The time-varying trade elasticity and the time-invariant welfare gains from trade *

Elizaveta Archanskaia†

March 2016

Abstract

Canonical trade theory predicts a mechanical link from increased trade integration to increased consumer purchasing power. We challenge the empirical validity of this claim. We propose a novel methodology that does not rely on trade cost data, and yet, enables us to obtain annual estimates of the elasticity of imports to variable trade costs. This parameter is dubbed the 'trade elasticity' in the recent literature, and it captures the strength of the incentive to trade. We find that the structural trade elasticity increases by 35% over 1995-2011. We use our estimates of the trade elasticity to track the purchasing power channel of the gains from trade. We find that access to foreign supply increases real income of the representative consumer by 15-30% at the interquartile range of our sample. This contribution is unchanged over 1995-2011 notwithstanding a 7-9 percentage point increase in reliance on foreign supply. We conclude that the increasing sensitivity of expenditure to trade costs wiped out the effect of reductions in trade barriers on consumer purchasing power.

Keywords: Trade elasticity, Armington elasticity, Welfare gains

JEL codes: F11, F14, F15

*This paper has greatly benefited from helpful discussions with Johannes van Biesebroeck, Lorenzo Caliendo, Thomas Chaney, Guillaume Daudin, Lionel Fontagné, Samuel Kortum, Jacques Le Cacheux, Thierry Mayer, Florian Mayneris, Yasusada Murata, Antonella Nocco, Mathieu Parenti, André Romahn, and Frank Verboven.

†KU Leuven. Email: liza.archanskaia@kuleuven.be. Address: Naamsestraat 69, 3000 Leuven, Belgium.

1 Introduction

The increased interdependence of countries within the world trade system has been documented extensively. This ongoing process of trade deepening raises an important question from the perspective of the consumer, namely whether increased reliance on foreign supply maps into higher real income. Yet, relatively few studies provide global quantifications of welfare gains associated to the recent spell of trade deepening.[1] One possible explanation is that such an exercise poses stringent data requirements. A further limitation is that the answer may be contingent on the parameterization of technology and preferences chosen by the researcher.[2]

Guided by trade theory, we nevertheless expect increased reliance on foreign supply to translate into increased real income for the representative consumer. Indeed, in canonical trade models, welfare gains are codetermined by the strength of the impediments and of the incentives to trade. The strength of the incentives to trade is reconductible to structural parameters that characterize the supply and demand sides of the economy (Costinot and Rodriguez-Clare (2014)). These parameters are generally assumed to be time-invariant. Holding the incentives to trade fixed, increased trade integration stems from a reduction in trade barriers that facilitates consumer access to foreign goods. We therefore expect a one-to-one mapping from trade deepening to consumer purchasing power (Arkolakis et al. (2012)).

The issue that is not addressed by canonical trade theory but may turn out to be empirically relevant is that the strength of the incentives to trade may change as impediments to trade are gradually removed. In particular, there is no empirical evidence to back the commonplace assumption that incentives to trade remained unchanged in the recent period in which the global economy underwent a process of structural transformation (Head and Mayer (2013)). This question is ultimately empirical in nature, and it constitutes the central focus of this paper.

Our main contribution consists in developing a theoretically grounded approach that does not rely on trade cost data, and yet, enables us to obtain annual estimates of the elasticity of imports to variable trade costs. This parameter is dubbed the 'trade elasticity' in the seminal paper of Arkolakis et al. (2012), and it captures the strength of the incentive to trade. Our estimates of the trade elasticity enable us to pin down the increase in real income of the representative consumer due to improved access to a fixed set of goods, i.e. the purchasing power channel of the gains from trade, and to track its evolution over 1995-2011.

[1] Caliendo et al. (2015) is one of the few papers that provide a global quantification of gains from tariff reductions.

[2] Results' sensitivity to functional form assumptions is discussed in Costinot and Rodriguez-Clare (2014), Melitz and Redding (2015), Head et al. (2014), Simonovska and Waugh (2014b), and Behrens et al. (2014).

Our main result is that the one-to-one mapping from trade deepening to consumer purchasing power is not borne out in the data because the incentive to trade has changed. The contribution of international trade to consumer purchasing power is stable at 15-30% of real income at the interquartile range of our sample, notwithstanding a 7-9 percentage point increase in reliance on foreign supply over 1995-2011. This stability is due to the 35% increase in the magnitude of the structural trade elasticity. Our results would be qualitatively different if we constrained the trade elasticity to be time invariant. We would then conclude that real income increased by 4-6 percentage points over 1995-2011.

To discipline our approach to estimation, we rely on a highly stylized model of the world economy that nests technological heterogeneity à la Eaton and Kortum (2002) within the canonical Anderson and van Wincoop (2003) model of trade. In our generalized Armington model, the set of goods is fixed across trade equilibria, and the gains from trade are determined by the increase in real income due to improved access to this set of goods. In this world, we can directly apply the sufficient statistic approach of Arkolakis et al. (2012) to quantify welfare gains from trade. Indeed, the structural trade elasticity and the extent of reliance on domestic supply suffice to determine their magnitude.

Arguably, our focus on the purchasing power channel of the gains from trade is restrictive. Indeed, gains from cheaper access to a fixed set of goods may be magnified (or dampened) by changes in product variety and product innovation (Melitz and Redding (2015), Sampson (2015), Perla et al. (2015)). The advantage of our approach is that this channel is particularly well-suited for quantification. Indeed, the magnitude of the gains from trade in this world is consistently defined over time, and it does not hinge on the intrinsic valuation of variety which may be impossible to recover (Behrens et al. (2014)). Thus, we do not claim that we pin down the magnitude of total gains from trade. Rather, we demonstrate that welfare analysis of the gains reaped from trade integration is highly sensitive to the assumption that the structural parameters of the model are time-invariant.

We work with a generalized Armington model with two-tier CES preferences. But we deviate from the canonical approach in which consumer sensitivity to cost differences is sector-specific (Feenstra et al. (2014), Imbs and Méjean (2015), Ossa (2015)). Instead, we allow the elasticity of demand to sector-specific shocks to differ from the elasticity of demand to economy-wide shocks. This wedge is assumed common to all sectors. The advantage of our set-up is its parsimony whereby two elasticities suffice to characterize consumer choice. Further, our parameterization delivers the gravity structure of Anderson and van Wincoop (2003) at the level of aggregate trade notwithstanding the fact that we incorporate cost heterogeneity in production. The upper tier demand elasticity captures the valuation of the 'made-in' effect common to all exported goods, and it directly determines the trade elasticity.

We may be worried that our deviation from the canonical two-tier demand set-up is not innocuous.[3] We put forward that the existence of a wedge in the sensitivity of demand to micro- and macro-level shocks is economically meaningful. Moreover, our parameterization is readily interpreted in the canonical set-up. Basically, we posit that demand for products of different origin is codetermined by product characteristics with the lower tier demand elasticity and by country characteristics with the upper tier demand elasticity.

On the supply side, we posit that only the best available technology is used in production. Technology is country and sector specific. This feature of the model rationalizes price variation within the set of goods delivered by each exporter to the world market. It also helps to reconcile truncation in disaggregate trade data with the prediction of our model that the set of goods is fixed across trade equilibria. We follow Baldwin and Harrigan (2011) who show that zeros may arise as a consequence of statistical thresholds for reporting trade. In our two-tier demand set-up, products associated to high cost draws carry marginal weight in consumption. By allowing for statistical thresholds, we get the prediction that trade flows at the product level are absent from trade statistics whenever their value falls below the threshold for reporting trade.[4]

We follow a stepwise approach to retrieve annual estimates of the structural trade elasticity. We use expenditure variation within product sets that the exporter delivers to each market to retrieve an estimate of the lower tier elasticity.[5] We use this estimate of the lower tier elasticity to consistently aggregate product prices. We then use expenditure variation at the level of aggregate trade together with the constructed price indices to retrieve an estimate of the upper tier elasticity.[6] Finally, we use the conditional distribution of exporter-specific sales to obtain an estimate of the wedge that truncation introduces between the structural and the measured trade elasticity (Head et al. (2014)). Combined with the estimate of the lower tier elasticity, this wedge determines the dispersion of technology draws. Combined with the estimate of the

[3] Different parameterizations of the demand system ultimately map into different error structures, i.e. different sets of assumptions on consumer heterogeneity that, by definition, is unobserved (Cardell (1997)).

[4] Archanskaia and Daudin (2014) also posit that statistical thresholds explain zeros in disaggregate trade data but they opt for a different parameterization of consumer choice and do not microfound cost heterogeneity.

[5] Labor market clearing in the exporting country generates a one-to-one mapping from the ranking of technology draws to the ranking of prices. We construct a proxy of technological ability and instrument product prices.

[6] Price indices are instrumented, and a proxy of unobserved quality is included in the estimation.

upper tier elasticity, this wedge determines the structural upper tier elasticity. The latter directly determines the structural trade elasticity in our model.

On the demand side, we find that the lower tier elasticity is in the 7-9 range while the upper tier elasticity is in the 2.5-3.5 range. These results confirm our prior that demand sensitivity to sector-specific shocks is significantly higher than demand sensitivity to economy-wide shocks. These estimates are directly comparable to recent studies that report wedges of similar magnitude between micro- and macro-level trade elasticities (Imbs and Méjean (2015)). We provide a new rationale for the existence of such wedges by showing that they may be capturing differences in demand sensitivity to different types of shocks. On the supply side, our estimates of the degree of dispersion in technology draws conform to magnitudes reported in Eaton and Kortum (2002), Costinot et al. (2012), and Caliendo and Parro (2015). We find that the shape parameter of the Fréchet distribution is in the 6-8 range.

Regarding the strength of the incentive to trade, we find that truncation increases the perceived elasticity of trade flows to trade costs. Specifically, the wedge between the measured and the structural trade elasticity is reduced from 15% in 1995 to 8% in 2011. Thus, the measured trade elasticity increases from 1.95 to 2.44 over 1995-2011 (+25%) while the structural trade elasticity increases from 1.7 to 2.3 (+35%). These magnitudes are significantly lower than the 3.7-5.5 range reported in Simonovska and Waugh (2014b) for the benchmark Armington and Ricardian models. We explain this discrepancy by the fact that our model accounts for multiple dimensions of heterogeneity in generating the data. Indeed, Simonovska and Waugh (2014b) show that more flexible models map into lower magnitudes of the trade elasticity.

Our paper belongs to the rapidly growing line of work on quantification of welfare gains in structural trade models. We contribute to the empirical literature on the estimation of trade elasticities by showing that there exists at least one parameterization of technology and preferences that allows obtaining annual estimates of the structural trade elasticity without relying on trade cost data.[7] We contribute to the ongoing debate on sufficient statistics that allow quantifying the gains from trade by showing that welfare analysis is not only sensitive to functional form assumptions but also to the assumption that structural parameters are time-invariant.[8]

The closest paper to our study is Archanskaia and Daudin (2014) who document an increase in perceived product substitutability over 1963-2009 and argue that it explains the non-

[7] Recent contributions to the estimation of trade elasticities are Caliendo and Parro (2015), Simonovska and Waugh (2014a,b), and Head et al. (2010) who build on Head and Ries (2001) and Eaton and Kortum (2002).

[8] Recent contributions are Melitz and Redding (2015), Head et al. (2014), Behrens et al. (2014), Bas et al. (2015).

decreasing distance elasticity of trade. Our paper is complementary to theirs in that we develop a theoretically grounded approach to the estimation of structural supply and demand parameters in a generalized Armington model and formally establish the linkage between the estimated parameters and the structural parameter that determines the strength of the incentive to trade.

The rest of the paper is organized as follows. In section 2 we go over the benchmark model, incorporate truncation, and present the stepwise approach to pin down the structural trade elasticity. In section 3 we report our estimates of the wedge between the measured and the structural trade elasticities. In section 4 we report our results on the lower tier elasticity and characterize the supply side of the economy. In section 5 we discuss the estimation strategy that delivers annual estimates of the upper tier demand elasticity and report our results on the structural trade elasticity. In section 6 we quantify the purchasing power channel of the gains from trade and evaluate the impact of the increasing trade elasticity. We conclude in section 7.

2 An Armington model with cost heterogeneity in production

2.1 The benchmark model

There is a continuum $k \in [0,1]$ of products ('sectors') in each country. Consumer preferences are represented by a two-tier CES utility function. At the lower tier, all products of the same national origin $Q_i(k)$ are combined into a country-specific composite good Q_i.

$$Q_i = \left[\int_0^1 Q_i(k)^{\frac{\sigma'-1}{\sigma'}} \,\mathrm{d}k \right]^{\frac{\sigma'}{\sigma'-1}} \tag{1}$$

At the upper tier, composite goods of different national origin are combined into an aggregate consumption good. Assuming $1 < \sigma \leq \sigma'$, overall utility is:

$$U = \sum_{i=1}^{N} \left\{ Q_i^{(\sigma-1)/\sigma} \right\}^{\sigma/(\sigma-1)} \tag{2}$$

The parameterization of technology mimicks the canonical Eaton and Kortum (2002) model. Production technology is constant returns to scale, non-proprietory within the country, and independently accumulated for each product k. Labor is the only factor of production. This restrictive assumption is due to the fact that we do not have information on the input-output structure of production at the HS 6-digit level (≈ 5000 distinct products).[9]

[9] The alternative approach followed in Caliendo et al. (2015) is to work with 15 big production sectors and to account for the full structure of input-output linkages in the world trade matrix.

The details on the parameterization of technology are provided in Appendix A. Here we directly derive the gravity structure of aggregate bilateral trade. In equilibrium, workers must be indifferent to being employed in any sector. Labor market clearing implies that the wage w_i is equalized across sectors. Denoting by $z_i(k)$ the best technique available for production in sector k, the factory gate price of product k is $P_i(k) = w_i/z_i(k)$. As in the canonical Ricardian model, prices of effectively produced goods map into inverse labor requirements in i. Ordering prices in increasing order $k = \{1,...,k,...,K\}$, we get:

$$P_i(1) < ... < P_i(k) < ... < P_i(K) \quad \Leftrightarrow \quad z_i(1) > ... > z_i(k) > ... > z_i(K) \tag{3}$$

The assumption of product differentiation by place of origin implies that all products survive and are exported. The ideal price index across the unit continuum of products is:

$$P_i = \left[\int_0^1 P_i(k)^{1-\sigma'} \, \mathrm{d}k \right]^{\frac{1}{1-\sigma'}} \tag{4}$$

As shown in Appendix A, the price index can be equivalently written as a function of the wage adjusted by the scale parameter of the Fréchet distribution $\tilde{z}_i$ that captures the expected technological ability of the exporter: $P_i/\chi = w_i/\tilde{z}_i$ where χ is a country-invariant scalar. The distribution of factory gate prices is invariant to trade costs. Under the assumption of iceberg trade costs t_{ij}, the landed price of product k delivered from i to j is $P_{ij}(k) = w_i\tau_{ij}/z_i(k)$ where $\tau_{ij} = 1 + t_{ij}$. Similarly, the landed price of the composite good delivered from i to j is:

$$P_{ij} = \chi \frac{w_i \tau_{ij}}{\tilde{z}_i} \tag{5}$$

We denote expenditure in j on the composite good delivered from i as $X_{ij} = P_{ij}Q_{ij}$. The value of bilateral trade is obtained by maximizing (2) subject to the constraint that total expenditure $Y_j = \sum_{i \in N} P_{ij}Q_{ij}$ does not exceed j's income.[10] The share spent on goods from i is:

$$\frac{X_{ij}}{Y_j} = \frac{\left(P_{ij}\right)^{1-\sigma}}{\sum_{n=1}^{N} \left(P_{nj}\right)^{1-\sigma}} \tag{6}$$

The gravity structure of aggregate bilateral trade replicates Anderson and van Wincoop (2003) whereby the magnitude of the trade elasticity is determined by the Armington elasticity σ that captures perceived substitutability of country-specific composite goods:

$$X_{ij} = \frac{Y_i Y_j}{Y_w} \left(\frac{\tau_{ij}}{\Pi_i \Phi_j} \right)^{1-\sigma} \tag{7}$$

[10] Total income is given by the landed value of exports from j to all partners: $\sum_{n \in N} P_{jn}Q_{jn}$.

where Y_w is world expenditure, $\Phi_j = \left[\sum_{n=1}^{N} \left(P_{nj}\right)^{1-\sigma}\right]^{1/(1-\sigma)}$ is the overall price index of the importer, $\Pi_i = \left[\sum_j s_j(\tau_{ij}/\Phi_j)^{1-\sigma}\right]^{1/(1-\sigma)}$ is the multilateral trade resistance term of the exporter, and $s_j = Y_j/Y_w$ is the expenditure share of each country.[11]

To sum up, our set-up rationalizes micro-level price heterogeneity among products of the same origin. Further, it delivers the prediction that the ranking of prices for exported products has a one-to-one mapping to the reverse ranking of technology in the exporting country. At the macroeconomic level, our set-up mimicks the canonical Armington model by predicting that each country produces a unique composite good Q_i, and that the supply of this composite good is perfectly inelastic. As the set of goods is fixed across trade equilibria, the structural trade elasticity is directly determined by the upper tier demand elasticity.

2.2 Truncated gravity

Only the intensive margin of trade is operational in our model. However, as we show in Appendix B, only a subset of products is reported as traded in the BACI dataset (Gaulier and Zignago (2010)) that we use in this paper. In our two-tier demand set-up, products characterized by high unit labor requirements carry marginal weight in consumption. Exports of such products generate small trade flows that may be omitted from trade statistics because their value falls below the thresholds for reporting trade. Baldwin and Harrigan (2011) discuss the importance of statistical thresholds in generating truncation in highly disaggregate US data. We know that the flow is not reported in the UN COMTRADE database from which the BACI dataset is built if its value falls below 1000 US\$. Consequently, we incorporate truncation in the model by allowing for statistical thresholds in reporting trade.

We posit the existence of a statistical threshold $\bar{X}$ common to all countries such that the nominal value of trade at the product-level is reported iff it is at least equal to this threshold. We characterize effective expenditure allocation among truncated composite goods by conditioning utility $\bar{U}_j$ to be derived from registered quantities $\bar{Q}_{ij}$ according to the truncated analog of (2). Total expenditure $\bar{Y}_j$ is set equal the sum of registered bilateral imports: $\bar{Y}_j = \sum_i \bar{X}_{ij}$ where $\bar{X}_{ij} = \sum_k X_{ij}(k) \left\{X_{ij}(k) : X_{ij}(k) \geq \bar{X}\right\}$. Basically, the solution to the non-truncated problem directly gives expenditure allocation in the truncated problem by conditioning on some

[11] See Anderson and van Wincoop (2003): use (6) and sum over i's partners to get income in i: $Y_i = \sum_j X_{ij} = \sum_j (\chi w_i \tau_{ij}/\tilde{z}_i)^{1-\sigma} \Phi_j^{\sigma-1} Y_j$. Solve for $(\chi w_i/\tilde{z}_i)^{1-\sigma} = Y_i \left[\sum_j (\tau_{ij}/\Phi_j)^{1-\sigma} Y_j\right]^{-1}$, plug this back into (6) to get $X_{ij} = Y_i Y_j \left(\frac{\tau_{ij}}{\Phi_j}\right)^{1-\sigma} \left[\sum_j (\tau_{ij}/\Phi_j)^{1-\sigma} Y_j\right]^{-1}$. Multiply and divide the RHS by Y_w and replace Π_i by its value.

threshold $\bar{X}$. A vector of trade deficits D_j equalizes truncated expenditure to truncated income: $\bar{Y}_j = \sum_n \bar{P}_{jn}\bar{Q}_{jn} + D_j$. Denoting by $\bar{\sigma}$ the parameter that measures the substitutability of truncated composite goods, we obtain the truncated analog of (6):

$$\frac{\bar{X}_{ij}}{\bar{Y}_j} \quad = \quad \frac{\left(\bar{P}_{ij}\right)^{1-\bar{\sigma}}}{\sum_{n=1}^{N}\left(\bar{P}_{nj}\right)^{1-\bar{\sigma}}} \tag{8}$$

The gravity equation for truncated trade is derived using the same procedure as for non-truncated trade. We denote $\bar{\Phi}_j = \left\{\sum_{n=1}^{N}\left(\bar{P}_{nj}\right)^{1-\bar{\sigma}}\right\}^{1/(1-\bar{\sigma})}$ the truncated price index in j. We replace the truncated bilateral price index $\bar{P}_{ij}$ in (8) by its value using (62), and we sum (8) across all i's partners to get truncated income in i that we denote $\bar{I}_i$:[12]

$$\bar{I}_i = \sum_j \bar{X}_{ij} = \sum_j \left\{ \bar{Y}_j \bar{\Phi}_j^{\bar{\sigma}-1} \left[\psi \left[Y_j \Phi_j^{\sigma-1} \right]^{\frac{\theta-\sigma'+1}{\sigma'-1}} \left[\frac{w_i \tau_{ij}}{\tilde{z}_i} \right]^{-(\theta-\alpha)} \right]^{\frac{1-\bar{\sigma}}{1-\sigma'}} \right\} \tag{9}$$

Next, we express the truncated price index as a function of the non-truncated price index:

$$\bar{\Phi}_j = \left\{ \sum_{n=1}^{N} \left(\bar{P}_{nj}\right)^{1-\bar{\sigma}} \right\}^{1/(1-\bar{\sigma})} = \psi^{\frac{1}{1-\sigma'}} \left[Y_j \Phi_j^{\sigma-1} \right]^{\frac{-(\theta-\sigma'+1)}{(\sigma'-1)^2}} \left[\sum_n \left(\frac{w_n \tau_{nj}}{\tilde{z}_n} \right)^{\frac{(\theta-\alpha)(1-\bar{\sigma})}{\sigma'-1}} \right]^{\frac{1}{1-\bar{\sigma}}} \tag{10}$$

We denote by $\bar{\bar{\Phi}}_j = \left\{ \sum_n \left(\frac{w_n \tau_{nj}}{\tilde{z}_n} \right)^{\frac{(\theta-\alpha)(1-\bar{\sigma})}{\sigma'-1}} \right\}^{1/(1-\bar{\sigma})}$ the truncated price index that is independent of market-specific characteristics whereby the last term on the RHS in (10) is simply $\bar{\bar{\Phi}}_j$. Replacing the truncated price index in (9) by its value in (10) and simplifying gives:

$$\bar{I}_i = \sum_j \bar{Y}_j \bar{\bar{\Phi}}_j^{\bar{\sigma}-1} \left[\frac{w_i \tau_{ij}}{\tilde{z}_i} \right]^{\frac{(\theta-\alpha)(1-\bar{\sigma})}{\sigma'-1}} \Leftrightarrow \left(\frac{w_i}{\tilde{z}_i} \right)^{\frac{(\theta-\alpha)(1-\bar{\sigma})}{\sigma'-1}} = \bar{I}_i \left[\sum_j \bar{Y}_j \left[\frac{\tau_{ij}^{\frac{\theta-\alpha}{\sigma'-1}}}{\bar{\bar{\Phi}}_j} \right]^{1-\bar{\sigma}} \right]^{-1} \tag{11}$$

Next, we replace the truncated bilateral price index $\bar{P}_{ij}$ by its value in (62) and rewrite truncated bilateral expenditure (8) as:

$$\bar{X}_{ij} \quad = \quad \bar{Y}_j \bar{\Phi}_j^{\bar{\sigma}-1} \left[\psi^{\frac{1}{1-\sigma'}} \left[Y_j \Phi_j^{\sigma-1} \right]^{\frac{-(\theta-\sigma'+1)}{(\sigma'-1)^2}} \right]^{1-\bar{\sigma}} \left[\frac{w_i \tau_{ij}}{\tilde{z}_i} \right]^{\frac{(\theta-\alpha)(1-\bar{\sigma})}{\sigma'-1}} \tag{12}$$

We simplify (12) by replacing $\bar{\Phi}_j$ by its value in (10). Further, we multiply and divide (11) by total truncated expenditure $\bar{Y}_w = \sum_j \bar{Y}_j$, and plug (11) in (12) to get:

$$\bar{X}_{ij} \quad = \quad \frac{\bar{Y}_j \bar{I}_i}{\bar{Y}_w} \left[\frac{\tau_{ij}^{\frac{\theta-\alpha}{\sigma'-1}}}{\bar{\bar{\Phi}}_j} \right]^{1-\bar{\sigma}} \left[\sum_j \frac{\bar{Y}_j}{\bar{Y}_w} \left(\frac{\tau_{ij}^{\frac{\theta-\alpha}{\sigma'-1}}}{\bar{\bar{\Phi}}_j} \right)^{1-\bar{\sigma}} \right]^{-1} \tag{13}$$

[12] See Appendix A: $\psi = \frac{\theta}{\theta-\sigma'+1} \left[\chi^{\sigma'-\sigma} \bar{X}^{-1} \right]^{\frac{\theta-\sigma'+1}{\sigma'-1}}$ and $\alpha = (\sigma'-\sigma)(\theta-\sigma'+1)/(\sigma'-1)$.

The last term on the RHS of (13) is a monotonic transformation of the truncated multilateral resistance term of the exporter $\bar{\bar{\Pi}}_i^{1-\bar{\sigma}} = \left[\sum_j \bar{s}_j (\tau_{ij}^{(\theta-\alpha)/(\sigma'-1)}/\bar{\bar{\Phi}}_j)^{1-\bar{\sigma}}\right]$ where $\bar{s}_j = \bar{Y}_j/\bar{Y}_w$ is the truncated expenditure share. Define $\eta = (\theta-\alpha)/(\sigma'-1)$. The truncated gravity equation is:

$$\bar{X}_{ij} = \frac{\bar{Y}_j \bar{I}_i}{\bar{Y}_w}\left[\frac{\tau_{ij}^{\eta}}{\bar{\bar{\Pi}}_i \bar{\bar{\Phi}}_j}\right]^{-(\bar{\sigma}-1)} \tag{14}$$

We denote the structural trade elasticity $\varepsilon = -(\sigma-1)$ and the measured trade elasticity $\bar{\varepsilon} = -\eta(\bar{\sigma}-1)$. By direct comparison of (14) to (7), we see that $\bar{\varepsilon}$ may deviate from ε through two channels. The first channel works at the micro-level and is captured through $\eta \in \left\{1, \frac{\theta}{(\sigma-1)}\right\}$.[13] As shown in Appendix A, this parameter captures how truncation increases the sensitivity of exporter-specific price indices to trade frictions. The second channel works at the macro-level and is captured through $\bar{\sigma} \geq \sigma$ (see below). This parameter captures how truncation increases the sensitivity of aggregate expenditure to differences in the prices of composite goods.

2.3 The wedge between the measured and the structural trade elasticity

We now characterize the wedge between the measured and the structural trade elasticity. Consider relative truncated expenditure for any pair of exporters i and i'. Expenditure on any product verifies $X_{ij}(k) = (P_{ij}(k)/P_{ij})^{1-\sigma'} X_{ij}$. Total observed expenditure is obtained by summing across reported product flows. In relative terms, we get:

$$\frac{\bar{X}_{ij}}{\bar{X}_{i'j}} = \frac{X_{ij} P_{ij}{}^{\sigma'-1} \int_0^{\bar{\upsilon}_{ij}} p^{1-\sigma'} f_i(p)\mathrm{d}p}{X_{i'j} P_{i'j}{}^{\sigma'-1} \int_0^{\bar{\upsilon}_{i'j}} p^{1-\sigma'} f_{i'}(p)\mathrm{d}p} \tag{15}$$

Consider the numerator on the right hand side of (15). The last component is equal to $\bar{P}_{ij}^{1-\sigma'}$ and is a monotonic transformation of the truncated price index. The second component is $P_{ij}^{\sigma'-1} = P_{ij}^{\sigma-1} P_{ij}^{\sigma'-\sigma}$. Since $X_{ij}/X_{i'j} = (P_{ij}/P_{i'j})^{1-\sigma}$, (15) simplifies to:

$$\frac{\bar{X}_{ij}}{\bar{X}_{i'j}} = \left[\frac{\bar{P}_{ij}}{\bar{P}_{i'j}}\right]^{1-\sigma'} \left[\frac{P_{ij}}{P_{i'j}}\right]^{\sigma'-\sigma} \tag{16}$$

We replace the truncated and the non-truncated price indices by their respective values in (62) and (5) and rearrange to get:

$$\frac{\bar{X}_{ij}}{\bar{X}_{i'j}} = \left[\frac{w_i \tau_{ij}/\tilde{z}_i}{w_{i'} \tau_{i'j}/\tilde{z}_{i'}}\right]^{-\theta \frac{(\sigma-1)}{(\sigma'-1)}} \tag{17}$$

[13] We establish that η monotonically increases from 1 to $\theta/(\sigma-1)$ as σ' decreases from $(\theta+1)$ to σ by rewriting $\eta = [\theta(\sigma-1)+(\sigma'-\sigma)(\sigma'-1)](\sigma'-1)^{-2}$ and evaluating this expression for $\sigma' \in [\sigma, (\theta+1)[$.

It is immediate from (17) that $\bar{\varepsilon}/\varepsilon = \theta/(\sigma'-1)$. This ratio is a sufficient statistic that captures the role of truncation in determining total trade. The intuition is that θ regulates how the number of observed products changes as a consequence of a change in trade costs while σ' determines the change in the price index associated with the inclusion of these marginal products (Chaney (2008)). Whenever σ' is sufficiently high, products associated to low technology draws have negligible weight in consumption. Consequently, truncated trade flows and price indices are nearly identical to their non-truncated analogs, and truncation has negligible impact on the magnitude of the trade elasticity.

An attractive feature of this framework is that it enables us to empirically evaluate how far the allocation of expenditure is from the canonical Armington and Ricardian worlds. Parameter restrictions of the model imply that $\sigma' \in [\sigma, (\theta+1)[$ whereby, for given θ and σ, $\bar{\varepsilon}/\varepsilon \in]1, \theta/(\sigma-1)]$. When σ' approaches its lower bound, our model mimicks the Ricardian world in which the lower and upper tier demand elasticities coincide, and the measured trade elasticity is determined by the degree of technology dispersion.[14] When σ' approaches its upper bound, our model mimicks the Armington world in which the impact of the extensive margin on aggregate trade is close to nil.

We pin down the wedge between the measured and the structural upper tier demand elasticity $\bar{\sigma}/\sigma$ with help of the second expression of relative truncated expenditure. The latter is obtained by taking the ratio of (8) for i and i' where we replace bilateral truncated price indices by their value in (62) to get:

$$\frac{\bar{X}_{ij}}{\bar{X}_{i'j}} = \left[\frac{w_i \tau_{ij}/\tilde{z}_i}{w_{i'}\tau_{i'j}/\tilde{z}_{i'}}\right]^{-\eta(\bar{\sigma}-1)} \tag{18}$$

From (17) and (18) we have $\theta(\sigma-1)(\sigma'-1) = \eta(\bar{\sigma}-1)$. Replacing η by its value, simplifying and rearranging gives:

$$\frac{\bar{\sigma}-1}{\sigma-1} = \frac{\theta(\sigma'-1)}{\theta(\sigma-1)+(\sigma'-\sigma)(\sigma'-1)} \tag{19}$$

We evaluate the ratio $(\bar{\sigma}-1)/(\sigma-1)$ for $\sigma' \in [\sigma, (\theta+1)[$. Given θ and σ, the ratio tends to 1 when σ' tends to its lower and upper bounds. The derivative of the ratio wrt σ' is:

$$\frac{\partial(\cdot)}{\partial\sigma'} = \frac{\theta(\sigma-1)+(\sigma'-1)(\sigma'-\sigma)-(\sigma'-1)(2\sigma'-(\sigma+1))}{\theta\left[(\sigma-1)+\theta^{-1}(\sigma-1)(\sigma'-\sigma)\right]^2} \tag{20}$$

The sign of the derivative in (20) is given by the sign of the numerator. We simplify the

[14] This situation corresponds to the maximal wedge between the structural and the measured trade elasticity.

expression in the numerator and rearrange to solve for σ' when the derivative is 0:

$$\frac{\partial(\cdot)}{\partial\sigma'} = 0 \ \Leftrightarrow \ \sigma' = 1 + \sqrt{\theta(\sigma - 1)} \tag{21}$$

Plugging (21) into (20) to solve for the maximum, we find that the ratio $(\bar{\sigma}-1)/(\sigma-1)$ increases from 1 to $\theta\left[2\sqrt{\theta(\sigma-1)}-(\sigma-1)\right]^{-1}$ as σ' increases from σ to $\left[1+\sqrt{\theta(\sigma-1)}\right]$ and thereafter decreases back to 1 as σ' increases from $\left[1+\sqrt{\theta(\sigma-1)}\right]$ to $(\theta+1)$.

2.4 A feasible approach to parameter estimation

When just a subset of products is reported as traded, the ratio $\theta/(\sigma'-1)$ determines the wedge $\bar{\varepsilon}/\varepsilon$. The magnitude of this ratio also determines the approach that needs to be followed to consistently aggregate product prices within exporter-specific product sets (A.3). In section 3.2 we pin down this ratio with help of the conditional distribution of exporter-specific sales. Here, we show how its magnitude affects the estimation strategy that delivers consistent estimates of the structural upper tier elasticity σ.

The special case is $\theta/(\sigma'-1) \approx 1$. Truncation has no impact on aggregate trade, and the structural upper tier elasticity σ coincides with the measured elasticity $\bar{\sigma}$. Observed expenditure shares can be used as weights to consistently aggregate product prices within exporter-specific product sets. This approximation preserves information on relative prices at the upper tier because constructed price indices differ from ideal price indices by a scalar that is invariant across exporters (A.3). In this special case, we retrieve a consistent estimate of the structural elasticity σ by estimating a standard demand equation at the level of aggregate trade without prior knowledge of the other parameters (σ', θ).

The general case is $\theta/(\sigma'-1) \gg 1$. An estimate of the lower tier elasticity σ' is required for consistent price aggregation within exporter-specific product sets (A.3). The structural elasticity σ is now defined by (19) for given σ', θ, and $\bar{\sigma}$. We obtain estimates of these three parameters by implementing a stepwise approach. We use the distribution of expenditure within exporter-specific product sets to retrieve an estimate of σ'. We use σ' to consistently aggregate product prices. We then estimate a standard demand equation at the level of aggregate trade to obtain an estimate of the measured upper tier elasticity $\bar{\sigma}$. Combined with the estimate of the wedge $\theta/(\sigma'-1)$, the estimate of $\bar{\sigma}$ enables us to solve for the structural elasticity σ.

3 Mapping the model to the data

3.1 Stability of exporter-specific price rankings across markets

A distinctive feature of our model is the prediction that the ranking of relative prices within the exporter-specific product set is stable across export markets. This prediction follows from the no arbitrage condition whereby labor market clearing in the exporting country implies that the reverse ranking of factory-gate prices maps into the ranking of unit labor requirements (3). Combined with our assumption that sectoral components of bilateral trade costs are white noise, the no arbitrage condition means that the ranking of landed prices relatively to some benchmark sector is common to all markets on which the exporter is active:

$$\frac{P_{ij}(2)}{P_{ij}(1)} < \ldots < \frac{P_{ij}(k)}{P_{ij}(1)} < \ldots < \frac{P_{ij}(K)}{P_{ij}(1)} \quad \Leftrightarrow \quad \left[\frac{z_i(2)}{z_i(1)}\right] > \ldots > \left[\frac{z_i(k)}{z_i(1)}\right] > \ldots > \left[\frac{z_i(K)}{z_i(1)}\right]$$

We evaluate the empirical relevance of this prediction for the set of exporters that deliver at least 250 products to the world market and at least 50 products to 10 or more markets in every year. This sample includes 92 exporters, and it covers 95-98% of world trade (Appendix B).

For each exporter we restrict the sample to 5 markets on which her variety coverage is largest and conduct the analysis on the set of products common to these 5 markets. We compute Spearman rank correlation coefficients ($\rho_{ijj't}$) for the ranking of relative prices for every exporter on every pair of markets $j \neq j'$. We obtain 10 observations per exporter in each year and 14,230 observations in total. Fig.1 pools results for all years while splitting the sample in half according to the number of products in the common set. Rankings are stable: the IQR is .59-.77 (.68-.81) for exporters with less (more) than 717 products.

In the model, product prices are expressed in normalized *per hour* terms while in the data product prices are measured in *per hour* terms multiplied by the number of hours needed to obtain one unit of the good. We may be worried that Fig.1 is picking up sectoral rather than exporter-sector characteristics, e.g. independently of its origin, a pen tends to cost less than a car. A related concern is that our data reports the unit value per kg instead of the price per item. Unit value rankings pick up differences in product bulkiness that may reflect sectoral rather than exporter-sector characteristics (Hummels and Schaur (2013)).

We therefore provide a sensitivity check for our results. We demean the data in the sectoral dimension and recompute rank correlation coefficients for demeaned price rankings. Fig.2 shows that this correction leads to a leftward shift of the distribution. Yet, rankings' independence is strongly rejected for the bulk of the sample: the IQR is .47-.80 (.45-.75) for exporters

Figure 1: Stability of own price rankings (by number of products in common set)

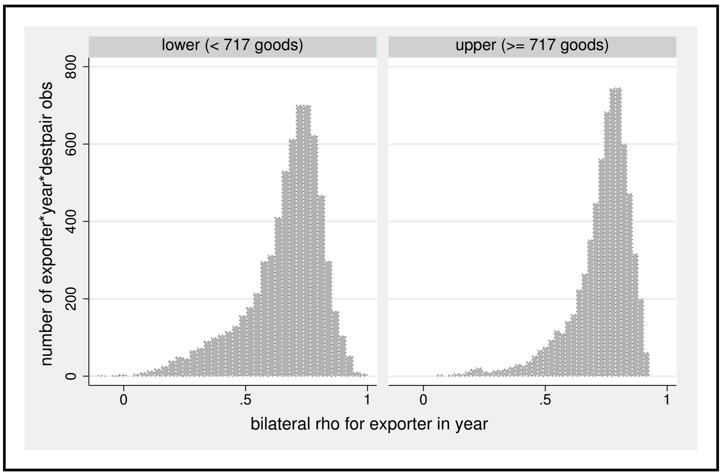

with less (more) than 717 products in the common set. Further, Fig.3 shows that the distribution of Spearman ρ for demeaned relative price rankings is stable over time.

Figure 2: Stability of demeaned price rankings (by number of products in common set)

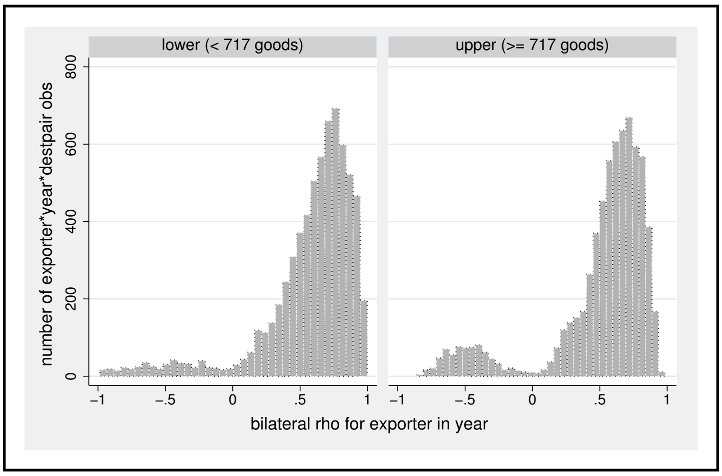

Demeaned price distributions are not perfectly correlated across markets. But the strength of the correlation that we document at the country-level is in line with patterns documented at the firm level by Eaton et al. (2011), Mayer et al. (2014), Arkolakis et al. (2014). The stability of demeaned price distributions supports the main prediction of our model that price variation in the product set carries information on the distribution of technological ability for each exporter.

Figure 3: Stability of demeaned price rankings (by year and number of products)

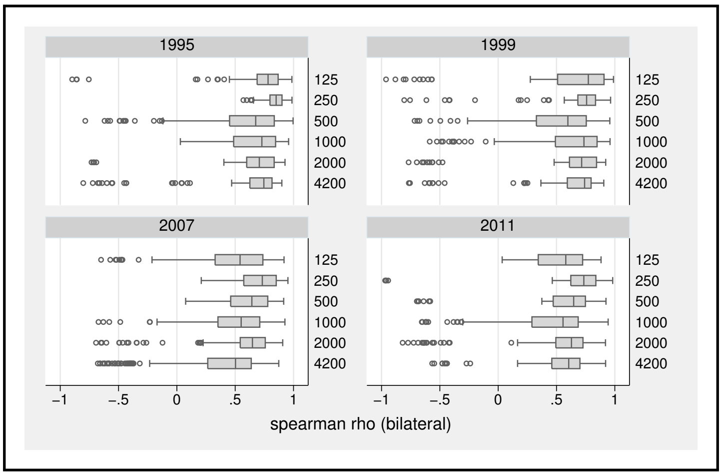

3.2 The tail exponent of the conditional distribution of sales

Our model delivers the prediction that the distribution of sales for any exporter i in any market j conditional on some threshold $\bar{x}$, is Pareto with parameter $\theta/(\sigma'-1)$. We use this prediction to pin down the magnitude of the wedge that statistical truncation introduces between the measured and the structural trade elasticity.[15]

Conditional on the cost threshold $\bar{\upsilon}$, the distribution of costs is $\Pr(\Upsilon \leq \upsilon | \Upsilon \leq \bar{\upsilon}) = (\upsilon/\bar{\upsilon})^{\theta}$ (Appendix A). We use the expression of sectoral demand (25) to define the conditional distribution of sales (X) in terms of the conditional distribution of costs (Υ):

$$\Pr[X \geq x | X \geq \bar{x}] = \Pr\left[\Upsilon \leq x^{\frac{1}{1-\sigma'}} X_{ij}^{\frac{1}{\sigma'-1}} P_{ij} \tau_{ij}^{-1} \;\middle|\; \Upsilon \leq \bar{x}^{\frac{1}{1-\sigma'}} X_{ij}^{\frac{1}{\sigma'-1}} P_{ij} \tau_{ij}^{-1}\right] \quad = \quad [x/\bar{x}]^{-\frac{\theta}{\sigma'-1}}$$

The conditional distribution of sales follows a power law. Denoting by R the number of observations in the tail and by $r = \{1,...,R\}$ the rank of each observation, the relationship between the log rank and the log value of the set of observations in the tail is expected to be approximately linear (Gabaix and Ibragimov (2011)):

$$\ln(r/R) \approx \theta/(\sigma'-1)\ln\bar{x} - \theta/(\sigma'-1)\ln(X_{ij})_{(r)} \tag{22}$$

where $(X_{ij})_{(1)} \geq ... \geq (X_{ij})_{(R)}$ denotes the set of ordered sales for some exporter-market pair ij.

Gabaix and Ibragimov (2011) demonstrate that direct implementation of (22) may deliver biased estimates of the tail exponent in small samples. We therefore implement the shifted

15 The trade flow is not reported in our dataset if it is below 1000 USD. In Appendix B we show that such small trade flows appear to be prevalent in the data. The wedge due to statistical truncation corresponds to $\theta/(\sigma'-1)$.

rank and the harmonic number estimators instead of the naive approach. Gabaix and Ibragimov (2011) argue that these alternative estimators minimize the small sample bias and are well-behaved in the presence of deviations from the power law in the tail.[16]

Our model posits that the tail exponent is common to all exporters while the number of observed products is specific to the pair ij. Hence, we can pool all data on conditional sales' distributions to estimate the tail exponent if we impose a common threshold $\bar{x}$ while allowing R_{ij} to be pair-specific. Equivalently, we can define the tail in a consistent way across the set of pairs ($R_{ij} = R$, $\forall \{i, j\}$) and include a pair dummy to control for the fact that $\bar{x}$ is now pair-specific. In practice, we combine these two criteria to further harmonize the set of distributions included in the estimation. Specifically, we implement the shifted rank and the harmonic number approaches on the set of ordered sales for all pairs $\{ij\}, j \neq i$: $(X_{ij})_{(1)} \geq ... \geq (X_{ij})_{(R)}$ that verify the definition of the tail R for a common threshold $\bar{x}$.

Trade reporting may be error-prone around the statistical threshold of 10^3 US\$. Consequently, we experiment with three alternative definitions of the threshold $\bar{x} = \left\{10^4, 10^5, 10^6\right\}$. For each threshold we drop sectoral flows smaller than $\bar{x}$ and restrict the sample to the set of distributions that cover $\geq 25\%$ of world variety (>1200 goods). Furthermore, we consider four alternative definitions of the tail $\{R = 125, 250, 500, 1000\}$.

The shifted rank estimator implements (22) while shifting r by .5. We estimate the following relationship:

$$\ln(r - 1/2) \quad = \quad a_{0t} + a_{ijt} - a_t \ln(X_{ijt})_{(r)} + (a_{ijt})_{(r)} \tag{23}$$

where $a_t = \theta_t/(\sigma_t - 1)$ is the estimated value of the tail exponent, a_{ijt} is the pair fixed effect, and $(a_{ijt})_{(r)}$ is the error term. The relationship is estimated separately in each year.

The harmonic number approach consists in defining the harmonic number associated with the rank as $H(r) = \sum_{i=1}^{r} \frac{1}{i}$ for $r \geq 1$ and $H(0) = 0$ and estimating the following relationship:

$$H(r-1) \quad = \quad a'_{0t} + a'_{ijt} - a'_t \ln(X_{ijt})_{(r)} + (a'_{ijt})_{(r)} \tag{24}$$

where $a'_t = \theta_t/(\sigma_t - 1)$ is the estimated value of the tail exponent, a'_{ijt} is the pair fixed effect, and $(a'_{ijt})_{(r)}$ is the error term. The relationship is estimated separately in each year.

Fig.4 reports the results for $\bar{x} = 10^5$ in the left pane and for $\bar{x} = 10^6$ in the right pane for

[16] Di Giovanni et al. (2011) use the shifted rank estimator to pin down the tail exponent of the sales distribution for US firms. Gabaix and Ibragimov (2011) use it to characterize the distribution of the population across US cities. Head et al. (2014) use the dual of this estimator to characterize exports of French and Chinese firms.

$R = 500$ because this definition of the tail delivers the best fit of the linear model to the data.[17]. The number of exporters varies between 17 and 25 (8 and 10) in the left (right) pane. The number of observations increases from 103,500 to 190,000 (from 18,000 to 40,500) in the left (right) pane over 1995-2011. The confidence intervals for our set of estimates indicate that the tail exponent $\theta/(\sigma'-1)$ is comprised between 1 and 1.3. The point estimates indicate that the tail exponent is weakly decreasing.[18]

Figure 4: Tail exponent for pairs with $\geq 25\%$ variety

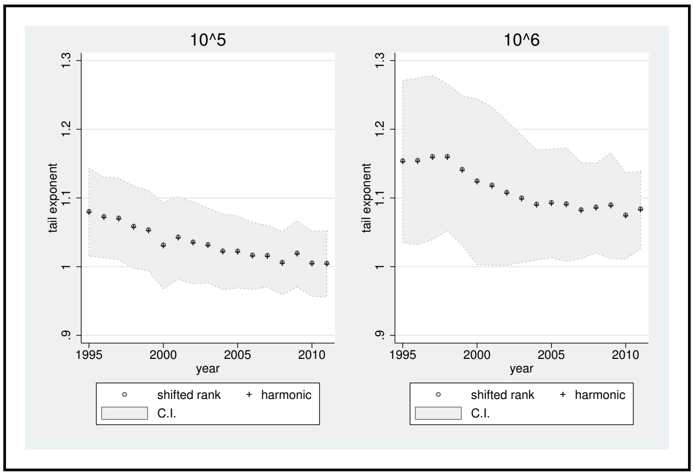

Figure 5: Tail exponent for pairs with $\geq 50\%$ variety

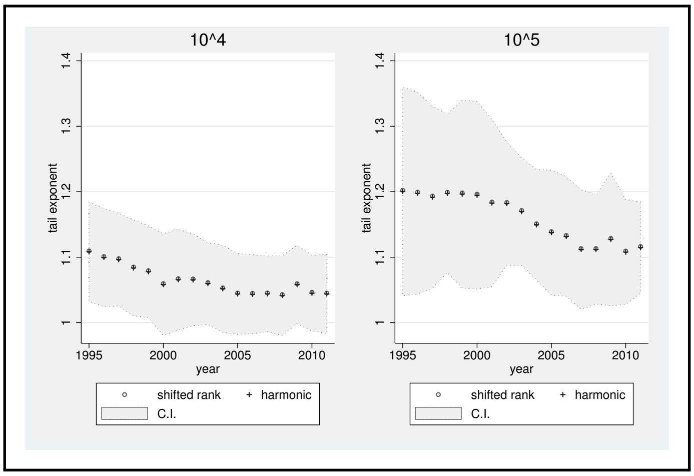

Fig.5 provides a sensitivity check. We report estimates of the tail exponent for $R = 500$

[17] The share of explained variance is 96-98%. In a subset of cases, the best fit is associated with $R = 250$. In all such cases, the confidence intervals for $R = 250$ comprise the confidence interval for $R = 500$.

[18] Results are qualitatively unchanged if the QQ-estimator of Head et al. (2014) is used instead: the tail exponent is .02-.04 higher, and it is weakly decreasing over 1995-2011.

obtained for the set of conditional sales distributions that cover $\geq 50\%$ of world variety for $\bar{x} = \{10^4, 10^5\}$.[19] Although results are qualitatively unchanged, the point estimate is sensitive to the choice of value thresholds that define the set of observations included the estimation. A qualitatively similar result is reported at the firm-level in Head et al. (2014). These authors find that the tail exponent increases when they restrict the sample to the biggest exporting firms.

Head et al. (2014) argue that the sensitivity of the point estimate to truncation puts into question the use of the Pareto assumption to model the distribution of firm-level sales. We point out that the Pareto assumption fits the data reasonably well at the country-level if we restrict the sample to exporters for whom we can define the tail of the distribution in a consistent way. To illustrate, Fig.6 plots the set of distributions used to estimate the tail exponent in the right pane of Fig.4 for the year 2007. The relationship is approximately linear, and the slope is similar for different exporters and for any given exporter across her set of export markets. These results are unchanged if we take any other year of the data.

Figure 6: Demeaned conditional sales' distributions (2007)

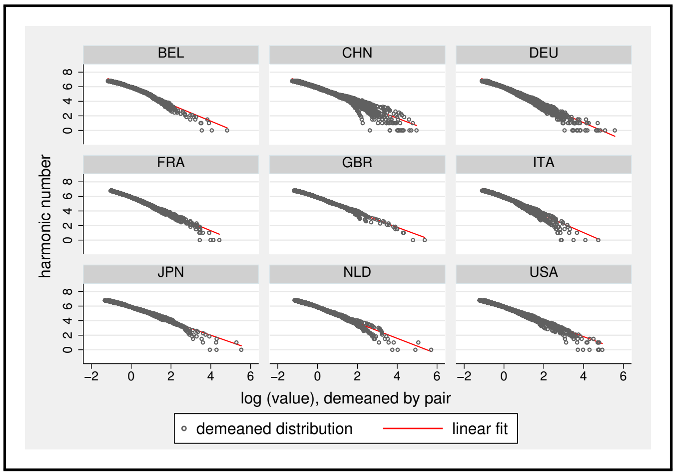

To sum up, conditional sales' distributions exhibit a pattern in the tail that conforms to a power law. We henceforth use estimates reported in the right pane of Fig.4 as the benchmark set of values for the tail exponent $\theta_t/(\sigma_t'-1)$. The point estimate of the tail exponent obtained in the benchmark specification indicates that the measured trade elasticity was 15% higher than the structural trade elasticity in 1995. This wedge was reduced to 8% by 2011.

[19] The number of exporters varies between 10 and 13 (4 and 5) in the left (right) pane. The number of observations increases from 78,500 to 124,000 (from 14,500 to 30,000) in the left (right) pane over 1995-2011.

4 The lower tier elasticity and the supply side of the economy

4.1 Identification of the lower tier demand elasticity: approach

Utility maximization at the lower tier of the CES (1) determines the allocation of expenditure among goods delivered by exporter i to market j in year t. Given total bilateral expenditure X_{ijt}, expenditure on product k is:

$$X_{ijt}(k) = \left(\frac{P_{ijt}(k)}{P_{ijt}}\right)^{1-\sigma'} X_{ijt} \tag{25}$$

To pin down the determinants of sectoral demand, we replace X_{ijt} by its value in (6), P_{ijt} by its value in (4), and the landed price of good k by its value $P_{ijt}(k) = w_{it}\tau_{ijt}/z_{it}(k)$. We get:

$$X_{ijt}(k) = \left(\frac{z_{it}(k)}{\tilde{z}_{it}}\right)^{\sigma'_t - 1} \left(\frac{w_{it}}{\tilde{z}_{it}}\right)^{-(\sigma_t - 1)} \tau_{ijt}^{1-\sigma_t} \Phi_{jt}^{\sigma_t - 1} Y_{jt} \tag{26}$$

The first term on the RHS of (26) captures a sectoral effect determined by the ratio of i's ability in sector k to i's expected ability. We see that the lower tier elasticity determines demand sensitivity to sectoral cost shocks. The second term captures an origin effect determined by the ratio of i's wage to i's expected ability. We see that the upper tier elasticity determines demand sensitivity to cost shocks that are common to all sectors. This equation underlines the distinctive feature of our demand set-up, namely the assumption that sectoral demand is more sensitive to sectoral than to economy-wide cost shocks.[20]

The main implication of this set-up is that demand shocks have no bearing on the ranking of prices within the exporter-specific product set.[21]. As long as the no-arbitrage condition on the labor market holds, demand shocks lead to changes in the equilibrium wage w_{it} but leave the ranking of prices unaffected. It follows that variation in sectoral expenditure within the exporter-specific product set directly maps into the variability of sectoral technology draws: $X_{ijt}(k)/X_{ijt}(k') = [z_{it}(k)/z_{it}(k')]^{\sigma'_t - 1}$. Consequently, we should be able to retrieve the lower tier elasticity by estimating a stochastic version of (25) after double demeaning the data.[22]

In practice, this simple approach fails for a number of reasons. First, instead of having direct information on prices, we have information on unit values in per kg terms. Notwithstanding extensive work by BACI data providers on improving the quality of unit value data (Gaulier

[20] The lower tier elasticity captures demand sensitivity to sectoral trade costs while the upper tier elasticity captures demand sensitivity to economy-wide trade costs. Our model thus provides a complementary explanation for the wedge between micro- and macro-level trade elasticities (Imbs and Méjean (2015), Ossa (2015)).

[21] We have shown in section 3 that price rankings are indeed stable across active export markets.

[22] Demeaning in the pair dimension controls for pair-specific characteristics. Demeaning in the sectoral dimension controls for sectoral characteristics that are independent of exporter-specific technology (section 3).

and Zignago (2010)), indirect price information is prone to measurement error (Hummels and Schaur (2013)). Second, labor market clearing may take time in which case demand shocks lead to a positive correlation between trade and unit values for some goods in some markets. Third, there may exist additional costs that capture the sectoral cost of producing quality.[23] Quality-related production costs can blur the mapping between the ranking of ability and the ranking of prices and lead to inconsistent estimates of σ_t' (Crozet et al. (2012)).

To deal with the first two concerns we construct a proxy of technological ability at the product level. Our model predicts that for any exporter on any market the ranking of exports relatively to some benchmark sector maps into the ranking of ability.[24] Our identification assumption is that demand shocks that occur in period t and market j situated on continent c $(j \in c)$ are independent of demand shocks that occurred in period $(t-1)$ and market j' situated on a different continent $(j' \notin c)$. As long as this assumption holds, the lagged ranking of sectoral exports in markets situated on continents other than c can be used to proxy the lagged ranking of technological ability and to instrument current prices in c. We expect this instrument to be strong because ability rankings are likely to be persistent.[25]

We split the world in three parts and obtain three datasets that each exclude a different part of the world $\{-c\}$.[26] We follow the approach of Costinot et al. (2012) and retrieve the ranking of sectoral ability for exporter i on the set of markets specific to dataset $-c$ by running the following specification separately for each exporter, year, and dataset:

$$\ln(X_{ijt}^{-c}(k)) \quad = \quad d_t^{-c} + d_{jt}^{-c} + d_t^{-c}(k) + \varepsilon_{jt}^{-c}(k) \tag{27}$$

where d_t^{-c} is the constant, d_{jt}^{-c} the market fixed effect, $d_t^{-c}(k)$ the sectoral fixed effect, and $\varepsilon_{jt}^{-c}(k)$ the error term. According to our model, the sectoral dummy $d_t^{-c}(k)$ captures sectoral ability relatively to some benchmark sector k': $d_t^{-c}(k) = \hat{d}_{it}^{-c}(k) = (\sigma_t' - 1)\ln(\hat{z}_{it}^{-c}(k)/\hat{z}_{it}^{-c}(k'))$. In each specification, we define the benchmark sector as the best observed draw for exporter i, i.e. we normalize by the good most frequently exported by i to the set of markets in $-c$.

To sum up, we instrument the ranking of period t prices on every European (resp.: American, Asian) market with help of lagged ability rankings estimated on all markets other than Europe

[23] Kugler and Verhoogen (2012) document that higher quality outputs require higher cost inputs.

[24] We show in Appendix C that the data provides support for the Ricardian prediction of our model: the ranking of relative sectoral exports for any pair of exporters is stable across markets.

[25] Hanson et al. (2015) document persistency of sectoral export capabilities at the country level.

[26] The excluded blocks $\{-c\}$ are: West and East Europe, North and South America, Asia and Rest of the World.

(resp.: America, Asia) in period $(t-1)$.[27] This approach is similar in spirit to Hummels and Schaur (2013) who use exporter-specific sales to all countries except the US as a predictor of latent product profitability on the US market.

To deal with the third concern, we model quality as in Aw and Lee (2014) and Crozet et al. (2012). We introduce the parameter $\delta_{it}(k)$ that captures the quality of product k. This parameter enters the lower tier utility function with exponent $\gamma_{jt}(k)$ that translates product quality into utility.[28] This parameter enters the cost function with exponent $\gamma_{it}(k)$ that captures the cost increment required to produce higher quality. The landed price becomes $P_{ijt}(k) = \left[\delta_{it}(k)^{\gamma_{it}(k)}/z_{it}(k)\right] w_{it}\tau_{ijt}$. As in Aw and Lee (2014), marginal costs increase with $\delta_{it}(k)$ iff quality requires more 'effective' labor units. The sectoral exports' equation (25) becomes:

$$X_{ijt}(k) = \left[P_{ijt}(k)/\delta_{it}(k)^{\gamma_{jt}(k)}\right]^{1-\sigma_t'} P^{*\,\sigma_t'-\sigma_t}_{ijt} \Phi^{*\,\sigma_t-1}_{jt} X_{jt} \tag{28}$$

where $P^*_{ijt} = \left\{ \int_0^1 \left[P_{ijt}(k)/\delta_{it}(k)^{\gamma_{jt}(k)}\right]^{1-\sigma_t'} \mathrm{d}k \right\}^{\frac{1}{1-\sigma_t'}}$ is the bilateral quality-adjusted price index and $\Phi^*_{jt} = \left[\sum_{n=1}^N \left(P^*_{nj}\right)^{1-\sigma}\right]^{1/(1-\sigma)}$ is the market-specific quality-adjusted price index.

We replace $P_{ijt}(k)$ by its value in (28), factor out $w_{it}\tau_{ijt}$ from the quality-adjusted price index, and define $\Delta_{ijt} = \left\{ \int_0^1 \left[\frac{\delta_{it}(k)^{\gamma_{it}(k)}}{z_{it}(k)\delta_{it}(k)^{\gamma_{jt}(k)}}\right]^{1-\sigma_t'} \mathrm{d}k \right\}^{\frac{1}{1-\sigma_t'}}$ to get:

$$X_{ijt}(k) = \left[\frac{z_{it}(k)}{\delta_{it}(k)^{\gamma_{it}(k)}}\right]^{\sigma_t'-1} \delta_{it}(k)^{\gamma_{jt}(k)(\sigma_t'-1)} \left[w_{it}\tau_{ijt}\right]^{1-\sigma_t} \Delta^{\sigma_t'-\sigma_t}_{ijt} \Phi^{\sigma_t-1}_{jt} Y_{jt} \tag{29}$$

We learn three things. First, the ranking of prices within the exporter-specific product set remains invariant to market characteristics. Second, our proxy of lagged sectoral ability is expected to map into quality-adjusted prices because it incorporates the cost of producing quality: $\hat{d}^{-c}_{it}(k) = (\sigma_t'-1)\left[\ln(\hat{z}^{-c}_{it}(k)/\hat{z}^{-c}_{it}(k')) - (\gamma_{it}(k))\ln(\hat{\delta}^{-c}_{it}(k)) + (\gamma_{it}(k'))\hat{\delta}^{-c}_{it}(k')\right]$. So long as consumer valuation of quality is destination-sector specific, the term $\delta_{it}(k)^{\gamma_{jt}(k)}$ is picked up by the residual of (27), and our proxy is not plagued by consumer valuation of quality. Third, consistent estimation of $(\sigma_t'-1)$ requires additional controls that extract $[\delta_{it}(k)]^{\gamma_{jt}(k)}$ from the residual of the sectoral exports' equation.

[27] We only retain precisely estimated rankings, i.e. those for which the correlation coefficient between standardized and non-standardized sectoral dummies exceeds .5.

[28] The lower tier utility function (1) becomes $Q_{ijt} = \left[\int_0^1 \left[\delta_{it}(k)^{\gamma_{jt}(k)} Q_{ijt}(k)\right]^{\frac{\sigma_t'-1}{\sigma_t'}} \mathrm{d}k \right]^{\frac{\sigma_t'}{\sigma_t'-1}}$.

We argue that we can consistently estimate the lower tier elasticity if we double-demean the data in the pair and sector dimensions and include the frequency variable as an additional control in the estimation. The frequency variable $f_{it}(k) = n_{it}(k)/N_{it}$ counts the number of times that exports of k exceed the statistical threshold $\bar{X}$ in the set of i's active markets. In the price equation, the frequency variable helps to disentangle the positive linkage between prices and quality from the negative linkage between prices and ability. In the value equation, the frequency variable helps to extract the term $\delta_{it}(k)^{\gamma_{jt}(k)}$ from the residual.

The intuition is the following. The predicted component of the product price $\hat{P}_{ijt}(k)$ is determined by the lagged effective unit labor requirement $\left[\delta_{i,t-1}(k)^{\gamma_{i,t-1}(k)}/z_{i,t-1}(k)\right]$ and is - by construction - orthogonal to product frequency $f_{it}(k)$. From our model we know that conditional on $\left[z_{it}(k)/\delta_{it}(k)^{\gamma_{it}(k)}\right]$, the probability of observing k in j is always increasing in consumer valuation of quality $\delta_{it}(k)^{\gamma_{jt}(k)}$. Consequently, we expect the deviation of frequency from the pair- and sector-specific mean to pick up pair-sector specific demand shifters.

4.2 Identification of the lower tier demand elasticity: estimation

In the first stage we estimate the following equation:

$$\ln\left[P^{c}_{ijt}(k)\right] = b_0 + b_{dt}\left[\hat{d}^{-c}_{i,t-1}(k)\right] + b_{ft}\ln\left[f_{it}(k)\right] + b_t(l) + b_{ijt} + \varepsilon^{c}_{ijt}(k) \quad (30)$$

where $b_t(l)$ is the sectoral fixed effect at the HS 4-digit level and b_{ijt} is the pair fixed effect. The number of observations in the estimation increases from 2.2 million in 1996 to 3.5 million in 2011. The share of explained variance in the price equation increases from 38-40% in 1996-1999 to 49-53% in 2000-2011.[29]

The coefficients of interest are b_{dt} and b_{ft}. We expect the proxy of lagged ability to pick up the negative relationship between prices and technology draws whereby $b_{dt} < 0$. The sign of b_{ft} is ambiguous. The coefficient is positive if the frequency variable primarily picks up the positive relationship between the extra cost of producing quality and the fact that higher quality goods generate more revenue. It is negative if the quality channel plays a minor role in determining whether the product is registered in trade statistics.

Fig.7 summarizes our results for b_{ft} in the left pane and for b_{dt} in the right pane. The coefficient on the frequency variable b_{ft} changes sign: from negative or insignificant in 1996-2002, it becomes positive and significant in 2003-2011. The coefficient on lagged technological

[29] Exporter-market pairs with fewer than 240 goods are dropped. Standard errors are clustered by exporter-HS 6-digit good. Kleibergen Papp statistics indicate that the instrument is a strong predictor of prices.

Figure 7: First stage results ($-c$)

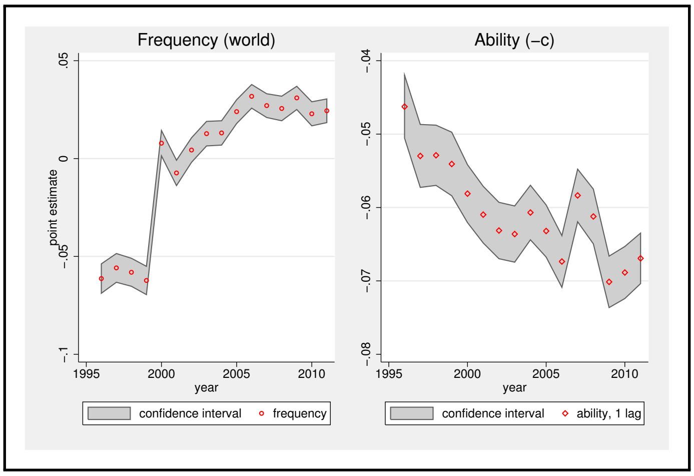

ability b_{dt} is negative and significant in all years.

In the second stage we estimate the following equation:

$$\ln\left[X^c_{ijt}(k)\right] \quad = \quad \beta_0 - \beta_{dt}\left[\hat{P}^c_{ijt}(k)\right] + \beta_{ft}\ln\left[f_{it}(k)\right] + \beta_t(l) + \beta_{ijt} + \varepsilon^c_{ijt}(k) \tag{31}$$

where $\beta_t(l)$ is the sectoral fixed effect at the HS 4-digit level and β_{ijt} is the pair fixed effect.

The coefficients of interest are β_{dt} and β_{ft}. As explained in the previous section, the coefficient on the instrumented price is expected to be negative and significant ($\beta_{dt} > 0$) while the coefficient on the frequency variable is expected to be positive and significant ($\beta_{ft} > 0$).

Figure 8: Second stage results ($-c$)

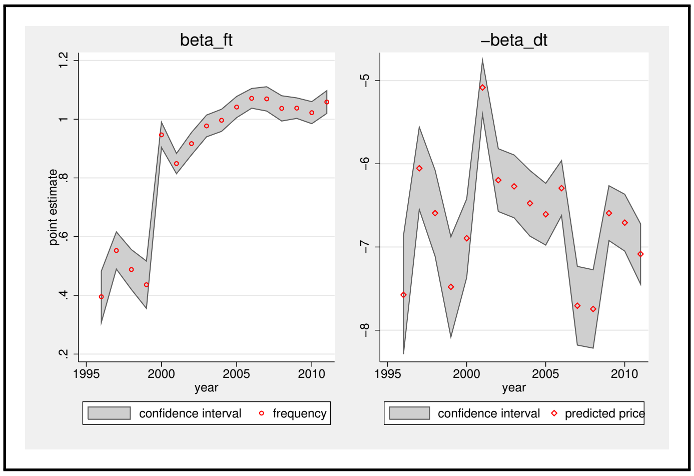

Fig.8 summarizes our baseline results for β_{ft} in the left pane and for $-\beta_{dt}$ in the right pane. The coefficient on the frequency variable is strictly positive. The coefficient on the instrumented price is always strictly negative. In most years, this coefficient is in the $\{-8, -6\}$ range. Our results for the second stage of the estimation remain qualitatively and quantitatively unchanged if we change the way of constructing the instrument or of defining the sample (Appendix D).

Figure 9: The lower tier elasticity ($\hat{\sigma}'_t$)

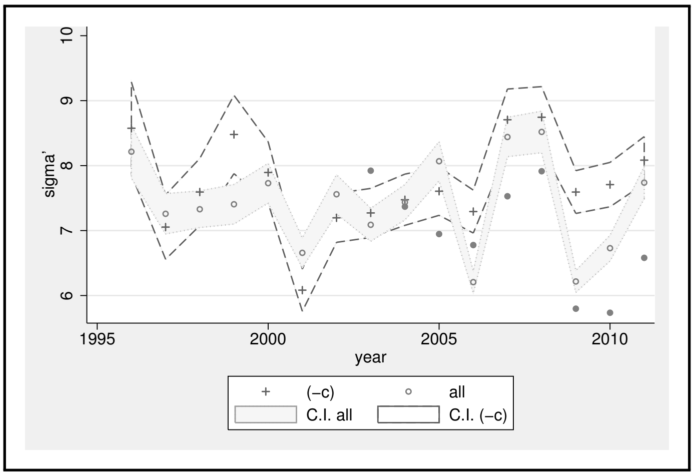

The price coefficient informs about the magnitude of the lower tier elasticity: $\beta_{dt} = \hat{\sigma}'_t - 1$. In Fig.9 we summarize our results for $\hat{\sigma}'_t$ in the benchmark specification (+) and in the two robustness checks (○, ●) discussed in Appendix D. In most years, $\hat{\sigma}'_t \in \{7, 9\}$.

In our model, the lower tier demand elasticity measures the percentage change in relative expenditure on HS 6-digit products of different origin following a one percent change in relative sectoral costs (26). To give an example, this elasticity captures how relative expenditure on boy's T-shirts produced in China and in the USA reacts to a change in the relative tariff imposed on boy's T-shirts of US and Chinese origin. Consequently, its magnitude should be compared to estimates of trade elasticities obtained at a highly disaggregate level. Parameter magnitudes obtained at comparable levels of disaggregation are in line with our estimates.[30]

4.3 The supply side of the economy

Our estimates of the lower tier elasticity help to characterize the supply side of the economy by retrieving the shape and scale parameters of the the Fréchet distribution.[31]

[30] See Imbs and Méjean (2015): Broda and Weinstein (2006) report an average Armington elasticity of 5.9 at SITC 3-digit but 11.7 at the TSUSA 10-digit level. Hummels (2001) finds 8.3 at the SITC 4-digit level. Head and Ries (2001) report $\{8, 11\}$ at the SIC 3-digit level. Imbs and Méjean (2015) report 6.4 for 56 ISIC sectors.

[31] We cannot separately identify the distribution of $z_i(k)$ and of $\delta_i(k)^{\gamma_i(k)}$. To keep things simple, we assume that the reciprocal of the effective unit labor requirement, $\hat{z}_i(k) = \left[z_i(k)/\delta_i(k)^{\gamma_i(k)}\right]$, is distributed Fréchet.

4.3.1 The shape parameter θ_t

The shape parameter θ_t is common to all exporters. We deduce its magnitude by combining the estimate of the tail exponent $\theta_t/(\sigma_t' - 1)$ with the estimate of σ_t'. Fig.10 summarizes our results for the benchmark specification (+) and for the two robustness checks (∘, •) discussed in Appendix D. We find that $\hat{\theta}_t \in \{6, 8\}$ in most years. These magnitudes are in line with estimates obtained in Eaton and Kortum (2002) and Costinot et al. (2012).

In the first robustness check (∘) we change the way of constructing the instrument. Specifically, we now use the full set of markets to construct the lagged ranking of ability for each exporter. We obtain parameter estimates that are almost everywhere comprised within the confidence interval of the benchmark specification. In the second robustness check (•) we reduce the sample by two orders of magnitude and focus on the set of pairs ij used in the benchmark estimation of the tail exponent in section 3.2. We see that the point estimate of the shape parameter remains in the range of the benchmark specification for the subset of years in which the instrument is sufficiently strong: $\hat{\theta}_t \in \{6, 8\}$ in 2003-2008 and $\hat{\theta}_t \in \{5, 6\}$ in 2009-2011.

Figure 10: The technology dispersion parameter ($\hat{\theta}_t$)

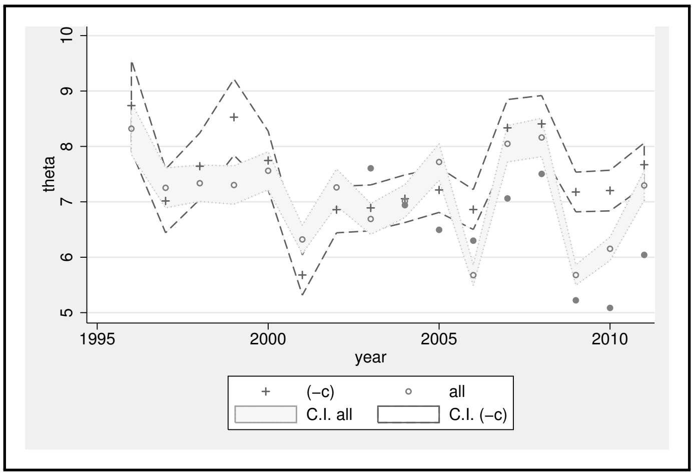

4.3.2 The scale parameters $\tilde{z}_{it}$

The scale parameter $\tilde{z}_{it}$ is specific to the exporter, and it captures her expected technological ability. We use the estimate of σ_t' together with ability rankings obtained for each exporter on the full set of export markets to retrieve the scale parameters $\tilde{z}_{it}$.

First, we retrieve the annual distribution of ability. We rescale the estimated constant $\hat{d}_t$ and the estimated sectoral dummies $\hat{d}_{it}(k)$ by $(\hat{\sigma}'_t - 1)$ and add the rescaled constant to the rescaled dummies. The rescaled constant provides an estimate of ability in the benchmark sector $\hat{d}_t/(\hat{\sigma}'_t - 1) = ln(\hat{z}_t(k'))$. The adjusted dummies provide ability estimates in the other sectors $\hat{d}_{it}(k)/(\hat{\sigma}'_t - 1) + ln(\hat{z}_t(k')) = ln(\hat{z}_t(k))$. Second, we invert the CDF of the ability distribution to show that any of its empirical quantiles can be used to deduce its scale parameter. Denoting by p the cumulative probability and by $q_{it}(p)$ the corresponding quantile, we have the following relationship for the Fréchet: $q_{it}(p) = \tilde{z}_{it}\left[-ln(p)\right]^{-1/\theta_t}$.

We observe just a subset of these quantiles in the data. We explicitly account for truncation by associating the smallest observed empirical quantile $q_{it}(p_{min})$ to the smallest estimated ability: $q_{it}(p_{min}) = \hat{z}_{it}(k_{min})$.[32] Denoting world variety by V_t and exporter-specific variety by V_{it}, the smallest observed cumulative probability is $p_{min} = 1 - V_{it}/V_t$. We obtain the other empirical quantiles by ranking estimated abilities in increasing order $k_{min} \leq k_v$ where $v \in \{1, ..., (V_{it} - 1)\}$ and we associate each element of this ranking to its cumulative probability: $p_v = 1 - (V_{it} - v)/V_t$.[33] We implement the QQ-estimator (Head and Mayer (2014), Kratz and Resnick (1996)) to estimate the following relationship separately for each country and year:

$$\ln\left[q_{it}(p_v)\right] = \ln\left[\hat{z}_{it}(k_v)\right] = \ln(\tilde{z}_{it}) + \ln\left\{\left[-\ln(p_v)\right]^{-1/\hat{\theta}_t}\right\} + z_{itv} \tag{32}$$

The exponentiated constant provides an estimate of the expected technological ability of exporter i in year t. The ranking of exponentiated constants provides the ranking of exporters in terms of their technological ability. This ranking is not sensitive to using the point estimate or any other value of σ'_t within the confidence interval of our estimates of the lower tier elasticity.

Fig.11 shows that the estimated ranking of technological ability explains most of the variation in export variety registered in trade statistics for our set of 92 exporters. This result conforms to the predictions of our model. Fig.11 also shows that this ranking is not sensitive to including all observed quantiles in the estimation (○) or restricting the set of quantiles to the upper half of the ability distribution (+). In Fig.12 we plot real GDP and real GDP per capita against the ranking of country-specific technological ability to illustrate that the latter is a better predictor of income per capita than of economic size.[34]

[32] Unobserved quantiles of the productivity distribution must be strictly smaller than any of the observed quantiles.

[33] In the estimation we follow Head et al. (2014) and adjust the definition of the cumulative probabilities in the following way: $p_v = 1 - (V_{it} - v + 1 - .3)/(V_t + .4)$.

[34] Real GDP in current and chained PP prices is obtained from the Penn World Tables 8.1 (Feenstra et al. (2013)).

Figure 11: Exporter-specific variety and technological ability ($\tilde{z}_{it}$)

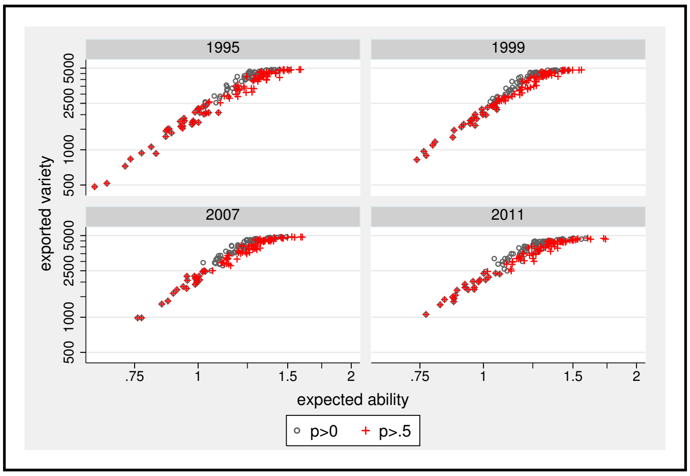

Figure 12: Real GDP per capita and technological ability ($\tilde{z}_{it}$)

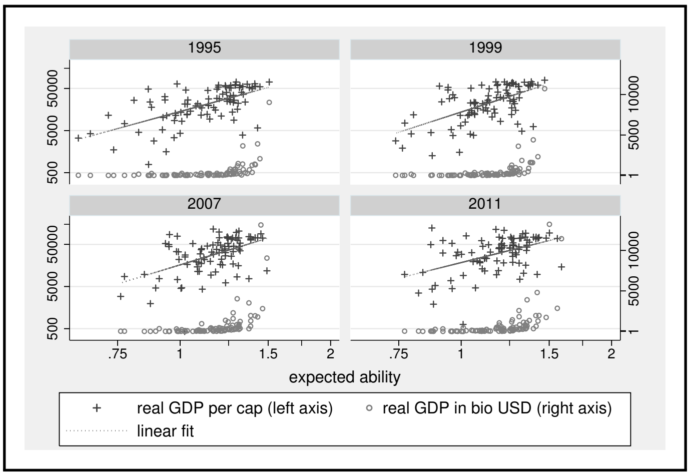

5 The upper tier demand elasticity

5.1 Perceived substitutability of truncated composite goods ($\bar{\sigma}_t$)

We use the point estimate of the lower tier demand elasticity $\hat{\sigma'}_t$ to aggregate observed product prices. Direct implementation of the CES formula delivers bilateral price indices of truncated exporter-specific composite goods: $\bar{P}_{ijt} = \left[\sum_k P_{ijt}(k)^{1-\hat{\sigma}'_t}\right]^{\frac{1}{1-\hat{\sigma}'_t}}$.

According to the benchmark model, we can retrieve an estimate of $\bar{\sigma}_t$ by direct implementation of the aggregate demand equation for truncated composite goods (8). However, this approach fails to deliver consistent estimates of $\bar{\sigma}_t$ in the presence of unobserved quality. Expenditure is then determined by quality-adjusted price indices that take into account consumer valuation of quality $\bar{P}^{*}_{ijt} = \left\{ \int\limits_{0}^{\bar{v}} \left[P_{ijt}(k) / \delta_{it}^{\gamma_{jt}(k)} \right]^{1-\sigma'_t} \mathrm{d}k \right\}^{\frac{1}{1-\sigma'_t}}$ whereby (8) becomes:

$$\frac{\bar{X}_{ijt}}{\bar{Y}_{jt}} = \frac{\left(\bar{P}^{*}_{ijt}\right)^{1-\bar{\sigma}}}{\sum_{n=1}^{N} \left(\bar{P}^{*}_{njt}\right)^{1-\bar{\sigma}}} \quad \Leftrightarrow \quad \bar{X}_{ijt} = \left[w_{it}\tau_{ijt}\right]^{1-\bar{\sigma}} \left(\bar{\Delta}\right)_{ijt}^{1-\bar{\sigma}} \left(\bar{\Phi}^{*}_{jt}\right)^{\bar{\sigma}_t - 1} \bar{Y}_{jt} \tag{33}$$

where $\bar{\Delta}_{ijt} = \left\{ \int\limits_{0}^{\bar{v}} \left[\delta_{it}^{\gamma_{it}(k)} / z_{it}(k) \delta_{it}^{\gamma_{jt}(k)} \right]^{1-\sigma'_t} \mathrm{d}k \right\}^{\frac{1}{1-\sigma'_t}}$ and $\bar{\Phi}^{*}_{jt} = \left[\sum_{n=1}^{N} \left(\bar{P}^{*}_{njt}\right)^{1-\bar{\sigma}} \right]^{\frac{1}{1-\bar{\sigma}_t}}$.

Bilateral price indices that we construct on the basis of observed product prices contain the cost of producing quality: $\bar{P}_{ijt} = w_{it}\tau_{ijt} \left\{ \int\limits_{0}^{\bar{v}} \left[\delta_{it}^{\gamma_{it}(k)} / z_{it}(k) \right]^{1-\sigma'_t} \mathrm{d}k \right\}^{\frac{1}{1-\sigma'_t}}$. Consequently, we expect these indices to be correlated with the residual of the exports' equation.

Consistent estimation of (33) requires instruments for bilateral price indices $\bar{P}_{ijt}$. We use bilateral distance and adjacency as predictors of bilateral trade costs. We use total population as a predictor of the unit labor cost. In both stages of the estimation, we include the proxy of the expected technological ability $\tilde{z}_{it}$ (section 4.3) to control for quality-adjusted exporter ability. We include market fixed effects d_{jt} to control for destination-specific characteristics.

The equation that we estimate in the first stage is:

$$\ln\left(\bar{P}_{ijt}\right) = b_0 + b_{1t}\ln[POP] + b_{2t}\ln[DIST] + b_{3t}[CONTIG] + b_{4t}\ln\tilde{z}_{it} + d_{jt} + \varepsilon_{ijt} \tag{34}$$

The share of explained variance in the price equation varies between 34 and 41%. The number of observations increases from 6,180 to 9,034 over 1995-2011. Fig.13 presents the results. The coefficients estimated on the three instruments are presented in the right pane: they all have the expected signs. Kleibergen Papp statistics indicate that they are strong predictors of prices. The coefficient on quality-adjusted ability is graphed in the left pane: it is insignificant in 1995-2002 but positive and significant in 2003-2011.

The equation that we estimate in the second stage is:

$$\ln\left(\bar{X}_{ijt}\right) = \beta_0 - \beta_{1t}\ln\left[\hat{\bar{P}}_{ijt}\right] + \beta_{4t}\ln\tilde{z}_{it} + d_{jt} + \varepsilon_{ijt} \tag{35}$$

Fig.14 presents the results for $-\beta_{1t}$ in the right pane and for β_{4t} in the left pane. The coefficient on the price variable is precisely estimated. It increases in magnitude from about 1.9 to 2.4 over 1995-2011. The coefficient on the ability variable is positive and significantly higher in magnitude than the coefficient on the price variable. One possible interpretation of this result is that the variable picks up consumer valuation of quality on top of exporter ability.

Figure 13: First stage results

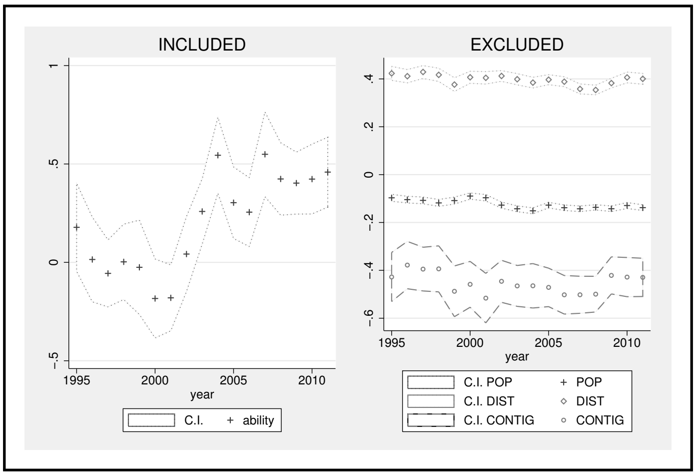

Figure 14: Second stage results

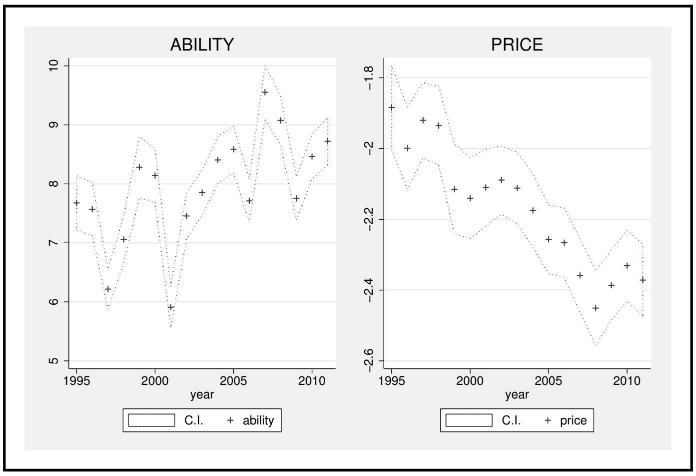

The price coefficient informs us about the perceived substitutability of truncated composite goods: $\beta_{1t} = \bar{\sigma}_t - 1$. In the right pane of Fig.15 we present our results for $\bar{\sigma}_t$ obtained in the benchmark specification. The central point in each year corresponds to the estimate of $\bar{\sigma}_t$ obtained when we use the point estimate of the lower tier elasticity $\hat{\sigma}_t'$ to compute the bilateral price index. We find that $\bar{\sigma}_t$ increases by 18% over 1995-2011, from 2.9 to 3.44. The two diamonds adjacent to each central point correspond to estimates of $\bar{\sigma}_t$ obtained when we use the upper and lower bounds of the lower tier elasticity estimates to compute the bilateral price index. We see that our estimates of $\bar{\sigma}_t$ are not sensitive to the choice of $\hat{\sigma'}_t$.

We represent the point estimate of $\bar{\sigma}_t$ with a cross when the set of instruments verifies overidentifying restrictions and with a circle when it does not. The set of instruments used

Figure 15: The measured upper tier demand elasticity $\bar{\sigma}_t$

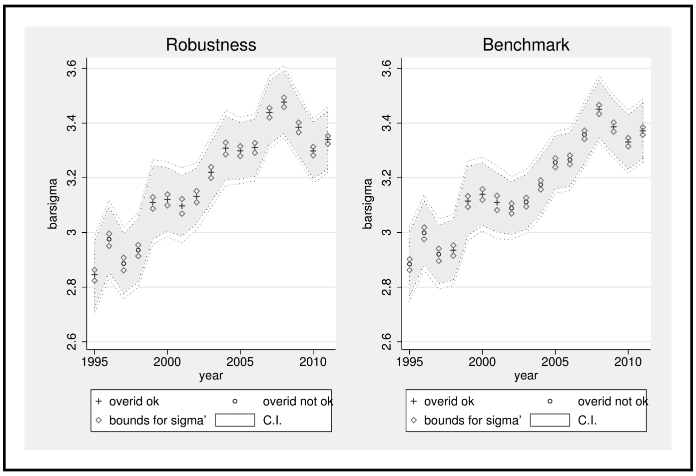

in the benchmark specification verifies overidentifying restrictions only in a subset of years. Consequently, in the left pane of Fig.15, we report the results ofa robustness check in which population is used as a control variable in the price and value equations. The two remaining instruments pass overidentifying restrictions in almost all years while results for $\bar{\sigma}_t$ remain unchanged: the coefficient increases from 2.91 to 3.48 over 1995-2011, i.e. by about 19.5%.

5.2 The structural upper tier demand elasticity (σ_t)

We rearrange (19) to solve for the structural upper tier demand elasticity σ_t as a function of the wedge $\theta_t/(\sigma'_t - 1)$, the lower tier elasticity σ'_t, and the measured upper tier elasticity $\bar{\sigma}_t$:

$$\sigma_t = \left[\frac{\theta_t}{(\sigma'_t - 1)}\frac{(\sigma'_t - \bar{\sigma}_t)}{(\bar{\sigma}_t - 1)} + \sigma'_t\right] / \left[\frac{\theta_t}{(\sigma'_t - 1)}\frac{(\sigma'_t - \bar{\sigma}_t)}{(\bar{\sigma}_t - 1)} + 1\right] \tag{36}$$

In the right pane of Fig.16 we present the results for the benchmark specification in which population is only used in the price equation. In the left pane of Fig.16 we present the results for the robustness check in which population is used in the price and value equations.

The central point in each year corresponds to the point estimate of $\hat{\sigma}_t$ that we obtain when we use the point estimates of $\theta_t/(\sigma'_t - 1)$, σ'_t, and $\bar{\sigma}_t$ to compute the structural upper tier demand elasticity. The structural demand elasticity $\hat{\sigma}_t$ increases from 2.7 to 3.3 over 1995-2011, i.e. by 22%. The two diamonds adjacent to this central point correspond to estimates obtained when we use the upper and lower bounds of $\hat{\sigma}'_t$ and the corresponding estimates of $\bar{\sigma}_t$ to compute $\hat{\sigma}_t$. In Appendix D.2 we show that our results on the evolution of the upper tier demand elasticity are not sensitive to the estimate of the tail exponent used to compute parameter magnitudes.

Figure 16: The structural upper tier demand elasticity ($\hat{\sigma}_t$)

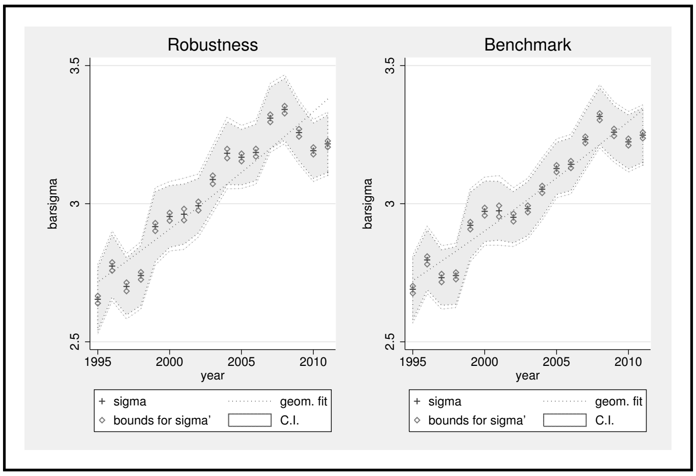

5.3 The structural trade elasticity (ε_t)

The magnitude of the structural trade elasticity is directly determined by the magnitude of the structural upper tier demand elasticity: $|\hat{\varepsilon}_t| = (\hat{\sigma}_t - 1)$. Fig.17 summarizes our results on $|\hat{\varepsilon}_t|$ that we deduce from our estimates of the structural upper tier demand elasticity (section 5.2).[35] In the right pane we present the results for the benchmark specification in which population is only used in the price equation. In the left pane we present the results for the robustness check in which population is used in the price and value equations. The structural trade elasticity $\hat{\varepsilon}_t$ increases from 1.7 to 2.3 over 1995-2011, i.e. by 35%.

Fig.17 also illustrates the magnitude of the wedge between the measured and the structural trade elasticity attributable to truncation. This wedge is determined by the tail exponent: $|\bar{\varepsilon}_t| = |\hat{\varepsilon}_t|\theta_t/(\sigma_t' - 1)$. The measured trade elasticity increases from 1.95 to 2.44, i.e. by about 25%.

These magnitudes are lower than the 3.7-5.5 range for the trade elasticity reported in Simonovska and Waugh (2014b) for the benchmark Armington and Ricardian models. This discrepancy may be due to the fact that our model takes into account the contribution of multiple dimensions of heterogeneity in generating the observed distribution of prices and expenditure while Simonovska and Waugh (2014b) study models that focus on a single dimension of heterogeneity. Our results are consistent with Simonovska and Waugh (2014b) to the extent that these authors show that more flexible models map into lower magnitudes of the trade elasticity.

[35] In Appendix D.2 we illustrate that our results on the evolution of the trade elasticity are qualitatively unchanged if we use the bounds of our estimates for the tail exponent to compute parameter magnitudes.

Figure 17: The trade elasticity ($\hat{\varepsilon}_t$)

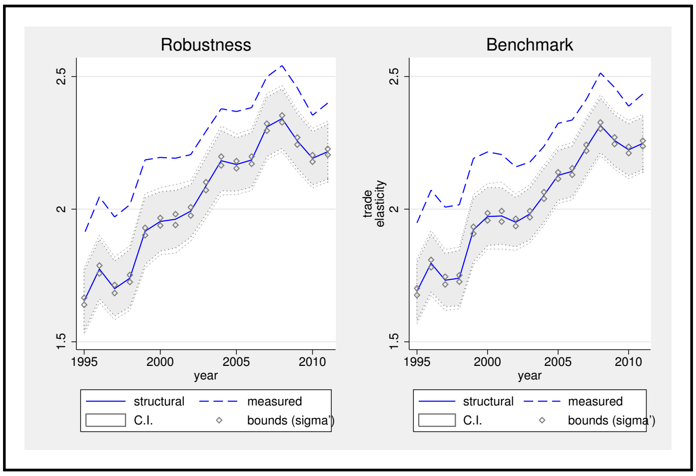

6 The purchasing power channel of the gains from trade

We illustrate the empirical relevance of our parametric approach to the estimation of the trade elasticity with help of an application. As shown by Behrens et al. (2014), any framework that keeps the set of goods fixed across trade equilibria belongs to the class of models discussed in Arkolakis et al. (2012) in which gains from trade can be computed with help of two sufficient statistics: the trade elasticity ε_t and the share of expenditure on domestic supply Λ_{iit}.

The set of goods is fixed in our model, and gains from trade are realized on the intensive margin. Reductions in trade barriers increase consumer purchasing power by improving access to a fixed basket of goods. Following Arkolakis et al. (2012), we compute the 'welfare cost of autarky', i.e. the percentage loss in real income incurred by moving to autarky from an equilibrium in which the country relies on foreign supply: $\Delta_{it} = 1 - \Lambda_{iit}^{1/|\varepsilon_t|}$.

First, we characterize reliance on domestic supply. Second, we combine this information with our estimates of the structural trade elasticity to quantify the increase in real income due to trade. Third, we show that relaxing the assumption of a time-invariant trade elasticity removes the mechanical link between reliance on foreign supply and consumer purchasing power.

6.1 Reliance on domestic supply

We use information on production (*PROD*), exports (X), and imports (M) reported at the ISIC 4-digit level in the UNIDO Supply and Demand database to compute the indicator of reliance on domestic supply (Λ_{iit}) in 1995-2010. We compute total domestic supply as $X_{iit} =$

$\sum_k \{PROD_{it}(k) - X_{it}(k)\}$ and total expenditure as $EXP_{it} = \sum_k \{PROD_{it}(k) - X_{it}(k) + M_{it}(k)\}$ whereby reliance on domestic supply is: $\Lambda_{iit} = X_{iit}/EXP_{it}$.

Fig.18 reports the interquartile range of reliance on domestic supply in the sample of 75 countries observed in the UNIDO SD database for 8 or more years. The IQR shifted from 56-71% in 1995 to 49-62% in 2010. Reliance on domestic supply at the median point of the sample decreased from about 65% to about 57% (+). In this figure we also plot the IQR of reliance on domestic supply in the sample of 61 countries observed in the UNIDO SD database for 8+ years that also belong to the sample of 92 exporters that we use to estimate the structural trade elasticity. The indicator is little sensitive to this change. In particular, reliance on domestic supply at the median point of the reduced sample decreases from about 65% to about 59% (∘).

Figure 18: Reliance on domestic supply (Λ_{iit})

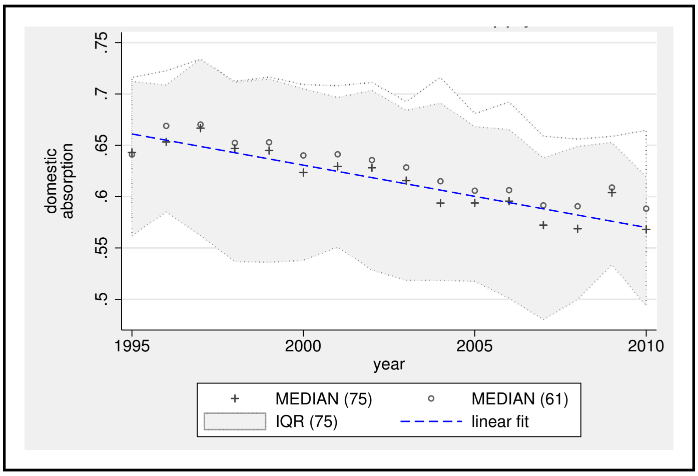

6.2 The purchasing power channel of the gains from trade

We compute the increase in real income associated to the purchasing power channel of the gains from trade by combining information on the structural trade elasticity $\hat{\varepsilon}_t$ with information on Λ_{iit}. We rely on estimates of the trade elasticity obtained when the benchmark set of instruments is used to estimate $\bar{\sigma}_t$ because this specification delivers the most conservative results on the evolution of the trade elasticity.[36] Further, we use the interquartile range of reliance on domestic supply in the sample of 75 countries present in UNIDO SD in ≥ 8 years.[37] To cover

[36] Results are qualitatively unchanged if we use estimates of the trade elasticity associated to the robustness check in which only distance and adjacency are used as instruments in the estimation of $\bar{\sigma}_t$.

[37] Results are not sensitive to using the interquartile range for the reduced 61-country sample.

the same timeframe as in the rest of the paper, we make the assumption that $\Lambda_{ii,2011} \approx \Lambda_{ii,2010}$.[38]

Figure 19: The purchasing power channel of the gains from trade

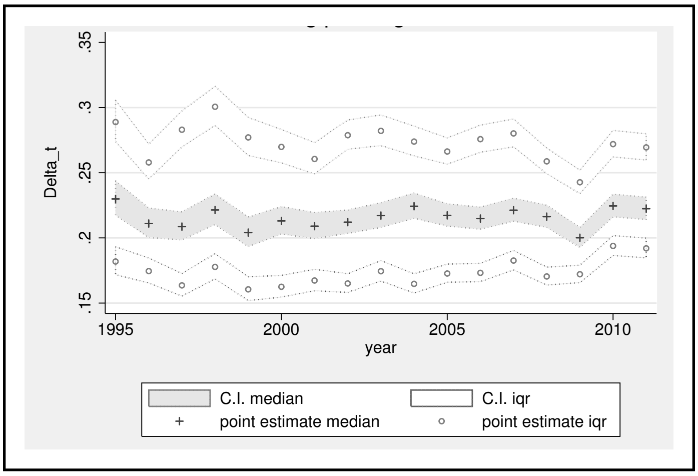

Fig.19 summarizes our results on the increase in real income associated to the purchasing power channel of the gains from trade at the median point of our sample (the central bar on the graph) as well as at the upper and lower bounds of the IQR (resp. the lower and upper bars on the graph). The confidence intervals for the increase in real income are obtained by using the confidence interval for our estimates of $|\hat{\varepsilon}_t|$.

Figure 20: The impact of a time-varying trade elasticity on the gains from trade

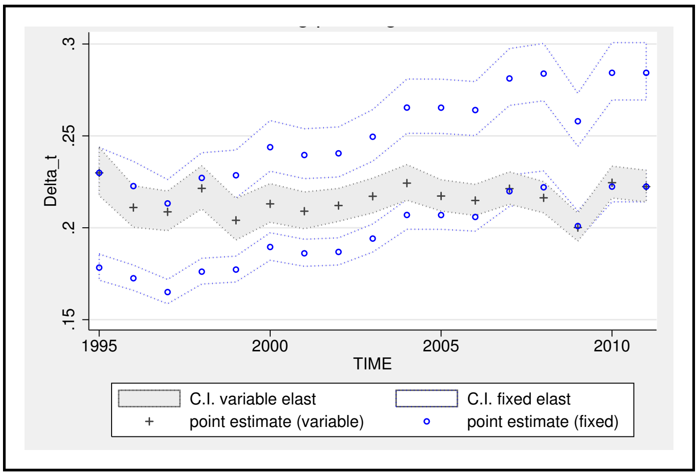

The purchasing power channel of the gains from trade is far from negligible. Fig.19 documents that access to foreign supply increases real income by 15-30% at the IQR of our sample.

38 Results are not sensitive to instead applying the annualized change in Λ_{iit} (-.4pp) to obtain $\Lambda_{ii,2011}$.

Further, increased reliance on foreign supply does not mechanically lead to an increase in real income. To give an example, the increase in real income associated to the purchasing power channel of the gains from trade remains stable at 22-23% at the median point of our sample.

In Fig.20 we document that this stability is fully attributable to the increase in the structural trade elasticity. The middle bar in Fig.20 reproduces our results for the sample median reported in Fig.19 where we allow the trade elasticity to vary over time. The upper (resp. lower) bar shows how the purchasing power channel of the gains from trade would have evolved if the trade elasticity were kept fixed at its initial (resp. final) value. If we assumed that the trade elasticity were time-invariant, we would conclude that increased reliance on foreign supply led to a 4-5 percentage point increase in real income over 1995-2011.

Figure 21: The tail exponent and the magnitude of welfare gains from trade

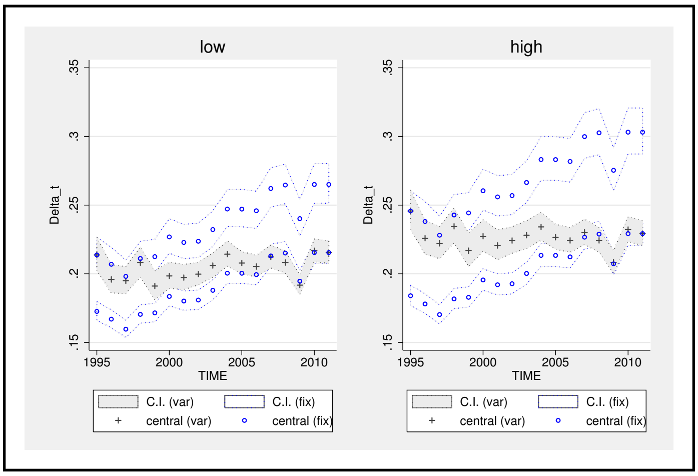

We may be worried that the numbers reported in Fig.20 are contingent on the choice of a particular value of the tail exponent. Indeed, up to this point, we have used the point estimate of $\theta_t/(\sigma_t' - 1)$ to compute the range of estimates for $\hat{\varepsilon}_t$. Consequently, we also use the bounds of the confidence interval for the tail exponent to compute the trade elasticity $\hat{\varepsilon}_t$. Given $\hat{\varepsilon}_t$, we recompute the magnitude of welfare gains at the median point of the sample.

In the left (resp. right) pane of Fig.21 we graph the evolution of welfare gains when the lower (resp. upper) bound of the confidence interval for the tail exponent is used to compute the trade elasticity. As in Fig.20, the middle bar corresponds to welfare gains associated to a time-varying trade elasticity while the upper (resp. lower) bar corresponds to a trade elasticity fixed at its initial (resp. final) value. Again, we find that real income increases by 4-6 percentage points if the trade elasticity is assumed to be time-invariant. The contribution of trade to consumer purchasing power is stable if the trade elasticity is allowed to vary over time.

7 Conclusion

The increased interdependence of countries within the world trade system has been documented extensively. But the important question from the perspective of the consumer is whether - and by how much - welfare gains from trade increased as a consequence of increased reliance on foreign supply. Recent work on this topic has argued that the answer to this question is contingent on parametric assumptions made on the supply and demand structure of the economy. In this paper we highlight a concomitant issue that makes it difficult to provide a general answer. Specifically, we argue that results may be highly sensitive to the commonplace assumption that structural parameters are time-invariant.

The magnitude of welfare gains from trade is co-determined by the strength of the impediments and of the incentives to trade. In this paper we show that both are subject to change over time. We make this point with help of a highly stylized model of the world economy that nests technological heterogeneity à la Eaton and Kortum (2002) within the canonical Armington model of trade of Anderson and van Wincoop (2003). We work with this particular parameterization of technology and preferences because it enables us to pin down the magnitude of one particular channel of the gains from trade and to track its evolution over time.

In our generalized Armington model, the set of goods is fixed across trade equilibria. Gains from trade are determined by the increase in consumer purchasing power associated to improved access to a fixed set of goods. The magnitude of welfare gains in such a model is independent of the intrinsic valuation of variety which may be impossible to recover. Moreover, two statistics suffice to determine their magnitude: the trade elasticity and the extent of reliance on domestic supply. Our parameterization of technology and preferences enables us to obtain annual estimates of the structural trade elasticity in the absence of information on trade frictions.

We find that access to foreign supply leads to a sizeable increase in consumer purchasing power because the structural trade elasticity is low in magnitude. Specifically, we document a 15-30% increase in real income at the interquartile range of the sample. Further, we find that the increasing sensitivity of expenditure to cost differences counterbalanced the effect of trade deepening on consumer purchasing power. Specifically, we document a 7-9 percentage point increase in reliance on foreign supply over 1995-2011 and a concomitant 35% increase in the magnitude of the trade elasticity. The interplay of these two processes kept the magnitude of purchasing power gains from trade unchanged over 1995-2011. We conclude that by relaxing the assumption of a time-invariant trade elasticity we remove the mechanical link between reliance on foreign supply and consumer purchasing power.

References

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

Archanskaia, E. and G. Daudin (2014). Product heterogeneity and the distance puzzle. *unpublished manuscript*.

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

Arkolakis, C., S. Ganapati, and M.-A. Muendler (2014). The extensive margin of exporting products: A firm-level analysis. *NBER Working Paper Series* (16641).

Aw, B. Y. and Y. Lee (2014). A model of demand, productivity, and foreign location decision among Taiwanese firms. *Journal of International Economics 92*, 304–316.

Baldwin, R. and J. Harrigan (2011). Zeros, quality and space: trade theory and trade evidence. *American Economic Journal: Microeconomics 3*, 60–88.

Bas, M., T. Mayer, and M. Thoenig (2015). From micro to macro: demand, supply, and heterogeneity in the trade elasticity. *CEPR Discussion Paper Series* (10637).

Behrens, K., Y. Kanemoto, and Y. Murata (2014). New trade models, elusive welfare gains. *CEPR Discussion Paper Series* (10255).

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

Caliendo, L., R. Feenstra, J. Romalis, and A. Taylor (2015). Tariff reductions, entry, and welfare: Theory and evidence from the last two decades. *unpublished manuscript*.

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

Cardell, S. N. (1997). Variance components structures for the extreme-value and logistic distributions with application to models of heterogeneity. *Economic Theory 13*, 185–213.

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

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

Costinot, A. and A. Rodriguez-Clare (2014). Trade theory with numbers: Quantifying the consequences of globalization. In G. Gopinath, E. Helpman, and K. Rogoff (Eds.), *Handbook of International Economics*, Volume 4, pp. 197–262. Amsterdam: Elsevier.

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

Di Giovanni, J., A. Levchenko, and R.Ranciére (2011). Power laws in firm size and openness to trade: Measurement and implications. *Journal of International Economics 85*(1), 42–52.

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

Eaton, J. and S. Kortum (2010). *Technology in the Global Economy: A Framework for Quantitative Analysis*. Unpublished manuscript.

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

Eaton, J., S. Kortum, and S. Sotelo (2013). International trade: Linking micro and macro. In Acemoglu, Arellano, and Dekel (Eds.), *Advances in Economics and Econometrics, Applied Economics: Tenth World Congress*, Volume 2. Cambridge: Cambridge University Press.

Feenstra, R., R. Inklaar, and M. Timmer (2013). The Next Generation of the Penn World Table. *NBER Working Paper Series* (19255).

Feenstra, R. C., M. Obstfeld, and K. N. Russ (2014). In Search of the Armington Elasticity. *NBER Working Paper series* (20063).

Gabaix, X. and R. Ibragimov (2011). Rank-1/2: A simple way to improve the OLS estimation of tail exponents. *Journal of Business & Economic Statistics 29*, 24–39.

Gaulier, G. and S. Zignago (2010). BACI: International trade database at the product level. *CEPII working papers* (23).

Hanson, G. H., N. Lind, and M.-A. Muendler (2015). The dynamics of comparative advantage. *NBER Working Paper Series* (21753).

Head, K. and T. Mayer (2013). What separates us? Sources of resistance to globalization. *Canadian Journal of Economics 46*(4), 1196–1231.

Head, K. and T. Mayer (2014). *Handbook of international economics*, Volume 4, Chapter Gravity Equations: workhorse, toolkit, and cookbook. Elsevier.

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

Head, K., T. Mayer, and M. Thoenig (2014). Welfare and trade without Pareto. *American Economic Review 104*(5), 310–316.

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

Hummels, D. (2001). Towards a geography of trade costs. *Unpublished manuscript, Purdue University*.

Hummels, D. and G. Schaur (2013). Time as a trade barrier. *American Economic Review 103*, 1–27.

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

Kratz, M. and S. Resnick (1996). The QQ-estimator and heavy tails. *Stochastic Models 12*(4), 699–724.

Kugler, M. and E. Verhoogen (2012). Prices, plant size, and product quality. *Review of Economic Studies 79*(1), 307–339.

Mayer, T., M. Melitz, and G. Ottaviano (2014). Market size, competition, and the product mix of exporters. *American Economic Review 104*(2), 495–536.

Melitz, M. and S. Redding (2015). New trade models, new welfare implications. *American Economic Review 105*(3), 1105–1146.

Ossa, R. (2015). Why trade matters after all. *unpublished manuscript*.

Perla, J., C. Tonetti, and M. E. Waugh (2015, January). Equilibrium technology diffusion, trade, and growth. *NBER Working Paper series* (20881).

Sampson, T. (2015). Dynamic selection: An idea flows theory of entry, trade, and growth. *Quarterly Journal of Economics (forthcoming)*.

Simonovska, I. and M. Waugh (2014a). The elasticity of trade: Estimates and evidence. *Journal of International Economics 92*(1), 34–50.

Simonovska, I. and M. Waugh (2014b). Trade models, trade elasticities, and the gains from trade. *NBER Working Paper Series* (20495).

Appendices

A The framework

A.1 Supply set-up

The world contains N countries that each have an endowment L_i of labor. Output is produced using labor, which is perfectly mobile across sectors and immobile across countries. Production technology is linear in labor, with unit labor cost w_i.

There is a unit continuum $k \in [0,1]$ of products ('sectors') in each country. Output can be produced using any of the production techniques for sector k available in country i. Production techniques vary in efficiency z. Techniques are drawn from a common distribution but these draws are independent in each sector. Varieties of product k produced within the same country are taken to be perfect substitutes.[39] Constant returns to scale and within-sectoral product homogeneity imply that the best available technique is used in production of each sector within the country. Nonetheless, techniques may differ across sectors within the country and across countries for any given sector.

In each sector technology improvement follows the Poisson process described in Eaton, Kortum, and Sotelo (2012) whereby at each point in time the number of techniques available for producing output with efficiency $Z > z$ follows a Poisson distribution with parameter $\lambda_i(t) = T_i(t)z^{-\theta}$ (Eaton et al. (2013)). This parameter is increasing in $T_i(t)$ which denotes the stock of technology accumulated by time t and in $1/\theta$ which denotes the extent of dispersion in technology draws.[40] This parameter maps fundamental exporter ability into the number of goods that can be produced with efficiency higher than any given threshold z.

Given a Poisson process for the arrival of ideas and a stock of technology T_i, the probability of no technique with efficiency $Z > z$ arriving in a unit interval is given by the Poisson density for $X = 0$, where X is the number of draws with efficiency higher than z:

$$\Pr[X=0] = \frac{(\lambda_i)^0 \exp\{-\lambda_i\}}{0!} = \exp\left\{-T_i z^{-\theta}\right\} \tag{37}$$

The probability that at least one technique of higher efficiency occurs is given by:

$$\Pr[Z>z] \quad = \quad 1 - \Pr[X=0] = 1 - \exp\left\{-T_i z^{-\theta}\right\} \tag{38}$$

[39] For consistency with the Armington set-up, varieties of k produced in different countries are perceived to be imperfect substitutes.

[40] θ is the shape parameter of the Pareto distribution from which efficiency is drawn. A lower θ corresponds to a distribution with a fatter tail, e.g. a higher probability of getting a high draw (Eaton and Kortum (2010)).

As the process of technology upgrading takes place independently in each sector, this probability distribution also characterizes the distribution of the best production techniques in the unit continuum of sectors. The structure of production thus replicates Eaton and Kortum (2002) wherein techniques effectively used in production are distributed Fréchet.

A.2 The price of the country-specific composite good

A.2.1 Derivation of the price index

The Armington assumption of product differentiation by place of origin means that all products survive and are exported to the world market.[41] Denoting F_i the distribution of prices across sectors within a country, we obtain the lower tier price index by aggregating across the distribution of realized prices:

$$P_i(p) = \left\{ \int_0^\infty p^{1-\sigma'} \mathrm{d}F_i(p) \right\}^{\frac{1}{1-\sigma'}} \tag{39}$$

Efficiency is the realization of the random variable Z with independent draws for each sector from the Fréchet distribution with parameter λ. The unit cost of producing k in i corresponds to the realization of the random variable $\Upsilon = w_i/Z$. Consequently, the number of techniques that allow producing output with cost lower than some threshold υ is distributed Poisson with parameter $\lambda_i = T_i(w_i/\upsilon)^{-\theta}$ (using $z = w_i/\upsilon$) where the time subscript is suppressed given our focus on expenditure allocation in cross-section. Applying (37) the probability of no technique allowing production with cost less than υ arriving in a unit interval is given by $\exp\{-\lambda_i\}$. Applying (38) the probability of at least one lower cost draw arriving is given by $1-\exp\{-\lambda_i\}$. The distribution of lowest costs is Weibull with parameter λ_i (Eaton and Kortum (2002)):

$$F(\upsilon) \quad = \quad \Pr[\Upsilon \leq \upsilon] = 1 - \exp\left\{-T_i {w_i}^{-\theta} \upsilon^{\theta}\right\} \tag{40}$$

and the corresponding pdf is:

$$f(\upsilon) \quad = \quad T_i w_i^{-\theta} \theta \upsilon^{\theta-1} \exp\left\{-T_i {w_i}^{-\theta} \upsilon^{\theta}\right\} \tag{41}$$

Perfect competition within each sector implies that the distribution of realized prices is directly given by the distribution of least costs:

$$P_i(p)^{1-\sigma'} \quad = \quad P_i(\upsilon)^{1-\sigma'} = \int_0^\infty \upsilon^{1-\sigma'} f(\upsilon)\, \mathrm{d}\upsilon \tag{42}$$

[41] Recall that in Eaton and Kortum (2002) countries supply homogeneous sectoral products, and the consumer only cares about the combination of least-cost products in the unit continuum.

Hence, we can apply Lemma 2 in Eaton and Kortum (2010) to compute the price of the country-specific composite good. Assuming that parameter restrictions $1 < \sigma \leq \sigma' < \theta + 1$ hold and denoting by $\gamma = (\theta + 1 - \sigma')/\theta$ the parameter of the Gamma function, we get:[42]

$$P_i = \left\{ T_i {w_i}^{-\theta} \right\}^{-1/\theta} \{\Gamma(\gamma)\}^{1/1-\sigma'} \quad (43)$$

To simplify notation, we denote $\chi = \{\Gamma(\gamma)\}^{1/(1-\sigma')}$ the source invariant scalar. Further, we use the definition of the mean of the Fréchet distribution $T_i^{1/\theta}\Gamma(1-1/\theta)$ to define the scale parameter of the Fréchet $\tilde{z}_i = T_i^{1/\theta}$. The price of the country-specific composite good is given by the unit cost of production, adjusted by the scale parameter of the Fréchet that is a sufficient statistic of expected technological ability of the country (Costinot et al. (2012)):

$$P_i = \chi \frac{w_i}{\tilde{z}_i} \quad (44)$$

A.2.2 Consistent aggregation of micro-level prices when all products are observed

Whenever the full set of products is reported as traded, the estimate of the lower tier demand elasticity is not needed for consistent price aggregation within exporter-specific product sets. In the absence of truncation, the distribution of market shares can be used for consistent price aggregation. This approach delivers an approximation to the ideal price index of the country-specific composite good up to a scalar that is invariant across exporters. Given that this approximation preserves information on ideal price indices in relative terms, we can obtain a consistent estimate of the upper tier demand elasticity without prior knowledge of the lower tier demand elasticity.

From (43) we know that σ' enters the expression of the ideal price index through the argument $\gamma = (\theta + 1 - \sigma')/\theta$ of the Gamma function and through the exponent to which $\Gamma(\gamma)$ is raised. Hence, even if σ' is unknown, the true theoretical price index $P_i(\sigma')$ can be recovered by appropriate rescaling of any other price index $P_i(\tilde{\sigma}')$:

$$P_i(\sigma') = \frac{\kappa(\sigma')}{\kappa(\tilde{\sigma}')} P_i(\tilde{\sigma}') \quad (45)$$

where $\tilde{\sigma}' \in \,]1, \theta+1[$ and $\kappa(\sigma') = \{\Gamma(\gamma)\}^{1/(1-\sigma')}$. The scalar $\delta(\tilde{\sigma}') = \kappa(\sigma')/\kappa(\tilde{\sigma}')$ is invariant across exporters. Hence, information on relative prices of exporter-specific composite goods is preserved if price indices are computed using a value of σ' that differs from the true value.

[42] See Eaton and Kortum (2010): plug (41) into (42); use the definition of λ to write $d\lambda = T_i {w_i}^{-\theta} \theta \upsilon^{\theta-1} d\upsilon$ and $(\lambda / T_i {w_i}^{-\theta})^{(1-\sigma')/\theta} = \upsilon^{1-\sigma'}$; change the variable of integration and rearrange (42) to get $P_i^{1-\sigma'} = \left\{ T_i {w_i}^{-\theta} \right\}^{-(1-\sigma')/\theta} \int_0^\infty \lambda^{(1-\sigma')/\theta} \exp\{-\lambda\}\, \mathrm{d}\lambda$. The latter integral is equal to $\Gamma[1 + (1-\sigma')/\theta]$.

This property is used to prove that the price index $\tilde{P}_i$ obtained by using observed expenditure weights to aggregate sectoral prices instead of the true value of σ' can also be used to back out the true theoretical price index $P_i(\sigma')$ by appropriate rescaling.

To see how the approximation works, we rearrange the lower tier demand equation to express the landed price of the non-truncated bundle in terms of any product k:

$$P_i = P_{ik}\left[\frac{X_{ik}}{X_i}\right]^{1/(\sigma'-1)} \tag{46}$$

Denote the total number of categories in the product classification by N and work with the discrete version of the ideal price index by summing (46) across these N categories:

$$P_i = \frac{1}{N}\sum_{k=1}^{N} P_{ik}\left[\frac{X_{ik}}{X_i}\right]^{1/(\sigma'-1)} \tag{47}$$

The approximation to the ideal price index is obtained by replacing the theoretical weights $[X_{ik}/X_i]^{1/(\sigma'-1)}$ with observed expenditure weights:

$$\tilde{P}_i = \sum_{k=1}^{N} P_{ik}\left[\frac{X_{ik}}{X_i}\right] \tag{48}$$

The ordering of bundle prices must be unchanged since (48) is a monotonic transformation of (47). The approximation also preserves the cardinal ranking. This is proved by showing that relative bundle prices are unchanged when expenditure weights generated by the true but unknown lower tier elasticity are used together with the assumption $\sigma' = 2$ in aggregation.

Denote sectoral expenditure associated with any given value of σ' by $X_{ik}(\sigma')$ and use the associated sectoral demand equation $X_{ik}(\sigma') = [P_{ik}/P_i(\sigma')]^{(1-\sigma')}X_i$ to rewrite observed sectoral weights in terms of weights associated with $\sigma' = 2$.

$$X_{ik}(\sigma') = \frac{[P_{ik}/P_i(\sigma')]^{(1-\sigma')}}{[P_{ik}/P_i(2)]^{(1-2)}} X_{ik}(2) \tag{49}$$

Use the result that $P_i(\sigma') = \delta(2)P_i(2)$ and simplify (49) to get:

$$X_{ik}(\sigma') = \left[\frac{P_{ik}}{P_i(\sigma')}\right]^{(2-\sigma')} \delta(2)X_{ik}(2) \tag{50}$$

Plug (50) in (48) to get:

$$\tilde{P}_i = \sum_{k=1}^{N} P_{ik}\left[\frac{P_{ik}}{P_i(\sigma')}\right]^{(2-\sigma')} \delta(2)\left[\frac{X_{ik}(2)}{X_i}\right] \tag{51}$$

Use the fact that $[X_{ik}(2)/X_i] = [P_{ik}/P_i(2)]^{(-1)}$ and $P_i(\sigma') = \delta(2)P_i(2)$ to rewrite (51):

$$\tilde{P}_i = \sum_{k=1}^{N} P_{ik}\left[\frac{P_{ik}}{P_i(\sigma')}\right]^{(2-\sigma')} \delta(2)\left[\frac{P_{ik}}{\delta(2)^{-1}P_i(\sigma')}\right]^{-1} \tag{52}$$

Simplifying (52) leaves two components: $\tilde{P}_i = P_i(\sigma')^{\sigma'-1} \sum_{k=1}^{N} P_{ik}^{(2-\sigma')}$. The second component is linked to the price index $P_i(\tilde{\sigma}')$, with $\tilde{\sigma}' = \sigma' - 1$. This component can be written $\sum_{k=1}^{N} P_{ik}^{(2-\sigma')} = \sum_{k=1}^{N} P_{ik}^{(1-\tilde{\sigma}')} = P_i(\tilde{\sigma}')^{1-\tilde{\sigma}'}$. Combining all of the above gives:

$$\tilde{P}_i \quad = \quad P_i(\sigma')^{\sigma'-1} P_i(\tilde{\sigma}')^{1-\tilde{\sigma}'} \tag{53}$$

The latter is linked to the true price index: $P_i(\sigma') = \delta(\tilde{\sigma}')P_i(\tilde{\sigma}')$ whereby:

$$\tilde{P}_i \quad = \quad \delta(\tilde{\sigma}')^{\sigma'-2} P_i(\sigma') \tag{54}$$

The approximation $\tilde{P}_i$ is obtained by appropriate rescaling of the true underlying price index. The scalar $\delta(\tilde{\sigma}')^{\sigma'-2}$ is invariant across exporters. The approximation preserves information on relative prices. Hence, knowledge of the lower tier elasticity is not a prerequisite for consistent price aggregation.

A.3 The price of the composite good when a subset of products is observed

Our model predicts that higher trade costs leave the set of goods unaffected whereby only the intensive margin of trade is operational. However, only a subset of products is reported as traded in the dataset. We reconcile the key prediction of our model with the prevalence of zeros in the data by showing that in our framework products in which the unit labor requirement is sufficiently high carry marginal weight in expenditure on the exporter-specific composite good. Such low technology draws are associated to small trade flows that may be omitted from trade statistics because of the existence of statistical thresholds for reporting trade.[43]

A.3.1 Derivation of the price index

Assume there exists a statistical threshold $\bar{X}$ common to all countries such that the nominal value of sectoral bilateral trade is registered iff it is at least equal to this threshold. Define $\bar{\upsilon}$ the maximal production cost associated with the smallest observed nominal value and apply (25) to the price of each sectoral good: $X_i(k) \geq \bar{X}$ implies $P_i(k) \leq \bar{\upsilon}$. The fraction of high cost draws determines overall product variety observed for each exporter on the world market. Bilateral trade frictions together with market-specific characteristics determine bilateral variation in product variety for every exporter.

[43] The threshold is 1000 current USD for the UN COMTRADE database from which the BACI dataset is built.

The cost threshold $\bar{\upsilon}$ is incorporated in the lower tier price index to obtain the price of the truncated product set $\bar{P}_{ij}$:

$$\bar{P}_{ij}(p) \quad = \quad \bar{P}_{ij}(\upsilon) = \left\{ \int_0^{\bar{\upsilon}} \upsilon^{1-\sigma'} f(\upsilon)\, \mathrm{d}\upsilon \right\}^{\frac{1}{1-\sigma'}} \tag{55}$$

Following Eaton and Kortum (2010), we derive the price index of the truncated set by rewriting (55) as the product of two terms: the expected number of bilateral draws below the threshold and the expected cost of such draws. The number of techniques that allow production with cost less than $\bar{\upsilon}$ is given by $\mu_i(\bar{\upsilon})$ (see A.1). Augmenting unit cost w_i with bilateral trade frictions τ_{ij} defines this statistic at the bilateral level. The expected bilateral number of draws is:

$$\mu_{ij}(\bar{\upsilon}) \quad = \quad \left(w_i \tau_{ij}/\tilde{z}_i\right)^{-\theta} \bar{\upsilon}^{\theta} \tag{56}$$

The expected cost of such draws is obtained by integrating the conditional density function over effectively observed cost draws. In Eaton and Kortum (2010) techniques are drawn from a Pareto distribution with parameter θ. Hence, the distribution of costs conditional on the cost threshold $\bar{\upsilon}$ is $F_c(\upsilon) = \Pr(\Upsilon \leq \upsilon | \Upsilon \leq \bar{\upsilon}) = (\upsilon/\bar{\upsilon})^{\theta}$. The corresponding conditional density function is $f_c(\upsilon) = \theta \upsilon^{\theta-1} \bar{\upsilon}^{-\theta}$. The lower tier price index for the truncated product set is:

$$\bar{P}_{ij}(\bar{\upsilon}) = \left\{ E\left[P_{ij}(k)^{1-\sigma'} | P_{ij}(k) \leq \bar{\upsilon} \right] \right\}^{\frac{1}{1-\sigma'}} = \left\{ \underbrace{\mu_{ij}(\bar{\upsilon})}_{\text{nbr draws}} \underbrace{\int_0^{\bar{\upsilon}} \upsilon^{1-\sigma'} f_c(\upsilon)\, \mathrm{d}\upsilon}_{\text{expected cost}} \right\}^{\frac{1}{1-\sigma'}} \tag{57}$$

Replacing $\mu_{ij}(\bar{\upsilon})$ and $f_c(\upsilon)$ and solving for the integral defines the price of the truncated product set as a function of exporter characteristics, trade frictions, and the cost threshold:

$$\bar{P}_{ij}(\bar{\upsilon}) \quad = \quad \left\{ \frac{\theta}{\theta - \sigma' + 1} \left(\frac{w_i \tau_{ij}}{\tilde{z}_i} \right)^{-\theta} \bar{\upsilon}^{\theta - \sigma' + 1} \right\}^{1/(1-\sigma')} \tag{58}$$

The next step is to derive the cost threshold. Using the two-tier structure of expenditure allocation, the landed value of sectoral trade that is effectively observed at the bilateral level is:

$$X_{ij}(k)\big|_{X_{ij}(k) \geq \bar{X}} \quad = \quad \left[\frac{P_{ij}(k)}{P_{ij}} \right]^{1-\sigma'} \left[\frac{P_{ij}}{\Phi_j} \right]^{1-\sigma} Y_j \tag{59}$$

The expression of the cost threshold is obtained by solving for the upper bound of the observed landed sectoral price in (59):

$$P_{ij}(k) \; \leq \; \left[\frac{Y_j}{\bar{X}} \Phi_j^{\sigma-1} \right]^{1/(\sigma'-1)} P_{ij}^{\;\frac{\sigma'-\sigma}{\sigma'-1}} \quad = \quad \bar{\upsilon} \tag{60}$$

Plug (5) into (60) to visualize the four components of the cost threshold:

$$\bar{\upsilon} \quad = \quad \left\{ \left[\chi^{\sigma'-\sigma} \bar{X}^{-1} \right] \left[Y_j \Phi_j^{\sigma-1} \right] \left[\frac{w_i}{\tilde{z}_i} \right]^{\sigma'-\sigma} \tau_{ij}^{\sigma'-\sigma} \right\}^{\frac{1}{\sigma'-1}} \tag{61}$$

Define $\psi = \frac{\theta}{\theta-\sigma'+1} \left[\chi^{\sigma'-\sigma} \bar{X}^{-1} \right]^{\frac{\theta-\sigma'+1}{\sigma'-1}}$ and $\alpha = (\sigma'-\sigma)(\theta-\sigma'+1)/(\sigma'-1)$. Plugging (61) into (58) gives the landed price of the truncated product set:

$$\bar{P}_{ij} \quad = \quad \left\{ \psi \left[Y_j \Phi_j^{\sigma-1} \right]^{\frac{\theta-\sigma'+1}{\sigma'-1}} \left(\frac{w_i \tau_{ij}}{\tilde{z}_i} \right)^{-(\theta-\alpha)} \right\}^{\frac{1}{1-\sigma'}} \tag{62}$$

Statistical thresholds for trade registration lead to the exclusion of prices associated to high unit labor requirements from the computation of the price index for the country-specific composite good. Parameter restrictions are $\theta+1 > \sigma' \geq \sigma > 1$ whereby $\theta > \alpha \geq 0$. Taking the derivative wrt σ', it is immediate that the exponent is monotonically decreasing from $\theta/(\sigma-1)$ to 1 as σ' increases from σ to $\theta+1$. Hence, truncation enhances the sensitivity of aggregate prices to exporter characteristics and bilateral trade frictions whenever $\theta \gg (\sigma'-1)$.[44]

A.3.2 Consistent aggregation of micro-level prices when a subset of products is observed

We now show that knowledge of the lower tier demand elasticity is a prerequisite for consistent price aggregation in the presence of truncation. Specifically, the severity of truncation is a function of exporter ability. Consequently, the approximation that consists in computing the price index with help of observed market shares instead of the lower tier elasticity fails to preserve information on relative prices of exporter-specific composite goods.

Throughout this section we condition prices and expenditure on market-specific characteristics and bilateral trade frictions so as to focus on the impact of differences in exporter ability. First, we show that the price of the truncated composite good exceeds the price of the non-truncated composite good for any exporter. Consider the partial derivative of the ideal price index for the truncated bundle wrt the cost threshold $\bar{\upsilon}$. The sign of this derivative is determined by the sign of the outer exponent $1/(1-\sigma')$:

$$\frac{\partial \bar{P}_i}{\partial \bar{\upsilon}} \quad = \quad \frac{\partial}{\partial \bar{\upsilon}} \left\{ \left[\int_0^{\bar{\upsilon}} p_i^{1-\sigma'} f(p_i) \mathrm{dp_i} \right]^{1/(1-\sigma')} \right\} < 0 \tag{63}$$

[44] Denote $\theta = \sigma'-1+\Delta$ and rewrite the exponent as $1+\Delta(\sigma-1)/(\sigma'-1)^2 > 1$. Direct comparison of the exponent in (62) and (5) gives $(\theta-\alpha)/(\sigma'-1) \geq 1$.

The solution to the integral in (63) is given by (58). Using (58) it is immediate that the second derivative is positive and that the first derivative tends to 0 in the limit:

$$\lim_{\bar{\upsilon}\to\infty} \partial \bar{P}_i / \partial \bar{\upsilon} = 0$$

The price of the truncated composite good is monotonically increasing in the extent of truncation. Moreover, the gap between the truncated and the non-truncated price index is monotonically decreasing in exporter ability. To see this, form the ratio of the two prices using (5) and (62) and consider the partial derivative with respect to $(\tilde{z}_i/w_i)$:

$$\frac{\partial(\bar{P}_i/P_i)}{\partial(\tilde{z}_i/w_i)} = \frac{\partial}{\partial(\tilde{z}_i/w_i)} \left\{ \left[\frac{\psi}{\Gamma(\gamma)}\right]^{\frac{1}{(1-\sigma')}} \left[\frac{\tilde{z}_i}{w_i}\right]^{\frac{(\theta-\sigma'+1)-\alpha}{(\sigma'-1)}} \right\} \tag{64}$$

The exponent on $(\tilde{z}_i/w_i)$ can be written $(1-\sigma)(\theta-\sigma'+1)/(\sigma'-1)^2 < 0$ establishing that the derivative is negative. For any two exporters i' and i such that $(\tilde{z}_{i'}/w_{i'})/(\tilde{z}_i/w_i) < 1$, the relative price of the truncated composite good exceeds the relative price of the non-truncated composite good for the less able exporter: $\bar{P}_{i'}/\bar{P}_i > P_{i'}/P_i$. Hence, truncation amplifies price dispersion relatively to structural price dispersion among exporters.

Next, we consider the approximation that uses observed expenditure weights to compute the price of the truncated composite good:

$$\tilde{\bar{P}}_i(\sigma', \bar{\upsilon}) = \sum_{\{k|P_{ik}\leq\bar{\upsilon}\}} P_{ik} \left[\frac{X_{ik}(\sigma')}{\bar{X}_i(\bar{\upsilon})}\right] \tag{65}$$

To relate expenditure weights to underlying price indices, we rearrange the sectoral demand equation and sum across all prices below the cost threshold:

$$\sum_{\{k|P_{ik}\leq\bar{\upsilon}\}} P_{ik}^{1-\sigma'} = \sum_{\{k|P_{ik}\leq\bar{\upsilon}\}} \frac{X_{ik}(\sigma')}{X_i} P_i(\sigma')^{1-\sigma'} = \frac{\bar{X}_i(\bar{\upsilon})}{X_i} P_i(\sigma')^{1-\sigma'} \tag{66}$$

Truncated expenditure $\bar{X}_i(\bar{\upsilon})$ is written as a function of truncated and non-truncated prices:

$$\bar{X}_i(\bar{\upsilon}) = \bar{P}_i(\sigma')^{1-\sigma'} P_i(\sigma')^{\sigma'-1} X_i \tag{67}$$

Plugging (67) and (50) in (65) gives:

$$\tilde{\bar{P}}_i(\sigma', \bar{\upsilon}) = \sum_{\{k|P_{ik}\leq\bar{\upsilon}\}} P_{ik}^{3-\sigma'} P_i(\sigma')^{-1} \delta(2) \bar{P}_i(\sigma')^{\sigma'-1} \left[\frac{X_{ik}(2)}{X_i}\right] \tag{68}$$

To simplify (68), use the fact that $[X_{ik}(2)/X_i] = [P_{ik}/P_i(2)]^{(-1)}$ and $P_i(\sigma') = \delta(2)P_i(2)$:

$$\tilde{\bar{P}}_i(\sigma', \bar{\upsilon}) = \sum_{\{k|P_{ik}\leq\bar{\upsilon}\}} P_{ik}^{2-\sigma'} \bar{P}_i(\sigma')^{\sigma'-1} \tag{69}$$

Define $\tilde{\sigma}' = \sigma' - 1$ and use the fact that $\bar{P}_i(\tilde{\sigma}')^{1-\tilde{\sigma}'} = \sum_{\{k|P_{ik} \leq \bar{v}\}} P_{ik}^{1-\tilde{\sigma}'}$ to get:

$$\tilde{\bar{P}}_i(\sigma', \bar{v}) = \bar{P}_i(\tilde{\sigma}')^{1-(\sigma'-1)} \bar{P}_i(\sigma')^{\sigma'-1} = \left[\frac{\bar{P}_i(\sigma')}{\bar{P}_i(\tilde{\sigma}')}\right]^{\sigma'-1} \bar{P}_i(\tilde{\sigma}') \tag{70}$$

The approximation delivers the truncated price index for $(\sigma' - 1)$ instead of σ'. This index is rescaled by an exporter-specific ratio of truncated price indices $\bar{P}_i(\sigma')/\bar{P}_i(\tilde{\sigma}')$.

Using (62) and defining $\zeta(\sigma') = (\theta - \alpha(\sigma'))/(\sigma' - 1)$, the ratio is given by:

$$\frac{\bar{P}_i(\sigma')}{\bar{P}_i(\tilde{\sigma}')} = \left\{\frac{\psi(\sigma')^{1/(1-\sigma')}}{\psi(\tilde{\sigma}')^{1/(1-\tilde{\sigma}')}}\right\} \left\{\frac{\tilde{z}_i}{w_i}\right\}^{\zeta(\tilde{\sigma}')-\zeta(\sigma')} \tag{71}$$

The first term on the RHS is a positive scalar invariant across exporters. The sign of the exponent in the second term determines the relationship of this ratio with exporter ability $(\tilde{z}_i/w_i)$. Using the definition $\alpha(\sigma') = (\sigma' - \sigma)(\theta - \sigma' + 1)/(\sigma' - 1)$, the derivative of $\zeta(\sigma')$ is:

$$\frac{d\zeta(\sigma')}{d\sigma'} = -\frac{(\sigma'-1)\left[2(\theta-\sigma')+\sigma+1\right]+2(\sigma'-\sigma)(\theta-\sigma'+1)}{(\sigma'-1)^3} < 0 \tag{72}$$

As $\tilde{\sigma}' = \sigma' - 1 < \sigma'$, (72) establishes that the exponent is positive: $(\zeta(\tilde{\sigma}') - \zeta(\sigma')) > 0$. Hence, the ratio is increasing in exporter ability: $\partial(\bar{P}_i(\sigma')/\bar{P}_i(\tilde{\sigma}'))/\partial(\tilde{z}_i/w_i) > 0$. Thus, the approximation underestimates price dispersion relatively to structural price dispersion among exporters. We conclude that the neutrality of the approximation with respect to relative prices is not preserved under truncation. Consequently, knowledge of the lower tier elasticity is a necessary prerequisite for obtaining an unbiased estimate of the upper-tier elasticity.

B The dataset: variety and market coverage

The full sample contains 213 exporters operating on 213 markets that make up 35,543 pairs active at least once. We reduce this sample to 92 exporters operating on 211 markets using the following criteria. First, each exporter must be active on at least 5% of all markets and deliver at least 5% of world variety to the set of markets on which she is active. Further, the exporter must deliver at least 1% of world variety to 10 or more individual markets. The resulting reduced sample that we use in the estimation covers between 95 and 98% of total trade between 1995 and 2011 and contains 10,443 pairs active at least once.[45]

Fig.22 illustrates that the 92 exporters in the reduced sample are the best performing exporters in terms of market coverage. We see in the left pane of the figure that 50% of exporters

45 This sample of pairs is not balanced. The corresponding balanced sample comprises 5,589 pairs and covers between 90 and 97% of total trade.

Figure 22: Market coverage in BACI

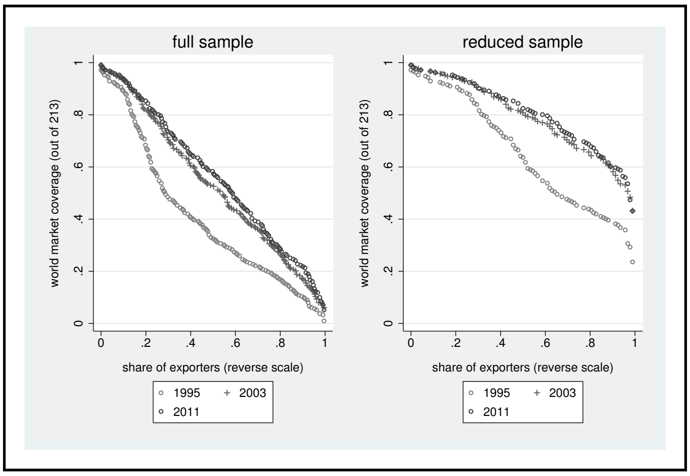

Figure 23: Variety coverage in BACI

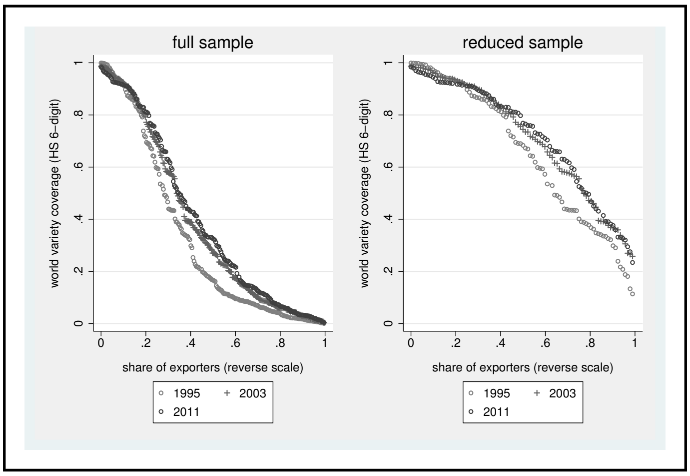

in the full sample cover fewer than 40% of all world markets in 1995. We see in the right pane of the figure that almost all exporters in the reduced sample cover at least 40% of all markets in 1995. Further, market coverage improves more strongly over time in the reduced sample. In particular, by 2011, almost all exporters in the reduced sample cover 60% of world markets while only about half of them did in 1995.

Fig.23 illustrates a similar pattern in terms of variety coverage. We see in the left pane of the figure that about 2/3 of all exporters in the full sample cover less than 40% of total world variety in 1995. We see in the right pane of the figure that 2/3 of exporters in the reduced sample cover more than 40% of world variety in 1995. Further, we see through the rightward shift of the distribution that variety coverage improves over time. Most of this increase takes place between the mid-1990s and the mid-2000s.

Figure 24: Product coverage for USA, Germany, China, and Mexico in 1995 and 2011

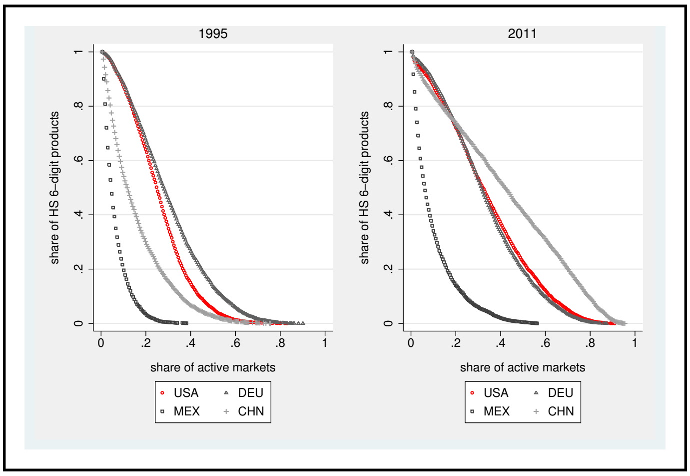

Further, we document that truncation remains prevalent in the reduced sample. For each exporter we compute the fraction of HS 6-digit products that reach at least $1, 2, \ldots, N$ markets in a given year. We scale this fraction by the total number of markets active in that year and obtain for each exporter the distribution of observed product variety across the set of world markets. Pooling information on the distribution of variety coverage for these 92 exporters in 1995-2011, we estimate the rate of decay in observed product variety associated with increasing the share of markets to which the product set is delivered. We find that the average rate of decay is 1.25, e.g. a 10% increase in the number of markets to which the set of goods is delivered is associated to a 25% reduction in observed product variety. On average, these 92 exporters improved their ability to maintain product variety on the set of world markets. In particular, the average rate of decay in product variety decreased from 1.30 in 1995 to 1.18 in 2011.

Fig.24 illustrates this pattern for a subset of exporters (USA, Germany, China, and Mexico) in 1995 and 2011. We see an improvement of variety coverage over time, in particular for the emerging economies. Nevertheless, even the best performing exporter in terms of product coverage, e.g. China in 2011, delivers less than half of its product set to 50% of world markets.

To sum up, even in the reduced sample only a fraction of the exporter-specific product set is delivered to most markets. In terms of our model, this finding means that most trade flows are too small to be reported in trade statistics.

C The Ricardian prediction of the model

Our model predicts that the ranking of relative sectoral exports for any pair of exporters maps into the pattern of structural comparative advantage. In this section we investigate to what extent this prediction is borne out in the data.

C.1 Relative sectoral demand for products of different origin

We rewrite the sectoral demand equation (26) for product k delivered by exporter i to market j in relative terms for any two exporters $\{i, i'\}$:

$$\frac{X_{ij}(k)}{X_{i'j}(k)} = \left[\frac{w_i \tau_{ij}/\tilde{z}_i}{w_{i'}\tau_{i'j}/\tilde{z}_{i'}}\right]^{-(\sigma-1)} \left[\frac{z_i(k)/\tilde{z}_i)}{z_{i'}(k)/\tilde{z}_{i'}}\right]^{\sigma'-1} \tag{73}$$

We use (73) to rewrite relative sectoral exports for the pair $\{i, i'\}$ relatively to some benchmark sector $\bar{k}$. Denoting $x_{ii'} = X_{i'}(\bar{k})/X_i(\bar{k})$ and $z_{ii'} = \left[z_{i'}(\bar{k})/z_i(\bar{k})\right]^{\sigma'-1}$ and ordering sectors according to decreasing ability in i, the ranking of relative exports is predicted to be stable across active export markets:

$$\begin{aligned} & x_{ii'}\frac{X_i(1)}{X_{i'}(1)} > \ldots > x_{ii'}\frac{X_i(k)}{X_{i'}(k)} > \ldots > x_{ii'}\frac{X_i(K)}{X_{i'}(K)} \\ \Leftrightarrow \; & z_{ii'}\left(\frac{z_i(1)}{z_{i'}(1)}\right)^{\sigma'-1} > \ldots > z_{ii'}\left(\frac{z_i(k)}{z_{i'}(k)}\right)^{\sigma'-1} > \ldots > z_{ii'}\left(\frac{z_i(K)}{z_{i'}(K)}\right)^{\sigma'-1} \end{aligned} \tag{74}$$

A qualitatively similar prediction is obtained in the multisectoral Ricardian model of Costinot et al. (2012) but the mechanism that generates rankings' stability is different. In Costinot et al. (2012), stability is obtained because the country is relatively more likely to be the lowest cost supplier across the full set of sectoral varieties if her expected technological ability in the sector is relatively high. In our model, stability is obtained because the best domestic technology is used to produce each product, and products of different origin are observed in each market.

C.2 Stability of the ranking of relative sectoral exports in the data

In practice, most exporters deliver a subset of goods to a subset of markets (Appendix B). The good is less likely to be observed in any market if its technology draw is in the left tail of the exporter-specific distribution (Appendix A). Hence, the ratios on the two extremes of (74) are likely to drop out of the sample. Consequently, we expect to evaluate the stability of relative exports' rankings in the medium range of (74).

The set of products common to any pair of exporters is market-specific. To focus on a stable product set, we conduct the analysis on the common set of products that two exporters deliver to 5 markets on which their variety coverage is largest. We exclude exporter pairs that deliver fewer than 50 common products to these 5 markets. For each pair of exporters ii' we construct the ranking of relative sectoral exports on each market as in (74) and compute the Spearman rank correlation coefficient ($\rho_{ii'jj't}$) for each market combination $j \neq j'$. We obtain 10 observations for each pair of exporters and report the median ($\rho_{ii't}$).

Figure 25: Stability of local rankings (by number products in common set)

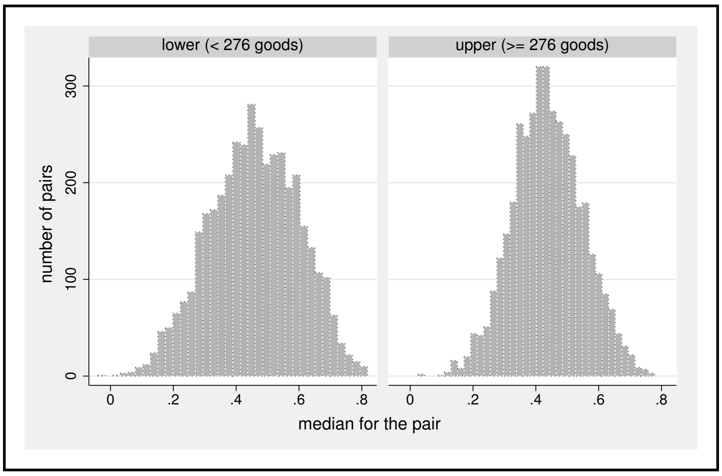

Fig.25 pools results for all years while splitting the sample of exporter pairs in half according to the number of products in the common set. The IQR for the distribution of Spearman rank correlation coefficients is .36-.56 (.36-.51) for exporter pairs with less (more) than 276 products. Thus, rankings' independence is rejected for the bulk of the sample but the correlation of local relative export rankings is strong for just about half of the pairs.

We provide a sensitivity check for this result by evaluating the stability of global and local relative export rankings. The working hypothesis is that the local ranking is more likely to deviate from the pattern of structural comparative advantage than the global ranking. This approach is in line with Mayer et al. (2014) who compute rank correlation coefficients for global and local rankings of product exports at the firm level.

We conduct the analysis on the set of common products that two exporters deliver to 10 markets on which their variety coverage is largest. We drop exporter pairs that deliver fewer than 50 comon products to these 10 markets. We compute Spearman rank correlation coefficients ($\rho_{ii'jt}$) for the global ranking of pair-specific relative sectoral exports with the local ranking on each of these markets $j = \{1, ..., 10\}$. We obtain 10 observations for each exporter pair and report the median ($\rho_{ii't}$).

Figure 26: Stability of global and local rankings (by number products in common set)

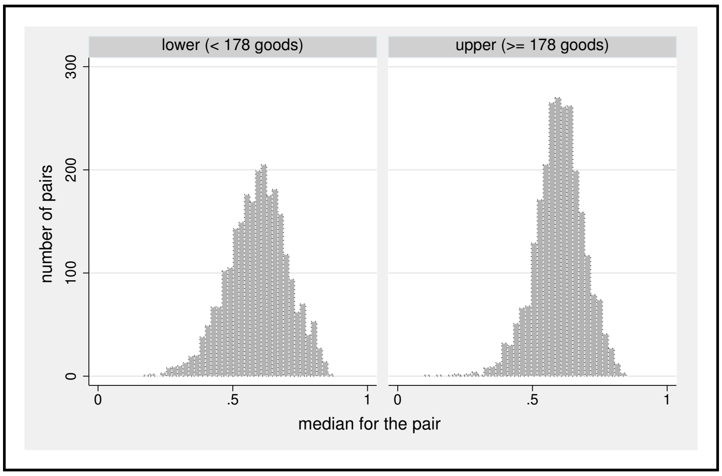

Figure 27: Stability of global and local rankings (by year and number of products)

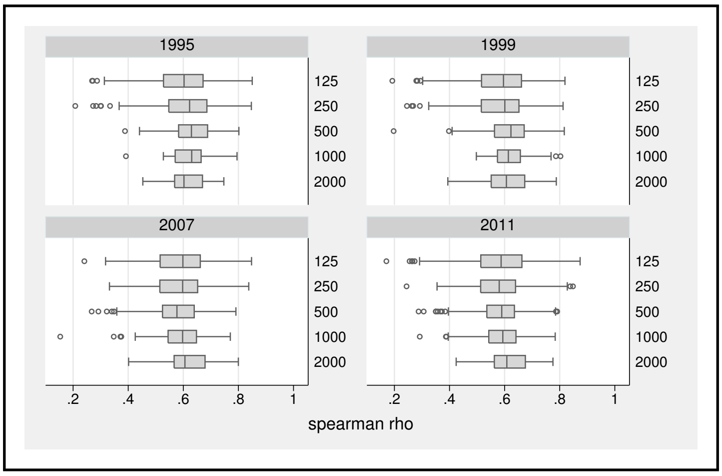

Fig.26 pools results for all years while splitting the sample in half according to the number of products in the common set. The distribution of Spearman ρ shifts rightward: the IQR is .52-.68 (.55-.66) for pairs with less (more) than 178 products in the common set. For most exporter pairs, the correlation of global and local rankings is strong. Fig.27 shows the annual distribution of Spearman ρ split by the number of products.[46] The distribution is stable over time and little sensitive to the number of products in the common set.

Quantitatively, our results are in line with Spearman ρ reported by Mayer et al. (2014) at the firm level. Qualitatively, our results are consistent with Hanson et al. (2015) who document the

[46] The numbers on the y-axis of each boxplot correspond to the upper bound of the included product set. Thus, the first (second,...) box corresponds to exporter pairs that share between 50 and 125 (125 and 250,...) goods.

stability of countries' export capabilities across markets and over time. We conclude that there is support for the Ricardian prediction of the model. Further, global relative export rankings better capture the pattern of comparative advantage than local relative export rankings.

D Robustness checks

D.1 The lower tier elasticity

We present the results of two robustness checks for the estimation of the lower tier elasticity.

We have shown in Appendix C that global export rankings better capture the pattern of structural comparative advantage than local export rankings. Consequently, in the first robustness check we use the full set of markets on which the exporter is active to construct the proxy of sectoral ability.

We denote d_t the constant, d_{jt} the market fixed effect, $d_t(k)$ the sectoral fixed effect, and $\varepsilon_{jt}(k)$ the error term. We retrieve the ranking of ability by running the following specification separately for each exporter in each year: $\ln(X_{ijt}(k)) = d_t + d_{jt} + d_t(k) + \varepsilon_{jt}(k)$. According to our model, the sectoral dummy captures the sectoral ability of exporter i relatively to some benchmark sector k': $d_t(k) = \hat{d}_{it}(k) = (\sigma_t' - 1)\ln(\hat{z}_{it}(k)/\hat{z}_{it}(k'))$. We normalize by the 'best' draw defined to be the most frequently observed product in the set of exported products.

We then proceed as in the core of the paper. In the first stage of the estimation we demean the data in the pair and sector dimensions and predict current prices with the first lag of the ability ranking. We include the frequency with which each product is observed on the world market as an additional control:

$$\ln\left[P_{ijt}(k)\right] = b_0 + b_{dt}\left[\hat{d}_{i,t-1}(k)\right] + b_{ft}\ln\left[f_{it}(k)\right] + b_t(l) + b_{ijt} + \varepsilon_{ijt}(k) \tag{75}$$

where $b_t(l)$ is the sectoral fixed effect at the HS 4-digit level and b_{ijt} is the pair fixed effect.

Fig.28 summarizes our results for b_{ft} in the left pane and b_{dt} in the right pane. We include first stage results obtained in the benchmark specification $(-c)$ for comparison purposes. The magnitude of the coefficients increases relatively to the benchmark specification but results are qualitatively unchanged.[47]

The equation that we estimate in the second stage is:

$$\ln\left[X_{ijt}(k)\right] = \beta_0 - \beta_{dt}\left[\hat{P}_{ijt}(k)\right] + \beta_{ft}\ln\left[f_{it}(k)\right] + \beta_t(l) + \beta_{ijt} + \varepsilon_{ijt}(k) \tag{76}$$

[47] The number of observations included in the regression increases from 1.5-1.8 million in 1996-1997 to 3.4 million by 2010-2011. Results are not sensitive to using the first or the second lag of the ability ranking.

Figure 28: First stage results: benchmark (-c) and robustness (all)

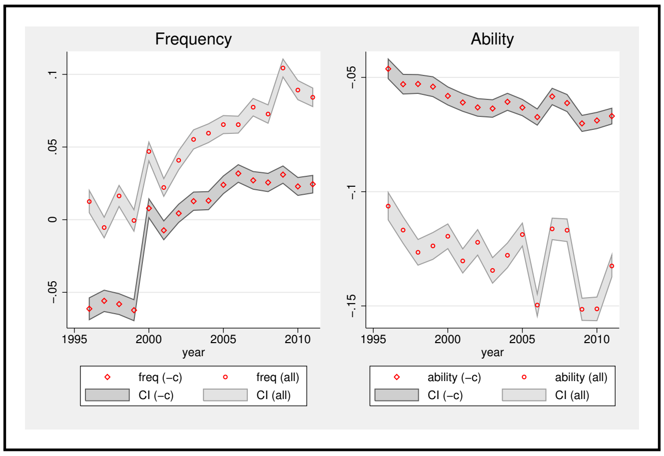

Figure 29: Second stage results: benchmark (-c) and robustness (all)

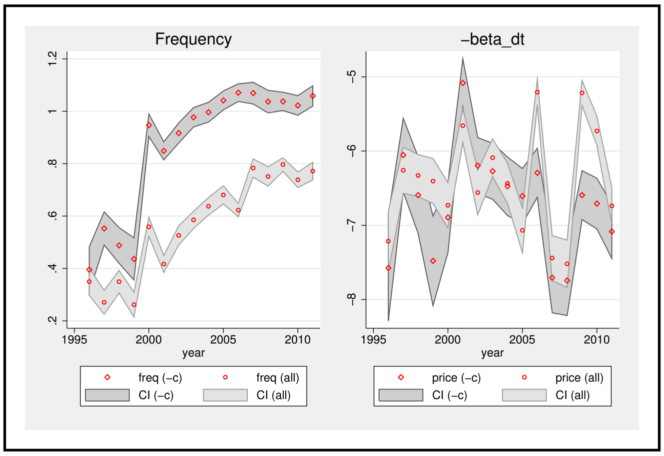

Fig.29 summarizes our results for β_{ft} in the left pane and for $-\beta_{dt}$ in the right pane. We include second stage results obtained in the benchmark specification ($-c$) for comparison purposes. The magnitude of the coefficient on the frequency variable is smaller in the robustness check. The estimate of the lower tier elasticity obtained in the robustness check is almost always contained within the bounds of the confidence interval of the benchmark specification.

In the second robustness check we focus on the set of observations that were used to obtain the benchmark estimate of the tail exponent $\theta_t/(\sigma_t' - 1)$ in section 3.2. This approach reduces the number of observations by two orders of magnitude because we restrict the sample to exporter-market pairs that have ≥ 1000 observations above 10^6 USD. Moreover, for each pair we only retain information on the 500 biggest trade flows because they correspond to our definition of the right tail of the pair-specific conditional sales distribution.[48] We implement the same two-stage approach as in the first robustness check.

Figure 30: First stage results for the second robustness check

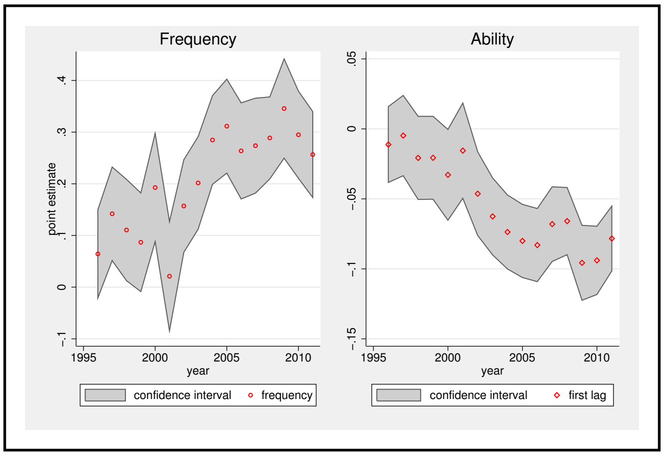

Fig.30 summarizes our results for the first stage of the estimation. The coefficient on the lagged ability ranking has the right sign in all years but it becomes a sufficiently strong predictor of current prices only in 2003-2011.[49] Consequently, we only report second stage results for 2003-2011 in Fig.31. Results of the first robustness check (*all*) are included for comparison. We see that the precision of the estimation is strongly reduced in the second robustness check. Nevertheless, the point estimate indicates that the magnitude of the lower tier elasticity in this reduced sample is comparable to estimates obtained in the much larger sample: $\hat{\sigma}'_t \in \{6, 8\}$.

D.2 The upper tier demand elasticity and the trade elasticity

We use the upper and lower bounds of our estimates of the tail exponent $\theta_t/(\sigma'_t - 1)$ to compute the upper tier elasticity σ_t and the trade elasticity ε_t. Fig.32 shows that our results for the structural upper tier demand elasticity $\hat{\sigma}_t$ are qualitatively unchanged when we use the upper and lower bounds of the tail exponent estimates to compute parameter magnitudes. Fig.33 documents a similar pattern for our results on the measured and structural trade elasticities.

[48] The number of observations varies between 13-18 thousand over 1996-2003. It thereafter increases from 25 to 40 thousand over 2004-2011.

[49] The Kleibergen Papp statistic is > 20 in 2003-2011, but < 10 in 1996-2001, and equal to 14 in 2002.

Figure 31: Second stage results of the two robustness checks: 2003-2011

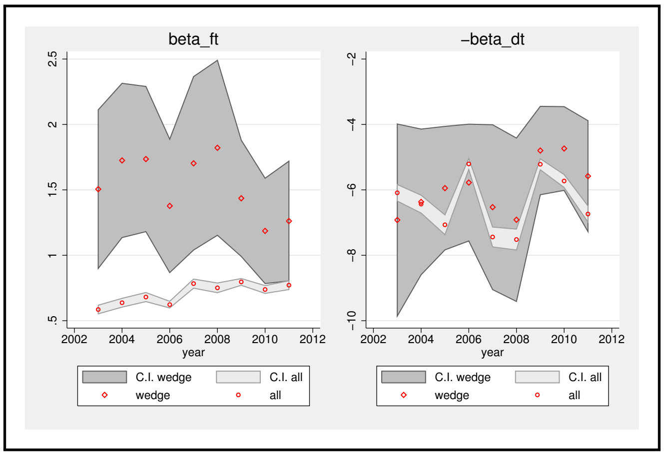

Figure 32: The upper tier demand elasticity ($\hat{\sigma}_t$)

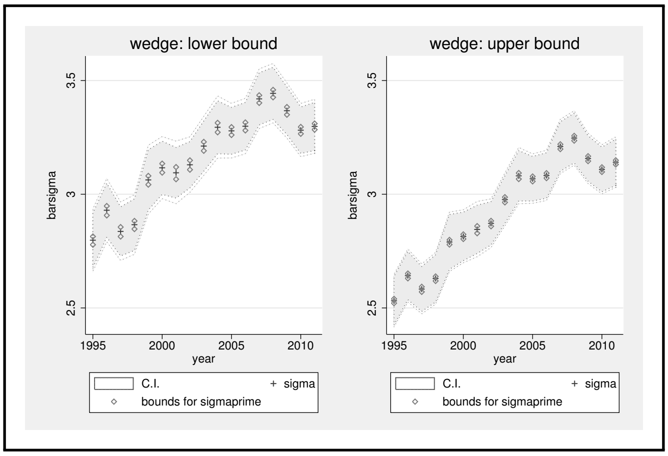

Figure 33: The measured and structural trade elasticity