), we proxy for $\ln A_n$ with GDP$_n$ and population$_n$, and for firm-level determinants α with the count of markets served by the firm. An origin country dummy for Chinese exporters account for all differences across the two groups, such as wages, w_i. Last, we ensure comparability by i) keeping the same sample of product-market combinations as in previous estimations using tetrads, ii) running the estimation with the same dimension of fixed effects (hs6). Each column of Tables A.10 and A.11 show the simple OLS (biased) estimates or the EK-tobit method run in the sample of product-market combinations by reference importing country. As in previous usages of that method, the OLS seems very severely biased, probably due to the extremely high selection levels observed (with all reference countries, slightly less than 14% of possible flows are observed). Strikingly, the EK tobit estimates are very comparable to the tetrad estimates shown until now , giving us further confidence in an order of magnitude of the firm-level trade

[24]For each product, we fill in with zero flows destinations that a firm found unprofitable to serve in reality. The set of potential destinations for that product is given all countries where at least one firm exported that good.

elasticity around located between -4 and -6.

Table A.10: Correcting for the selection bias.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ref. country:	Australia		Canada		Germany		Italy	
	OLS	EK Tobit	OLS	EK Tobit	OLS	EK Tobit	OLS	EK Tobit
Applied Tariff	1.26^a	-6.28^a	1.33^a	-5.05^a	2.09^a	-3.78^a	1.41^a	-6.39^a
	(0.16)	(0.87)	(0.18)	(0.93)	(0.17)	(0.71)	(0.16)	(0.74)
RTA	-0.67^a	1.37^a	-0.55^a	2.36^a	-0.63^a	2.27^a	-0.67^a	2.42^a
	(0.06)	(0.33)	(0.06)	(0.34)	(0.06)	(0.31)	(0.05)	(0.30)
Distance	0.01	-0.40^a	0.05^b	0.14	0.02	-0.21^c	0.04^b	-0.06
	(0.02)	(0.13)	(0.02)	(0.12)	(0.02)	(0.11)	(0.02)	(0.11)
Common language	0.07^c	3.32^a	0.30^a	4.95^a	0.21^a	3.94^a	0.22^a	4.14^a
	(0.04)	(0.20)	(0.04)	(0.23)	(0.03)	(0.17)	(0.03)	(0.17)
Contiguity	0.07^b	1.55^a	0.06^b	1.23^a	0.15^a	1.91^a	0.10^a	1.85^a
	(0.03)	(0.12)	(0.03)	(0.12)	(0.02)	(0.10)	(0.02)	(0.11)
Colony	0.64^a	4.29^a	0.69^a	2.50^a	0.85^a	2.65^a	0.79^a	2.23^a
	(0.10)	(0.55)	(0.10)	(0.55)	(0.09)	(0.43)	(0.08)	(0.43)
GDP_n	0.16^a	1.57^a	0.15^a	1.84^a	0.19^a	1.65^a	0.16^a	1.48^a
	(0.01)	(0.05)	(0.01)	(0.06)	(0.01)	(0.05)	(0.01)	(0.05)
Population_n	0.03^b	0.71^a	0.07^a	0.61^a	0.02	0.67^a	0.04^a	0.84^a
	(0.01)	(0.06)	(0.02)	(0.07)	(0.01)	(0.05)	(0.01)	(0.06)
Chinese exporter dummy	0.26^a	0.44^a	0.39^a	0.58^a	0.51^a	1.44^a	0.37^a	1.11^a
	(0.03)	(0.13)	(0.04)	(0.16)	(0.03)	(0.15)	(0.03)	(0.14)
# of dest. by firm	0.19^a	2.06^a	0.14^a	1.98^a	0.13^a	1.96^a	0.14^a	2.04^a
	(0.01)	(0.03)	(0.01)	(0.03)	(0.01)	(0.03)	(0.01)	(0.03)
Observations	617231	5518719	571508	4386238	730720	5828685	735346	5917116
R^2	0.058		0.051		0.059		0.053	
Pseudo R^2		0.082		0.086		0.075		0.077

Notes: All estimations include fixed effects for each hs6 product level. Standard errors are clustered at the hs6-destination-origin country level. Applied tariff is the logarithm of applied tariff plus one at the hs6 product level and destination country. a, b and c denote statistical significance levels of one, five and ten percent respectively.

Table A.11: Correcting for the selection bias.(cont.)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ref. country:	Japan		UK		New Zealand		Poland	
	OLS	EK Tobit	OLS	EK Tobit	OLS	EK Tobit	OLS	EK Tobit
Applied Tariff	1.13^a	-5.45^a	1.73^a	-4.99^a	0.62^a	-4.34^a	1.22^a	-4.60^a
	(0.23)	(1.09)	(0.16)	(0.78)	(0.22)	(1.14)	(0.20)	(0.87)
RTA	-0.28^a	3.12^a	-0.73^a	2.62^a	-0.51^a	1.65^a	-0.52^a	1.98^a
	(0.08)	(0.37)	(0.05)	(0.29)	(0.08)	(0.40)	(0.08)	(0.35)
Distance	0.18^a	0.38^a	0.05^a	-0.08	0.07^b	-0.18	0.08^a	-0.03
	(0.03)	(0.13)	(0.02)	(0.11)	(0.03)	(0.14)	(0.03)	(0.12)
Common language	0.28^a	3.70^a	0.25^a	4.31^a	0.31^a	4.01^a	0.25^a	4.21^a
	(0.04)	(0.20)	(0.03)	(0.16)	(0.06)	(0.29)	(0.05)	(0.24)
Contiguity	0.02	1.46^a	0.14^a	1.31^a	0.03	1.67^a	0.09^b	1.75^a
	(0.03)	(0.13)	(0.02)	(0.11)	(0.04)	(0.19)	(0.04)	(0.16)
Colony	0.56^a	3.91^a	0.73^a	2.51^a	0.64^a	4.17^a	0.71^a	3.41^a
	(0.11)	(0.79)	(0.08)	(0.51)	(0.11)	(0.98)	(0.09)	(0.57)
GDP_n	0.13^a	1.71^a	0.17^a	1.57^a	0.14^a	1.80^a	0.19^a	1.58^a
	(0.01)	(0.07)	(0.01)	(0.05)	(0.02)	(0.08)	(0.02)	(0.07)
$Population_n$	0.09^a	0.73^a	0.03^b	0.71^a	0.08^a	0.52^a	0.02	0.59^a
	(0.02)	(0.08)	(0.01)	(0.06)	(0.02)	(0.09)	(0.02)	(0.07)
Chinese exporter dummy	0.21^a	1.10^a	0.40^a	1.40^a	0.15^a	0.10	0.34^a	1.08^a
	(0.04)	(0.15)	(0.03)	(0.14)	(0.04)	(0.20)	(0.05)	(0.18)
# of dest. by firm	0.18^a	1.91^a	0.14^a	1.95^a	0.25^a	2.43^a	0.22^a	2.18^a
	(0.01)	(0.03)	(0.01)	(0.03)	(0.01)	(0.05)	(0.01)	(0.04)
Observations	414135	2732178	735485	5901262	343141	2864115	361503	2891673
R^2	0.046		0.061		0.048		0.052	

Notes: All estimations include fixed effects for each hs6 product level. Standard errors are clustered at the hs6-destination-origin country level. Applied tariff is the logarithm of applied tariff plus one at the hs6 product level and destination country. a, b and c denote statistical significance levels of one, five and ten percent respectively.

A.1.5. Non-constant trade elasticity

Table A.12: Non-constant trade elasticity

	2000			2006		
Dependent variable:	Tot.	# exp.	Avg.	Tot.	# exp.	Avg.
	(1)	(2)	(3)	(4)	(5)	(6)
Applied Tariff$_{n,\mathrm{FR}}$	-5.74[a]	-3.41[a]	-2.34[a]	-4.31[a]	-3.52[a]	-0.79[b]
	(1.02)	(0.69)	(0.72)	(0.53)	(0.57)	(0.37)
Applied Tariff$_{n,\mathrm{CN}}$	5.08[a]	2.55[a]	2.53[a]	2.16[a]	1.55[a]	0.61
	(0.99)	(0.71)	(0.66)	(0.57)	(0.58)	(0.38)
Observations	99745	99745	99745	218036	218036	218036
R^2	0.357	0.590	0.093	0.339	0.594	0.072
rmse	2.42	1.00	2.02	2.34	0.98	2.01

Notes: All estimations include a product and reference country fixed effects and the four components (nFR, nCN, kFR, and kCN) of each gravity control (distance, common language, contiguity and colony). In all estimations standard errors are clustered at the destination-reference country.

Appendix 2: Theoretical Mean-to-Min ratios under Pareto and Log-Normal distributions

In general, the shape of the distribution of firms' productivity matters for the aggregate trade elasticity, generating heterogeneous responses across country pairs to the same trade costs shock In this Appendix we consider two different distributions of the (rescaled) productivity: i) Pareto, which turns out to be a quite special case where heterogeneity washes out, and ii) Log-normal which maintains the mapping between heterogeneous productivity and heterogeneous trade elasticities.

More precisely, the central relationship (12) makes it clear that the heterogeneity of aggregate trade elasticity comes entirely from the term γ_{ni} that stems from endogenous selection of firms into export markets (see equation 13). In turn, γ_{ni} depends on the cost-performance index V_{ni} as defined by (10). We therefore need to understand how these γ and V terms behave under alternative distribution assumptions.

If productivity is Pareto then the rescaled unit input requirement a has PDF $g(a) = \theta a^{\theta-1}/\bar{a}^{\theta}$, which translates into

$$V_{ni}^{\mathrm{P}} = \frac{\theta a_{ni}^{*\theta-\sigma+1}}{\bar{a}^{\theta}(\theta-\sigma+1)}. \tag{A.3}$$

The elasticity of V_{ni} with respect to a^* is

$$\gamma_{ni}^{\mathrm{P}} = \theta - \sigma + 1 > 0. \tag{A.4}$$

Hence, Pareto makes all the γ_{ni} terms be the same, and therefore transforms an expression generally yielding heterogeneous trade elasticities into a one-parameter elasticity $\frac{d\ln X_{ni}}{d\ln \tau_{ni}} = \theta$, that is related to the supply side of the economy only.

When productive efficiency is distributed log-normally, things are very different. For $\varphi \sim$log-$\mathcal{N}(\mu,\nu)$, the distribution of rescaled unit input requirements is $a \sim$log-$\mathcal{N}(-\mu,\nu)$, and we can write

$$V_{ni}^{\mathrm{LN}} = \exp[(\sigma-1)\mu + (\sigma-1)^2\nu^2/2]\Phi[(\ln a_{ni}^* + \mu)/\nu + (\sigma-1)\nu], \tag{A.5}$$

where $\Phi()$ denotes the CDF of the standard normal distribution. Differentiating $\ln V_{ni}$ with respect to $\ln a_{ni}^*$,

$$\gamma_{ni}^{\mathrm{LN}} = \frac{1}{\nu} h\left(\frac{\ln a_{ni}^* + \mu}{\nu} + (\sigma-1)\nu\right), \tag{A.6}$$

where $h(x) \equiv \phi(x)/\Phi(x)$, the ratio of the PDF to the CDF of the standard normal. Thus γ_{ni} is no longer the constant $1-\sigma+\theta$ which obtains for productivity distibuted Pareto with shape parameter θ. Bilateral elasticities therefore write as

$$\varepsilon_{ni}^{\mathrm{P}} = -\theta, \qquad \text{and} \qquad \varepsilon_{ni}^{\mathrm{LN}} = 1 - \sigma - \frac{1}{\nu} h\left(\frac{\ln a_{ni}^* + \mu}{\nu} + (\sigma - 1)\nu\right). \tag{A.7}$$

The $\mathcal{H}$ function is a central element of our calibration procedure, as summarized by relationship (15), that reveals cutoffs and therefore aggregate bilateral elasticities. Comparing (10), (13) and (14) we see that $\mathcal{H}$ and γ are closely related

$$\gamma_{ni} \times \mathcal{H}(a_{ni}^*) = a_{ni}^* \frac{g(a_{ni}^*)}{G(a_{ni}^*)} \tag{A.8}$$

With Pareto, we make use of (A.4) to obtain

$$\mathcal{H}^{\mathrm{P}}(a_{ni}^*) = \frac{\theta}{\theta - \sigma + 1}, \tag{A.9}$$

With a log-normal productivity, equation (A.6) leads to

$$\mathcal{H}^{\mathrm{LN}}(a_{ni}^*) = \frac{h[(\ln a_{ni}^* + \mu)/\nu]}{h[(\ln a_{ni}^* + \mu)/\nu + (\sigma - 1)\nu]}, \tag{A.10}$$

An attractive feature of our quantification procedure relates to the small number of relevant parameters to be calibrated. Under Pareto, equations (A.4) and (A.9) show that only the shape parameter θ matters. Similarly, under a log-normal, only the calibration of the second-moment of the distribution, ν, is necessary for inverting the $\mathcal{H}$ function to reveal the cutoff and for quantifying the aggregate elasticity: This last point stems from the fact that shifting the first moment, μ, affects (A.6) and (A.10) in an identical way and so has no impact on the quantification.

Appendix 3: Monte Carlo Simulations

In Section 5.3 we find that the macro-based estimate of the aggregate trade elasticity is quantitatively close to the cross-dyadic average of the micro-based estimates when heterogeneity is calibrated as being log-normal. We interpret this finding as an empirical support in favor of this distributional assumption. In this section we substantiate this last statement by embracing a more theoretical perspective. This is an important step in the argument because the theoretical relationship between the macro- and the micro-based estimates of the elasticities is unknown (except under Pareto where they are unambiguously equal). Hereafter we provide simulation-based evidence that the similarity between micro- and macro-based estimates is not accidental, even under log-normal heterogeneity.

We proceed with Monte Carlo (MC) simulations of our generic trade model with heterogeneous firms. In the baseline simulations we generate fake bilateral trade for 10 countries and 1 million active firms per country. Our data generating process uses the firms' sales in equation (2). Firm-level heterogeneity in terms of rescaled labor requirement, $a \equiv \alpha \times b(\alpha)$ is assumed to be Pareto or Log-normally distributed with a set of parameters identical to the ones used in our empirical analysis (section 5.2). We also retain $\sigma = 5$ as the parameter for the intensive margin. Without loss of generality, in this partial equilibrium framework, we normalize the nominal wage, $w = 1$, and we draw A_n/f_{ni}, i.e. the dyadic ratio of destination n attractiveness over entry cost from a log-normal distribution. This distribution is calibrated such as to match an average dyadic share of exporting firms of 10 percent. Finally the applied-tariffs $\tau_{ni} = 1 + t_{ni}$ are drawn from an uniform distribution over the range $[1, 2]$.

In each MC draw, we first generate a matrix of firm-level trade flows that are non-zero when sales exceed the bilateral entry cost, i.e. $x_{ni}(a) > \sigma w f_{ni}$. In a first stage we infer from this fake trade dataset the micro-based estimates of the aggregate trade elasticities by applying the methodology of Section 5.2:

We first retrieve min-to-mean ratios for all country-pairs and then compute the corresponding set of theoretical dyadic elasticities (equations 16 and 17). In a second stage, we turn to the macro-based estimates of the trade elasticity. To this purpose we collapse firm-level trade flows at the country-pair level to construct a matrix of bilateral aggregate trade. We then run gravity regressions (both using country fixed effects and tetrads) and retrieve the point estimate of applied tariffs. Hence, for each draw, we obtain one macro-based estimate of the trade elasticity that we compare to the cross-dyadic average of the micro-based elasticities. This procedure is replicated 50 times. Notice that it is computationally demanding as we have to manipulate very large trade matrices (1 million firms $\times$ 10 origin countries $\times$ 10 destination countries) .

Table A.13: Monte Carlo results: elasticities wrt to a change in trade costs

Distribution:	Log-Normal				Pareto			
# firms per country:	1K	10K	100K	1M	1K	10K	100K	1M
total exports (micro)	-4.69	-4.57	-4.56	-4.55	-5.74	-5.36	-5.23	-5.18
	(0.60)	(0.38)	(0.34)	(0.33)	(0.87)	(0.37)	(0.21)	(0.13)
total exports (macro/tetrads)	-4.80	-4.60	-4.59	-4.58	-5.55	-5.37	-5.21	-5.18
	(0.66)	(0.29)	(0.09)	(0.03)	(0.81)	(0.56)	(0.35)	(0.22)
total exports (macro/FE)	-4.65	-4.57	-4.55	-4.55	-5.59	-5.31	-5.22	-5.17
	(0.20)	(0.09)	(0.03)	(0.01)	(0.29)	(0.16)	(0.15)	(0.08)
nb exporters (macro/FE)	-3.20	-3.16	-3.16	-3.16	-5.19	-5.15	-5.14	-5.13
	(0.09)	(0.03)	(0.01)	(0.00)	(0.13)	(0.05)	(0.02)	(0.01)
avg. exports (macro/FE)	-1.45	-1.41	-1.40	-1.39	-0.40	-0.15	-0.08	-0.04
	(0.17)	(0.08)	(0.03)	(0.01)	(0.26)	(0.16)	(0.15)	(0.08)

Notes: 50 replications for each cell, parameters on fixed costs of exports and size of the demand term have been calibrated so the share of exporters averages to 10-11% in all simulations. For each elasticity, the first line reports the average value. Standard deviations are in parentheses. For the micro elasticity, the number in parentheses is the average of standard deviations of the elasticity in each draw (quantifying the degree of heterogeneity in bilateral elasticities). For the macro elasticities, we report the standard deviation of elasticities across the 50 replications.

The simulation results are displayed in Table A.13 for log-normal (col.1-col.4) and Pareto (col.5-col.8) and for different degrees of firm scarceness (from 1000 to 1 million firms per country). Each column reports averages and standard errors across replications.

Our baseline simulation under Pareto (col. 8) shows that the simulated economy with 1 million firms conforms to the theoretical prediction of a model with a continuum of firms. The micro-based estimates of the aggregate trade elasticity are relatively homogeneous across dyads (the second row reports the mean value of the standard deviation within each draw) and their average (first row) is close to the macro-based estimates of the elasticities retrieved from tetrad-like specification (third row) or standard gravity (fourth row). Finally the elasticity of the average export (last row) is not significantly different from zero, as expected from the theoretical prediction associated with Pareto heterogeneity and a continuum of firms. We conclude from this exercise that scarceness does not seem to play a central role in our fake sample of 1 million firms with 10 percent of exporters.

From the baseline simulation under log-normal (Column 4) we see that the macro-based estimate of the aggregate elasticity and the cross-dyadic average of the micro-based estimates are quantitatively very close - i.e. equality cannot be rejected. This constitutes the main result of our Monte Carlo approach. It confirms that the similarity between micro- and macro-based estimates in section 5.3 can be safely interpreted as supportive of the log-normal distribution. Notice that the magnitude of the simulation results on the three macro-based elasticities (total exports, count of exporters and average exports) is also close to what we obtain with the sample of French and Chinese firms. This is remarkable given that our Monte Carlo approach is minimal and shares only few features with the true data, i.e. the parameters of firm-level heterogeneity and the share of exporters.