Trade Elasticities*

Jean Imbs[†] Isabelle Mejean[‡]

Abstract

Conventional aggregate trade elasticity estimates hardly vary across countries. We introduce an aggregate elasticity that is implied by theory: It is the value that equates the welfare gains from trade as implied by one- and multi-sector versions of the model in Arkolakis et al. (2012). These estimates are predicated on sector-level values for trade elasticites, which we provide at 3-digit levels for 28 developed and developing countries. The values for this aggregate elasticity vary greatly across countries, and they do so because of countries' patterns of production, and because a given sector-level elasticity displays considerable cross-country heterogeneity.

JEL Classification Numbers: F32, F02, F15, F41

Keywords: Price Elasticity of Imports, Sectoral Estimates, Trade Performance, Heterogeneity.

1 Introduction

Estimates of the price elasticity of imports obtained form aggregate data are close to zero, and hardly display any differences across countries. For instance, Houthakker and Magee (1969) estimate price elasticities of imports for 15 developed economies, and find no two estimates

*Thanks are due to an anonymous referee, to Brent Neiman, and to Andrei Levchenko, for their discussions of an early version of this paper. We are grateful to audiences at the September 2010 San Francisco Fed Pacific Basin Research Conference and the January 2011 American Economic Association Meetings. All errors are our own.

[†]Paris School of Economics (CNRS) and CEPR. Corresponding author: Paris School of Economics, 106-112 boulevard de l'Hopital, Paris, France 75013. jeanimbs@gmail.com, www.jeanimbs.com

[‡]Ecole Polytechnique, CREST and CEPR, isabelle.mejean@polytechnique.edu, http://www.isabellemejean.com

are significantly different from each other. Little has changed since then: the aggregate price elasticity of imports is a small number, close to zero from below, and it tends to be so for most countries. This conflicts with the view that the resilience of domestic production varies across sectors, depending for instance on the goods' quality and substitutability, or on the productivity of the firms producing in each sector. Under such a view, it is natural and intuitive to expect that the aggregate price elasticity of imports ought to vary significantly across countries, depending on the overall specialization of their production across sectors. For instance, the aggregate elasticity should be different in a developing country specialized in commodities, and in a mature diversified economy that produces high-quality goods and services. Yet, to our knowledge, there are no estimates of aggregate elasticities that display any meaningful cross-country heterogeneity.

Of course, this fact could come from difficulties in estimating such a parameter in the aggregate: because of issues of measurement, of endogeneity, or of aggregation. In this paper, we propose to infer the aggregate trade elasticity on the basis of a model. We build from the well known result in Arkolakis et al. (2012) that the welfare gains from trade depend on two parameters: the price elasticity of trade, and the degree of openness to foreign trade. Estimates of the welfare gains ought to be the same irrespective whether they are computed on sector-level or aggregate data. We use this equivalence to introduce an "aggregate" trade elasticity such that the welfare gains from trade implied by the one-sector version of Arkolakis et al. (2012) are equal to the welfare gains implied by sector-level estimates of trade elasticities. By definition, the aggregate trade elasticity is given by an adequately weighted average of sector-specific trade elasticities. The weights depend on the composition of sector expenditures, and on the distribution of openness at sector level. As such, the aggregate trade elasticity we introduce depends directly on the specialization of countries' production and trade.

The exercise enables us to decompose the sources of cross-country differences in aggregate trade elasticities, into differences in the sizes of sectors, their openness, and their trade elasticity. We find that aggregate trade elasticities are close to the US in most developed countries, with the exceptions of the UK, Australia, and Canada, where they are substantially larger in absolute value. Most emerging markets have much larger elasticities (in absolute value) than the US, e.g., China, Turkey, Chile, or Slovakia. In both cases, the differences come from high values of trade elasticities at sector level. Sector-level heterogeneity in trade elasticities has first-order consequences on the aggregate trade performance of countries.

We conduct the exercise in a panel of 28 developed and developing countries. The approach builds on 3-digit sector-level estimates of trade elasticities for all countries in the panel. The micro estimates are obtained following the methodology in Imbs and Mejean (2015), where we focused on the US. The estimation is based on a structural model, and thus immune to the conventional endogeneity issues that plague elasticity estimates.[1] The resulting estimates of the aggregate elasticity ε_j range from -3.4 to -9.9, with lowest values in Cyprus, Chile, and China. The cross-country average equals -5.9. Developed economies have estimates around -5, -4.9 in the US, although Canada, the UK, Australia, and Greece have substantially larger values (in absolute value), closer to -6, or even -9 in Greece. The values of ε_j are clearly significantly different from classic trade elasticity estimates obtained from aggregate data.

We decompose these international differences into three components: (i) the dispersion in sector-level elasticities across countries, (ii) differences in sectoral openness to trade, and (iii) differences in the sectoral allocation of expenditures. Trade in most economies in Western Europe is about as elastic as in the US (-4.9), and the differences are minimal across all three components of our decomposition: France, Germany, Hong Kong, Japan have similar aggregate elasticities, and also similar sector-level patterns. Some exceptions are Norway, Sweden, the UK, Australia, Canada, and Greece, that all have substantially larger elasticities in absolute value. In all these cases, the differences reflect the fact that sectoral trade is more elastic than in the US, especially in large and open sectors. This seems to correspond to relatively specialized developed economies - either in commodities, or perhaps in the financial sector. An exception is Austria, whose estimated elasticity is close to the US (-4.8). On average,

[1] All elasticity estimates are available on our websites, along with the codes used to obtain identification.

sector-level elasticities in Austria are substantially higher than in other developed countries. But this is offset by the fact that those sectors with relatively low elasticities are in fact the big sectors of Austria, both in terms of domestic expenditures and of trade - i.e the aggregate trade elasticity is low.

Most developing economies have aggregate elasticities that are larger than in the OECD. In most cases, this happens because sector-level elasticities are larger in absolute value. For instance, the Chinese trade elasticity is -6.9, and this happens because in China, large and open sectors are relatively more elastic than in the US. The same is true of Chile, Turkey, or Slovakia. An interesting exception is Malaysia, whose elasticity is -3.4. This happens not because trade is inelastic at sector level, but rather because the bulk of final expenditures falls on closed sectors. As a result, the domestic price index is relatively insulated from foreign shocks - i.e. the trade elasticity is low.

The international differences uncovered in this paper point to the importance of sectoral specialization in explaining the aggregate elasticity of trade, and ultimately welfare. We show that the dispersion in sector-specific elasticities has first-order effects on the aggregate trade elasticity. This stands in stark contrast with estimates of a country's trade elasticity arising from macroeconomic data, that are virtually identical across countries. Our paper takes direct inspiration from Arkolakis et al. (2012), whose model has recently been extended to cases allowing for sector-level heterogeneity. Ossa (2015) and Costinot and Rodriguez-Clare (2013) discuss the implications of heterogeneous trade elasticities. Levchenko and Zhang (2014) allow for sector-specific degrees of openness. In this paper, we combine both dimensions. We have two objectives in doing so: to introduce an aggregate trade elasticity that depends on countries' economic specialization, and to evaluate the empirical relevance of either source of heterogeneity in driving cross-country differences in trade performance. Even though we focus on the aggregate trade elasticity, it should be clear that our conclusions extend readily to measures of the welfare gains from trade.

The rest of the paper is structured as follows. Section 2 discusses our measure of aggregate trade elasticity, which relies on the equivalence of welfare formulas arising from the one-sector and the multi-sector versions of Arkolakis et al. (2012). Section 3 describes our estimation of sector-level elasticities, and data sources. Section 4 computes the aggregate trade elasticities ε_j implied by the multi-sector model, and compares them with macroeconomic estimates. The section closes with a decomposition of international differences in trade elasticities. Section 5 concludes.

2 Measuring aggregate trade elasticity

In this section, we introduce a theoretical measure of the aggregate trade elasticity. We build on the class of models used in Arkolakis et al. (2012) and Costinot and Rodriguez-Clare (2013), which imply a direct mapping between the welfare gains from trade and the price elasticity of trade. We describe the mapping in both a one-sector and a multi-sector vesion of the model. The aggregate trade elasticity is then defined as the value that equals the welfare gains in both cases.

2.1 Welfare in One and Multi-Sector Models

This section establishes that the one-sector version of Arkolakis et al. (2012) is nested into the multi-sector model they develop in their section 5.1. This holds true under perfect competition as well as, under some mild restrictions, under monopolistic competition. By definition, changes in aggregate welfare W_j^{MS} associated with moving to autarky in the multi-sector model are given by

$$d \ln W_j^{MS} = d \ln Y_j - d \ln P_j \tag{1}$$

where Y_j is aggregate income in country j and P_j is the price index.

Labor markets clear. Assuming balanced trade, $d \ln Y = d \ln w = 0$, where the second equality comes from the choice of labor as the numeraire. The change in welfare corresponding to a change in trade costs, e.g. a move to autarky, is entirely driven by a change in prices. Assuming Cobb-Douglas preferences across sectors, the price index is given by

$$P_j = \prod_s \left(P_j^s\right)^{\beta_j^s} \tag{2}$$

where β_j^s denotes the expenditure share in sector s, and P_j^s is the sector-specific price index. The welfare loss associated with a move to autarky is thus given by

$$d\ln W_j^{MS} = -\sum_s \beta_j^s \; d\ln P_j^s \tag{3}$$

Equation (3) says that the magnitude of the gains from trade in a multi-sector context depends on the extent of sectoral price adjustments, the gains being all the stronger (in absolute value) as the price in large sectors (in terms of expenditure) adjusts more. One would expect the magnitude of such sectoral price adjustments to be highly sensitive to the market structure and the assumption on the cost structure. As shown in Costinot and Rodriguez-Clare (2013), sectoral price adjustments can however be summarized using a single, simple, formula if i) preferences are CES within each sector, ii) trade is balanced, iii) the demand for imports is consistent with the gravity equation and iv) factors of production are used in the same way across all activities in all sectors. Under these assumptions, Costinot and Rodriguez-Clare (2013) show that sectoral price adjustments can be written as follows:

$$d\ln P_j^s = \frac{-1}{\varepsilon_j^s}\left[d\ln \lambda_{jj}^s - \delta_j^s d\ln r_j^s\right] \tag{4}$$

where δ_j^s is a dummy variable that characterizes the market structure in sector s of country j: It is equal to one under monopolistic competition with free entry and zero under perfect or Bertrand competition, or when there is monopolistic competition but restricted entry. r_j^s denotes the share of total revenues in country j generated from sector s. In autarky, it is equal to the share of sector s in expenditures, β_j^s, but not in an open-economy context. Under monopolistic competition with free entry, r_j^s impacts sectoral prices since entry into one sector entails gains from new varieties.

The main insight of Arkolakis et al. (2012) is that the response of sectoral prices to a foreign shock ultimately depends on the magnitude of terms-of-trade adjustments, whether those terms-of-trade adjustments take place at the intensive or the extensive margin. To measure those terms-of-trade adjustments ex-post, it is sufficient to quantify the impact that the shock has had on the domestic share in consumption λ_{jj}^s and multiply it by the inverse of the price elasticity of trade to convert the quantity adjustment into a welfare equivalent (i.e., prices). As underlined in Costinot and Rodriguez-Clare (2013), market structure is shown to matter for price adjustment in the multi-sector case, contrary to the one-sector model. This is because the cross-sectional mobility of factors ($d \ln r_j^s$) can partially compensate for any reduction of trade induced by a positive foreign shock.

Combining equations (3) and (4) gives a measure of welfare gains from trade, in a multi-sector context:

$$d \ln W_j^{MS} = \sum_s \frac{\beta_j^s}{\varepsilon_j^s} \left[d \ln \lambda_{jj}^s - \delta_j^s d \ln r_j^s \right] \tag{5}$$

In the special case of moving to autarky:

$$d \ln W_j^{MS} = \sum_s \frac{\beta_j^s}{\varepsilon_j^s} \left[- \ln \lambda_{jj}^s + \delta_j^s \ln \frac{r_j^s}{\beta_j^s} \right]$$

Costinot and Rodriguez-Clare (2013) use this framework to quantify the gains from trade in a sample of 32 countries. They notably compare the numbers obtained under different assumptions about market structure, and $\delta_j^s = 0$ or $\delta_j^s = 1$. The differences are negligible: welfare gains are roughly similar, averaging 14% under monopolistic competition and 15.3% under the alternative assumptions. Based on this result, the rest of our analysis focuses on the case $\delta_j^s = 0$ for all j, s, assuming the reallocation of revenues across sectors has negligible consequences.

Consider now a one-sector version of this model. By definition, the one-sector version imposes unique parameters, i.e. $\varepsilon_j^s = \varepsilon_j$ and $\lambda_{jj}^s = \lambda_{jj}$. While ε_j is still undefined at that point, $\lambda_{jj} \equiv \frac{\sum_s X_{jj}^s}{\sum_s Y_j^s}$, the aggregate share of domestic expenditures, where X_{jj}^s and Y_j^s denote domestic and total expenditures in sector s, is a weighted average of sectoral shares, defined as:

$$d\ln\lambda_{jj} = \frac{\sum_s \beta_j^s \; d\lambda_{jj}^s}{\lambda_{jj}} = \sum_s \frac{\lambda_{jj}^s}{\lambda_{jj}} \beta_j^s \; d\ln\lambda_{jj}^s \tag{6}$$

Constraining sector-level heterogeneity away, the expression for the welfare gains of trade in a multi-sector environment becomes

$$\frac{1}{\varepsilon_j} \sum_s \beta_j^s \; d\ln\lambda_{jj}^s = \frac{1}{\varepsilon_j} d\ln\lambda_{jj} = d\ln W_j^{OS} \tag{7}$$

where the first equality makes use of equation (6). W_j^{OS} is welfare in the one-sector model. The one sector model is a special case of the multi-sector version, with heterogeneity assumed away.

2.2 Aggregate trade elasticity in the one-sector model

Inasmuch as they stem from the same theory, the two versions must have the same welfare implications provided ε_j and λ_{jj} are calibrated adequately. This property implies a definition of the aggregate trade elasticity ε_j as a function of sector-level ones.[2] We define ε_j as being the value of the elasticity that results in equal welfare gains in the one- and multi-sector versions of the model. In particular:

$$\varepsilon_j = \frac{d\ln\lambda_{jj}}{d\ln W_j^{MS}} = \left(\sum_s \beta_j^s \frac{\ln\lambda_{jj}^s}{\ln\lambda_{jj}} \frac{1}{\varepsilon_j^s} \right)^{-1} \tag{8}$$

An aggregate trade elasticity is defined as a weighted average of sector-level elasticities, with weights that depend on expenditure shares and openness across sectors. The paper conducts the following experiment: (i) use sector-level data to estimate ε_j^s, and calibrate β_j^s, λ_{jj}^s and λ_{jj}, and (ii) use these numbers to infer the value of ε_j that must be used if welfare gains are to be

[2] An alternative would be to estimate aggregate trade elasticities directly from aggregate data. As shown in Imbs and Mejean (2015), estimating ε_j from aggregate data can deliver a value that is significantly different from the corresponding weighted average of ε_j^s, obtained from sectoral data. This is explained by a heterogeneity bias.

identical across the two versions.

Having recovered the (aggregate) elasticity which is consistent with the evidence at sector level, it is possible to analyze the cross-country heterogeneity in aggregate elasticities (and in aggregate welfare) that is attributable to various dimensions of heterogeneity. There are three potential drivers of cross-country heterogeneity, which potentially interact with each other: i) cross-country differences in the structure of consumption (in the $\{\beta_j^s\}$), ii) cross-country differences in the sensitivity of sectoral trade to price adjustments (in the $\{\varepsilon_j^s\}$), and iii) cross-country differences in the openness of different sectors (in the $\{\lambda_{jj}^s\}$). Section 4 discusses how these various sources of heterogeneity matter for cross-country estimates.

3 Estimation and Data

In equation (8), the only parameters that are not directly observed from the data are the sectoral elasticities. We now summarize the approach used to estimate them using import data observed at sector level. The empirical strategy is inspired by Feenstra (1994) and detailed in Imbs and Mejean (2015). That paper also implements alternative approaches, notably a gravity-type regression consistent to Caliendo and Parro (2015). Results presented for the US were consistent across empirical strategies, so that we focus in this paper on the structural strategy described below. The section closes with a review of the data needed for estimations and welfare computations.

3.1 Estimation

The strategy consists in estimating structurally an equilibrium model of bilateral trade flows. The demand-side of the model features Constant Elasticity of Substitution between varieties of (disaggregated) products exported by various countries. The import demand equation writes as follows:

$$d \ln s_{ijt}^s = \varepsilon_j^s \; d \ln P_{ijt}^s + \Phi_{jt}^s + \xi_{ijt}^s \tag{9}$$

where i denotes a variety, i.e. an origin country and t is a time indicator. s^s_{ijt} is the market share of country i in expenditures on good s of country j, at time t. The intercept Φ^s_{jt} is time-varying and common across origin countries. Finally, ξ^s_{ijt} is an error term combining preference shocks and trade costs. The shocks are assumed to be independent and identically distributed across sectors and countries. To account for the endogeneity of prices, Feenstra (1994) imposes a simple supply structure:

$$P^s_{ijt} = \exp(\upsilon^s_{ijt}) \left(C^s_{ijt}\right)^{\frac{\omega^s_j}{1-\omega^s_j}}$$

where C^s_{ijt} is real consumption of good k imported from country i, and ω^s_j maps into the price elasticity of supply in sector s. The technology shock υ^s_{ijt} is independent and identically distributed across sectors and countries. After rearranging, this implies

$$d\ln P^s_{ijt} = \omega^s d\ln s^s_{ijt} + \Psi^s_{jt} + \delta^s_{ijt} \tag{10}$$

where Ψ^s_{jt} is a time-varying intercept common across origin countries, and $\delta^s_{ijt} = \left(1 - \omega^s_j\right) dv^s_{ijt}$ is an error term that depends on supply shocks. Solve equation (9) for ξ^s_{ijt}, and equation (10) for δ^s_{ijt}, express both in deviations from a reference country r, and multiply term for term to obtain:

$$Y^s_{ijt} = \psi^s_{1j}\ X^s_{1ijt} + \psi^s_{2j}\ X^s_{2ijt} + e^s_{ijt} \tag{11}$$

where $Y^s_{ijt} = (d\ln P^s_{ijt} - d\ln P^s_{rjt})^2$, $X^s_{1ijt} = (d\ln s^s_{ijt} - d\ln s^s_{rjt})^2$, $X^s_{2ijt} = (d\ln s^s_{ijt} - d\ln s^s_{rjt})(d\ln P^s_{ijt} - d\ln P^s_{rjt})$, and $e^s_{ijt} = -\left(\xi^s_{ijt} - \xi^s_{rjt}\right)\left(\delta^s_{ijt} - \delta^s_{rjt}\right)\frac{1}{\varepsilon^s}$.

Feenstra (1994) observes that the time average of e^s_{ijt} is zero, provided the shocks ξ^s_{ijt} and δ^s_{ijt} are orthogonal to each other. The time averages of X^s_{1ijt} and X^s_{2ijt} constitute therefore appropriate instruments in equation (11), since $cov_{ijt}(\bar{X}^s_{1ij}, e^s_{ijt}) = cov_{ijt}(\bar{X}^s_{2ij}, e^s_{ijt}) = 0$. They solve the issue of endogeneity present in the import demand equation.[3] Since they are averages over time, identification is effectively obtained across countries.

[3] In practice, Common Correlated Effects are included in equation (11) to avoid double counting in a cross-country panel, and an intercept is included to account for the measurement error arising from the unit values used to approximate prices. Given the origin of potential measurement error, the intercept is allowed to vary at the most disaggregated level, i.e. for each HS6 category.

The procedure in Feenstra (1994) consists in estimating equation (11) and recovering the structural parameters $\hat{\varepsilon}_j^s$ and $\hat{\omega_j^s}$ from the estimated coefficients, $\hat{\psi}_{1j}^s$ and $\hat{\psi}_{2j}^s$. For some combinations of the estimated coefficients, however, the recovered values are not theoretically consistent.[4] In such circumstances, we follow Broda and Weinstein (2006). We apply a grid search algorithm over all the theoretically-consistent values for $(\varepsilon_j^s, \omega_j^s)$ and select the combination of parameters which minimizes the root mean square error. Because we do not want this procedure to create a bias, we restrict the grid search to values of ε_j^s higher than -29.[5]

3.2 Data

Sectoral information is needed on bilateral imports and unit values (i.e., prices) at sector-level for a cross-section of countries. We use the CEPII-BACI database documented in Gaulier and Zignago (2010). The data reports multilateral trade at the 6-digit level of the harmonized system (HS6), and cover around 5,000 products for a large cross-section of countries. The universe of products is partitioned into sectors according to the 3-digit ISIC (revision 2) level, which makes for a maximum of 27 sectors. Price elasticities are estimated for each ISIC sector of each importing country, but the data are collected at the most disaggregated (HS6) level. The data are yearly between 1995 and 2004. Before 1995, the number of reporting countries is unstable, and the unit values reported in BACI experience a structural break in 2004.

Identification requires that the cross-section of countries be wide enough for all sectors, and remain so over time. We retain goods for which a minimum of 20 exporting countries are available throughout the period. Both unit values and market shares are notoriously plagued by measurement error. We compute the median growth rate at the sector level for each variable, across all countries and years. When growth rates exceed five times this median value in one sector-country pair, this sector-country pair is dropped. The resulting sample covers about 85

[4] This happens when ψ_{1j}^s is not significantly positive.

[5] Equation (8) implies that $\partial\varepsilon_j/\partial\varepsilon_j^s$ is proportional to $\left(\frac{\varepsilon_j}{\varepsilon_j^s}\right)$.It is therefore conservative to limit the absolute value of ε_j^s, so that the sectors whose elasticities are estimated using a grid search continue to matter for the aggregate. They matter less if we use instead the lower bound used by Broda and Weinstein (2006), set at -79. (Although in results available upon request we verified that our main conclusions are not altered).

percent of world trade. Table 1 presents some summary statistics for the 28 countries with available data. The number of sectors (and the number of estimated elasticities ε_j^s) ranges from 10 to 27.

The main data constraint concerns the weights that enter equation (8). Both β_j^s and λ_{jj}^s require information on domestic consumption at sectoral level that must be compatible with the trade data in BACI. The constraint raises issues of concordance since information is needed on both production and trade at the sectoral level. This is what reduces the coverage to 28 countries. We use a dataset built by di Giovanni and Levchenko (2009) who merge information on production at the 3-digit ISIC (revision 2) level from UNIDO and on bilateral trade flows from the World Trade Database compiled by Feenstra et al. (2005). Domestic consumption at the sectoral level is computed as production net of exports, and overall consumption is production net of exports but inclusive of imports. We define

$$\beta_j^s \equiv \frac{Y_j^s - X_j^s + M_j^s}{\sum_k \left(Y_j^s - X_j^s + M_j^s\right)}$$

where X_j^s (M_j^s) denotes country j's exports (imports) in sector s and Y_j^s the value of its production. And

$$\lambda_{jj}^s \equiv \frac{Y_j^s - X_j^s}{Y_j^s - X_j^s + M_j^s}$$

To focus on meaningful computations, a minimum of 10 sectors is imposed for all countries. The constraint tends to exclude small or developing economies, such as Panama or Poland. The UNIDO data are in USD, and available at a yearly frequency. The values of β_j^s and λ_{jj}^s are computed over five-year averages in order to limit the consequences of cyclical fluctuations in trade. Two sets of estimations have been considered. In the main text, we use average weights between 1991 and 1995. For robustness, we have also considered averages between 1996 and 2000. Results are very similar and are available upon request.

The UNIDO dataset is focused on manufacturing goods only. The vast majority of traded goods are manufactures, so that the truncation remains minimal. We have experimented with the values for β_j^s and λ_{jj}^s implied by the OECD Structural Analysis database (STAN), which provides information on all sectors of the economy. For countries covered by both datasets, i.e. OECD members, the end elasticities were in fact virtually identical. At least for OECD members, this suggests the sampling issue caused by the UNIDO dataset is kept to a minimum. The last column in Table 1 reports the fraction of total trade covered by UNIDO data. The coverage is below 40 percent for small open economies such as Hong Kong, Cyprus, or Chile, but above 70 percent for large developed economies such as the US, France or Spain. Coverage is clearly limited for small open, developing economies. But, contrary to OECD data, UNIDO leaves the door open to some analysis for the developing world, not least China where coverage is above 50%.

4 Trade Elasticities in the One-Sector Model

We report the estimates of ε_j^s implied by sectoral data for the 28 countries with the required data, and discuss the corresponding values of ε_j. One-sector elasticities are compared with conventional macroeconomic estimates, and then with a weighted average of sectoral elasticities ε_j^s. The Section closes with a decomposition of the international differences in estimates of ε_j.

4.1 Sector-Level and Aggregate Trade Elasticities

Table 2 presents some summary statistics of the estimates of ε_j^s implied by BACI data between 1995 and 2004. There is considerable heterogeneity in mean sectoral elasticities across countries. Developed countries display average values around -5: Germany at -4.6, France at -4.8, or the US at -5.9. In contrast, developing exporting economies present estimates at least twice larger. Cyprus has the largest mean sectoral elasticity, equal to -14.4, closely followed by Chile, Indonesia and Guatemala.

Sectoral heterogeneity is sizeable within countries as well. The distribution of estimates tends to be most disparate and skewed in developing economies. For instance, estimates of ε_j^s

range between −1.8 and −29.0 in Chile, with a median of −8.2, substantially below the mean of −12.2. In Indonesia, estimates range from −2.3 to −29.0.[6] Ranges tend to be narrower for European developed countries, such as France, Germany, Italy or the UK. The distributions tend to be more symmetric also, with mean and median elasticities closer together.

Country and sector effects each explain approximately 10 percent of the cross-country dispersion in estimates of ε_j^s. Close to 80 percent of the variance in ε_j^s must therefore correspond to international differences in the trade elasticity for each sector s. The result is apparent from Table 3, where some sectoral estimates are drastically different from one country to the next. For instance, the elasticity for Fabricated Metal products is −3.5 in France, but −26.6 in Indonesia. Imports of Potteries are inelastic in Australia ($\hat{\varepsilon}_j^s = -1.9$), but elastic in Sweden ($\hat{\varepsilon}_j^s = -29.0$). Such disparities may correspond to differences in the very nature of the goods imported. For instance, Metal products imported by France are likely to be of higher quality than those imported by Indonesia. Sector elasticities correlate positively with import penetration (0.22, significant at the one percent confidence level), which is consistent with pro-competitive effects of trade, as in Edmond et al. (2015) or Chen et al. (2009). There is no significant correlation between sector elasticities and expenditure shares.

Values for the aggregate trade elasticities implied by those sectorestimates are reported in Table 4. The first two columns in Table 4 report the calibrated values of λ_{jj}, and the corresponding estimates of ε_j for the 28 countries with data. Standard errors are obtained using the Delta method detailed in the Appendix. There are considerable cross-country differences in aggregate elasticity estimates. They range from around −3.4 in Malaysia, down to −10.0 in Cyprus. Intermediate values between −4 and −5 are found for developed economies, with −4.9 for the US or −5.4 for the UK. No obvious correlate of ε_j is apparent from Table 4, as developing economies can be found at either extreme of the range of estimates.

[6] Note that these intervals are somewhat misleading because of the lower bound imposed on estimated elasticities obtained using a grid search procedure. For instance, the range of estimated elasticities is quite large in Table 2 for the US, $[-29.0, -3.0]$ but is strongly reduced once the single elasticity equal to -29 is neglected, to $[-5.1, -3.0]$. This is typically not the case for developing countries. For instance, once the single value of −29 is dropped, the intervals are equal to $[-24.6, -1.8]$ in Chile, and $[-24.3, -2.3]$ in Indonesia.

This heterogeneity maps directly into the welfare gains from trade implied by the multi-sector version of Arkolakis et al. (2012). Column (3) in Table 4 reports the welfare loss $d\ln W_j^{MS}$ associated with a move to autarky, given by:

$$d\ln W_j^{MS} = -\sum_s \frac{\beta_j^s}{\varepsilon_j^s} \ln \lambda_{jj}^s$$

whose estimation requires calibrated values for β_j^s and λ_{jj}^s. The welfare losses from autarky are highest in small open economies, like Hong Kong (36.3% of real income) or Malaysia (28.2%). They are lowest in large, closed economies, such as Japan (1.2%) , India (2.3%), China (3.0%) or the US (3.6%). On average, the losses are estimated around 7.5% of real income for developed, West European economies.

The ranking correlates with measures of overall openness, as reflected in the aggregate share of domestic expenditures $\lambda_{jj} \equiv \frac{Y_j - X_j}{Y_j - X_j + M_j}$. It takes lowest values in small open economies, like Hong Kong or Malaysia, and highest in large or closed countries, such as Japan, the US or India. But openness is not the sole determinant of welfare: $d\ln W_j^{MS}$ also decreases with trade elasticities, and depends on their distribution across sectors. The comparison of Greece and Austria is illustrative of two countries which are roughly as open to trade, spending around half of their consumption on imported goods, but with different sectoral specialization and elasticities, leading to different levels of welfare losses. The welfare loss from moving to autarky is twice as large in Austria than in Greece (respectively 13.5 and 7.8%). The difference is in part attributable to lower average sectoral elasticities in Austria, as illustrated in Table 2. It also comes from the cross-sector correlation between λ_{jj}^s and ε_j^s. For given average openness and average trade elasticity, the welfare loss $d\ln W_j^{MS}$ takes higher (absolute) value if open sectors tend to display low elasticities. This tends to happen in Austria, an open economy on average, whose imports are specialized in sectors with low trade elasticities. The specialization of trade matters for the aggregate trade elasticity, and therefore for welfare.

4.2 Comparisons

Table 4 does suggest an important result: the estimates of ε_j are unusual for one-sector models. We now compare our estimates of ε_j with alternative candidates. Trade elasticity estimates that arise from aggregate data are first considered. Aggregate data are the most natural source when it comes to estimating parameters that enter one-sector models. It is self-evident from Table 4 that our estimates of ε_j are significantly different from the conventional values for import price elasticities obtained in macroeconomics.[7] For instance, Figure 1 reproduces the estimates obtained in Houthakker and Magee (1969) for 15 developed economies. No point estimates are below -2, some are positive, and 10 out of 15 are not significantly different from zero. In fact, virtually no two estimates are significantly different from each other. For instance, the US price elasticity of imports is -0.5, Japan's is -0.78, and Canada's is -1.5.

Estimates of the aggregate trade elasticity have not changed sizeably in the considerable literature that followed Houthakker and Magee (1969). We confirm this conducting a similar estimation using trade data, which we aggregated to country level in order to estimate a gravity equation. Using the notation from section 3, we estimate

$$\Delta \ln P_{ijt}C_{ijt} = A_{ij} + (\varepsilon_j^A + 1)\ \Delta \ln P_{ijt} + \tilde{\nu}_{ijt} \tag{12}$$

where $\Delta X_t = X_t - X_{t-1}$, $P_{ijt}C_{ijt} = \sum_s P_{ijt}^s C_{ijt}^s$ and $\Delta \ln P_{ijt} = \frac{1}{2}\sum_s \left(\frac{P_{ijt}^s C_{ijt}^s}{P_{ijt}C_{ijt}} + \frac{P_{ijt-1}^s C_{ijt-1}^s}{P_{ijt-1}C_{ijt-1}}\right) \Delta \ln P_{ijt}^s$ is a Tornqvist price index. Identification is obtained through time variation. Column (4) in Table 4 reports the estimates of ε_j^A. In comparison with ε_j, the estimates of ε_j^A are much closer to zero. Of course, equation (12) is acutely problematic, as changes in prices are endogenous. But it is unlikely a correction for endogeneity would imply estimates of ε_j^A close to ε_j. At the very least, no existing estimates using aggregate data come even close.

Estimates of ε_j can therefore not be reproduced from aggregate data. Rather, they are

[7] See for instance the estimates reported in Francis et al. (1976). The book contains summaries of available aggregate elasticities from the literature.

given by a weighted average of sector level estimates ε_j^s, with weights given by $\beta_j^s \frac{\ln \lambda_{jj}^s}{\ln \lambda_{jj}}$:

$$\varepsilon_j = \left(\sum_s \beta_j^s \frac{\ln \lambda_{jj}^s}{\ln \lambda_{jj}} \frac{1}{\varepsilon_j^s} \right)^{-1}$$

The weights reflect the relative openness to trade and each sector's importance in overall consumption. The specialization of the economy matters in two ways: First, sectors that compose a large fraction of total expenditures receive a small weight. For a given shock and a given sectoral elasticity, a large value of β_j^s implies a large response of the overall price index, i.e. low aggregate trade elasticity. For the same reason, relatively open sectors enter with a small weight. For a given shock to traded quantities, a large value of $\frac{\ln \lambda_{jj}^s}{\ln \lambda_{jj}}$ means a large response of the sectoral price index. The response of the aggregate price index is accordingly large, which means low aggregate trade elasticity.

4.3 International Differences

International differences in trade elasticities are absent from estimates obtained from aggregate data. In macroeconomics, trade elasticities are customarily assumed to be identical across countries, and thus invariant to differences in the specialization of trade across countries. This is an undesirable property in light of anecdotal and journalistic arguments that the specialization of production or trade has direct implications on countries' external performance.

The trade elasticity introduced in this paper does not share this property. Cross-country estimates of ε_j display considerable heterogeneity, and theory can be used to identify its sources. Using its definition, it is easy to show how ε_j decomposes. In particular, a Taylor expansion of equation (8) around a reference country r implies

$$\frac{\varepsilon_j - \varepsilon_r}{\varepsilon_r} = \underbrace{-\sum_s Sh_r^s \frac{\beta_j^s - \beta_r^s}{\beta_r^s}}_{B_j} \underbrace{-\sum_s Sh_r^s \frac{\Delta\lambda_j^s - \Delta\lambda_r^s}{\Delta\lambda_r^s}}_{L_j} + \underbrace{\sum_s Sh_r^s \frac{\varepsilon_j^s - \varepsilon_r^s}{\varepsilon_r^s}}_{E_j} \tag{13}$$

$$\text{where} \qquad Sh_r^s \equiv \frac{\beta_r^s \frac{\ln \lambda_{rr}^s}{\ln \lambda_{rr}} \frac{1}{\varepsilon_r^s}}{\sum_s \beta_r^s \frac{\ln \lambda_{rr}^s}{\ln \lambda_{rr}} \frac{1}{\varepsilon_r^s}} \quad \text{and} \quad \Delta\lambda_j^s \equiv \frac{\ln \lambda_{jj}^s}{\ln \lambda_{jj}}$$

Equation (13) implies that the international dispersion in trade elasticities is determined by three terms. The first (B_j) reflects international differences in the sectoral composition of expenditures. The second one (L_j) reflects differences in sectoral openness, and the third (E_j) reflects differences in sectoral trade elasticities. In absolute value, ε_j is relatively high if (i) consumers spend less (relative to the reference country) in open and inelastic sectors, (ii) large and inelastic sectors are closed (relative to the reference), and (iii) sectors that are elastic (relative to the reference) also tend to be large and open.

Performing the decomposition described in equation (13) is straightforward, given the data requirements involved in computing ε_j. For reference, Figure 2 reproduces the cross-country estimates of ε_j as implied by Table 4. Figure 3, panel (a), illustrates how they compare with the elasticity obtained for the US ($\varepsilon_j - \varepsilon_{US}$). Finally, Figure 3, panel (b), reports the approximate decomposition in equation (13) using the US as reference country. To be precise, it applies the decomposition in equation (13) to the country-specific gaps reported in Figure 3, panel (a).[8]

It is interesting to note that high average estimates of ε_j^s, which tend to happen in the developing world as shown in Table 2, do not necessarily translate into large values for E_j. For instance, Chile or Greece have large positive values of E_j, whereas they are negative in Indonesia. As is obvious from equation (13), there is no correlation between sectoral averages of ε_j^s and the value of E_j. International differences in ε_j arise because the sectoral distributions of ε_j^s, β_j^s and $\Delta\lambda_{jj}^s$ change from one country to the next. Figure 3 suggests these international

[8]Equation (13) is an approximation, which creates discrepancies between the actual gap $\varepsilon_j - \varepsilon_{US}$ and its value implied by the sum of B_j, L_j, and E_j. But the approximation is inconsequential since the correlation between $\varepsilon_j - \varepsilon_{US}$ and its approximate components is above 90 percent across countries.

differences are smallest as regards β_j^s, as the B_j term tends to be the least important element of $\frac{\varepsilon_j-\varepsilon_r}{\varepsilon_r}$, except perhaps in Malaysia. The main reason why estimates of ε_j vary across countries appears to be summarized in E_j, i.e., in cross country differences in the sector-estimates of trade elasticities. Some differences in L_j do exist as well, but they are much smaller, except perhaps in Austria.

Several results are of interest. In most cases, the estimates of ε_j are larger (in absolute value) or similar to the US. Amongst developed countries, only Germany has an estimate closer to zero than in the US, at -4.3; France's estimate is -4.6. Among developing countries, elasticity estimates are observably larger in absolute value, except in Guatemala (-4.5), Indonesia (-4.4), and Malaysia (-3.4). Quite a few countries display aggregate elasticities that are estimated substantially above that in the US. For instance, Greece and Chile both display values of ε_j around -9; Canada's estimate is -8, Slovakia's -7.8, while China's is -6.9, Hungary's, Portugal's, and Turkey's are -6.4. As is apparent from this list, there is no systematic correlation beween income levels and elasticity estimates: even though most developed countries have estimates of ε_j in the US ballpark (Germany, France, Austria, South Korea, Italy, or Japan all have estimates around -5), there are exceptions.

Figure 3 does however reveal a systematic pattern across countries: high estimates of ε_j typically arise because of a high value for E_j. Countries with elastic trade are ones that tend to import more of (relatively) high-elasticity goods. In the conventional view of trade elasticities, this corresponds to heterogeneous price elasticities of demand for a given sector across countries. In the model of Eaton and Kortum (2002), international differences in estimates of ε_j^s correspond to differences in the dispersion of firm technologies at sector level. Deciding which of these two interpretations dominates in the data is beyond the scope of this paper. But the fact that high values of ε_j arise both in developed (Canada, Australia) and developing (China, Turkey) countries is suggestive that technology-based explanations can play a role. The fact that few countries have estimates of ε_j lower than the US suggests dispersion in firm technology is in fact highest in the US in our sample, i.e., that the US constitutes a legitimate benchmark.

Among the countries with largest estimates, China, Slovakia, and Canada stand out: in all three cases, the aggregate elasticity would be the same as in the US if sector-level estimates of trade elasticities were those of the US, i.e. if E_j were zero. Thus, large aggregate differences come from international differences in estimates of the elasticity of a given sector, consistent with Table 3. Australia, Chile, Greece, and to a lesser extent Turkey all share the same property, but the end effect on ε_j is mitigated by negative values of L_j: In these small open economies, inelastic sectors tend to be much more open to international trade than in the US, which has mitigating consequences on the estimates of ε_j. Australia, in particular, would have a trade elasticity as high as Canada (-8) if the relative openness of its sectors was the same as in the US.

Most European elasticities are similar to the US, but this masks some important differences. The United Kingdom, for instance, would have a much higher elasticity if large sectors were more open, closer to that of Turkey ($-6, 5$). This is because its importing sectors do tend to display high values of ε_j^s. This is also true of Norway and Sweden, where the effects of E_j and L_j on ε_j work in opposite directions. Austria is an extreme case of the same pattern: based on the estimates of ε_j^s there, the value of E_j in Austria would imply an aggregate elasticity close to -10, instead of the -4.8 we estimate. This illustrates the importance of letting both trade elasticities and the extent of openness vary by sector.

Interestingly, Germany tends to display similar estimates of ε_j^s than the US: but its aggregate elasticity is closer to zero (-4.3) because both L_j and B_j take negative values. Thus, the relatively low elasticity of German trade comes not from especially low trade elasticities at sector level, but rather from the structure of final consumption, and of openness.

A few developing countries display elasticity estimates in the US ballpark, sometimes even closer to zero. It is especially the case of Malaysia and Indonesia, Guatemala to a lesser extent. These constitute interesting exceptions. Malaysia's estimate of $\varepsilon_j = -3.4$ is the closest to zero in our sample. Figure 3 reveals this happens strictly for structural reasons: while $E_j > 0$ in Malaysia, just like it is in most other developing countries, B_j and L_j are both negative, quite sizeably so. $B_j < 0$ in particular reflects the structure of final expenditures in Malaysia, that tends to fall on relatively closed sectors, so that the response of aggregate quantities to international prices is muted. Malaysia provides an illuminating illustration of the decomposition introduced in this paper, emphasizing that, in principle, the structure of the economy matters as much as sector-level elasticities. Guatemala constitutes a similar example, with $E_j > 0$ but $L_j < 0$. Finally, Indonesia is an outlier, as it is the only developing country that display negative values for E_j, i.e. sector elasticity estimates that are closer to zero than in the US.

The decomposition of ε_j is relevant to understanding the international dispersion in trade elasticities. It is of course also important for welfare. The welfare gains from trade decrease in the trade elasticity, so that large estimates of ε_j mean lower welfare than what is implied by aggregate data. For instance, Figure 3 suggests the welfare gains from trade in China would be substantially higher if the distribution of sectoral elasticities were closer to the US. They would similarly be higher in Canada. To our knowledge, there is no alternative methodology that implies such a close mapping between the sectoral specialization of consumption and production, the elasticity of trade, and ultimately the welfare gains from trade.

5 Conclusion

Estimates of the aggregate elasticity of trade are computed for 28 countries, on the basis of the multi-sector model developed by Costinot and Rodriguez-Clare (2013). One-sector and multi-sector versions of the model should have identical welfare predictions, which implies an expression for the aggregate trade elasticity given by a weighted average of sector-level elasticities. We estimate structurally 3-digit trade elasticities for 28 developed and developing countries.

Estimates of aggregate trade elasticities are significantly different from conventional, macroeconomic trade elasticities. They are larger in absolute value, and heterogeneous across countries, with values ranging between -3.4 and -9.9. China has low estimates, -6.9. Western Europe and the US display estimates around -5, but Canada, Chile and Greece are closer to -9. The lowest values are found for small-open specialized economies. Using the theory, a decomposition of this international dispersion is introduced. Trade elasticities can differ because of the specialization of consumption, of production, or because of international differences in sector-level trade elasticities. Most countries have elasticity estimates larger (in absolute value) than the US because sector-level elasticities are themselves larger in absolute value. Inasmuch as welfare depends on the trade elasticity, these decomposition carry through to welfare.

References

Arkolakis, Costas, Arnaud Costinot, and Andres Rodriguez-Clare, "New Trade Models, Same Old Gains?," *American Economic Review*, February 2012, *102* (1), 94–130.

Broda, Christian and David E. Weinstein, "Globalization and the Gains From Variety," *The Quarterly Journal of Economics*, 2006, *121* (2), 541–585.

Caliendo, Lorenzo and Fernando Parro, "Estimates of the Trade and Welfare Effects of NAFTA," *Review of Economic Studies*, 2015, *82* (1), 1–44.

Chen, Nathalie, Jean Imbs, and Andrew Scott, "The Dynamics of Trade and Competition," *Journal of International Economics*, 2009, *77* (1), 50–62.

Costinot, Arnaud and Andrés Rodriguez-Clare, "Trade Theory with Numbers: Quantifying the Consequences of Globalization," in Gita Gopinath and Elhanan Helpman and Kenneth Rogoff, ed., *Handbook of International Economics, Vol. 4*, 2013, chapter 4.

di Giovanni, Julian and Andrei A. Levchenko, "Trade Openness and Volatility," *The Review of Economics and Statistics*, August 2009, *91* (3), 558–585.

Edmond, Chris, Virgiliu Midrigan, and Daniel Xu, "Competition, Markups, and the Gains from International Trade," *American Economic Review*, 2015, *105* (10), 3183–3221.

Feenstra, Robert C, "New Product Varieties and the Measurement of International Prices," *American Economic Review*, March 1994, *84* (1), 157–77.

Feenstra, Robert C., Robert E. Lipsey, Haiyan Deng, Alyson C. Ma, and Hengyong Mo, "World Trade Flows: 1962-2000," NBER Working Papers 11040, National Bureau of Economic Research, Inc January 2005.

Francis, J., B. Schumacher, and R. Stern, *Price Elasticities in International Trade*, Macmillan Press, 1976.

Gaulier, Guillaume and Soledad Zignago, "BACI: International Trade Database at the Product-Level. The 1994-2007 Version," CEPII Working Papers 2010-23, CEPII 2010.

Houthakker, H. and S. Magee, "Income and Price Elasticities in World Trade," *The Review of Economics and Statistics*, 1969, *51* (2), 111–25.

Imbs, Jean and Isabelle Mejean, "Elasticity Optimism," *American Economic Journal: Macroeconomics*, July 2015, *7* (3), 43–83.

Levchenko, Andrei A. and Jing Zhang, "Ricardian productivity differences and the gains from trade," *European Economic Review*, 2014, *65* (C), 45–65.

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

A Appendix: Variances

A.1 The variance of ε_j

Consider a Taylor expansion of $\varepsilon_j = \left(\sum_s \beta_j^s \frac{\ln \lambda_{jj}^s}{\ln \lambda_{jj}} \frac{1}{\varepsilon_j^s}\right)^{-1}$ around its estimated value $\hat{\varepsilon}_j$. We have:

$$\begin{aligned}\varepsilon_j &= \hat{\varepsilon}_j + \sum_s \left.\frac{\partial \hat{\varepsilon}_j}{\partial \varepsilon_j^s}\right|_{\varepsilon_j^s = \hat{\varepsilon}_j^s} (\varepsilon_j^s - \hat{\varepsilon}_j^s) \\ &= \hat{\varepsilon}_j + \sum_s \left(\frac{\hat{\varepsilon}_j}{\hat{\varepsilon}_j^s}\right)^2 \beta_j^s \frac{\ln \lambda_{jj}^s}{\ln \lambda_{jj}} (\varepsilon_j^s - \hat{\varepsilon}_j^s)\end{aligned}$$

Assuming no correlation in estimated sectoral elasticities, the variance is therefore given by

$$Var(\hat{\varepsilon}_j) = \sum_s \left[\left(\frac{\hat{\varepsilon}_j}{\hat{\varepsilon}_j^s}\right)^2 \beta_j^s \frac{\ln \lambda_{jj}^s}{\ln \lambda_{jj}}\right]^2 Var(\hat{\varepsilon}_j^s)$$

A.2 The variance of $d\ln W_j^{MS}$

A Taylor expansion of $d\ln W_j^{MS} = \frac{\ln \lambda_{jj}}{\varepsilon_j}$ around its estimated value $d\hat{\ln W}_j^{MS}$ implies

$$d\ln W_j^{MS} = d\hat{\ln W}_j^{MS} - \frac{\ln \lambda_{jj}}{(\hat{\varepsilon}_j)^2}(\varepsilon_j - \hat{\varepsilon}_j)$$

The variance is therefore given by

$$Var(d\hat{\ln W}_j^{MS}) = \left(\frac{\ln \lambda_{jj}}{(\hat{\varepsilon}_j)^2}\right)^2 Var(\hat{\varepsilon}_j)$$

Figure 1: Houtakker and Magee (1969) elasticity estimates

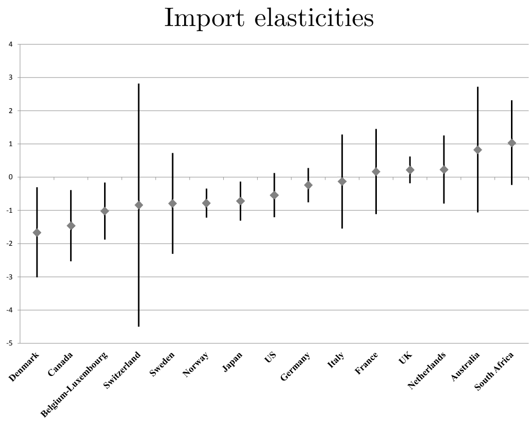

Note: The grey circles are the point estimates found in Houtakker and Magee (1969). Lines around the circles correspond to the confidence interval, at the 5% level.

Figure 2: Elasticity estimates in the one-sector model

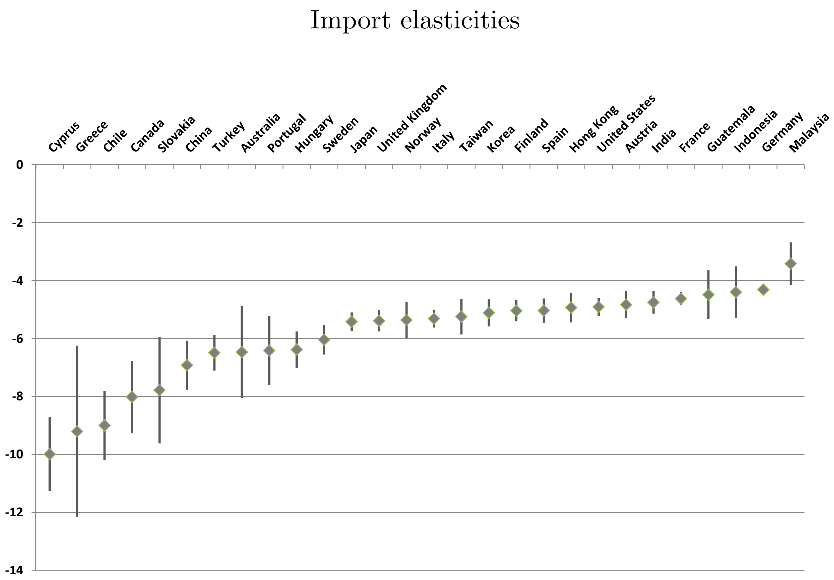

Note: The grey circles are the point estimates reported in Table 4. Lines around the circles correspond to the confidence interval, at the 5% level.

Table 1: Summary Statistics

	# sect	% Trade
Australia	17	0.471
Austria	24	0.717
Canada	24	0.645
Chile	17	0.335
China	20	0.512
Cyprus	18	0.244
Finland	26	0.654
France	26	0.784
Germany	21	0.508
Greece	17	0.428
Guatemala	18	0.369
Hong Kong	11	0.169
Hungary	19	0.471
India	18	0.337
Indonesia	15	0.425
Italy	25	0.726
Japan	26	0.611
Korea	26	0.586
Malaysia	18	0.504
Norway	20	0.498
Portugal	22	0.623
Slovakia	10	0.279
Spain	26	0.733
Sweden	25	0.729
Taiwan	20	0.401
Turkey	24	0.575
United Kingdom	26	0.811
United States	27	0.743

Notes: The first column reports the number of sectors under study. The second column is the percentage of the country's aggregate imports which the dataset covers.

Table 2: Summary statistics on estimated sectoral elasticities

Country	Count	Mean	Median	Min	Max
Australia	17	-9.9	-6.3	-29.0	-1.9
Austria	24	-6.7	-5.5	-15.4	-2.5
Canada	24	-9.0	-7.5	-29.0	-2.9
Chile	17	-12.2	-8.2	-29.0	-1.8
China	20	-7.0	-5.7	-29.0	-3.2
Cyprus	18	-14.4	-9.3	-29.0	-2.8
Finland	26	-5.7	-4.2	-22.4	-2.2
France	26	-4.8	-4.6	-9.8	-2.8
Germany	21	-4.6	-4.3	-11.1	-2.0
Greece	17	-11.1	-8.8	-29.0	-4.6
Guatemala	18	-11.5	-7.0	-29.0	-2.6
Hong Kong	11	-5.2	-5.1	-7.9	-3.6
Hungary	19	-8.0	-6.0	-29.0	-1.5
India	18	-6.0	-5.0	-21.2	-2.2
Indonesia	15	-12.2	-6.9	-29.0	-2.3
Italy	25	-5.8	-5.6	-11.7	-2.4
Japan	26	-6.4	-4.9	-25.9	-3.5
Korea	26	-5.8	-5.3	-14.2	-3.0
Malaysia	18	-7.8	-5.2	-29.0	-2.7
Norway	20	-6.5	-5.4	-17.8	-2.4
Portugal	22	-9.1	-7.8	-29.0	-2.5
Slovakia	10	-7.7	-7.2	-12.1	-4.2
Spain	26	-6.7	-5.9	-26.3	-3.2
Sweden	25	-9.9	-5.9	-29.0	-2.7
Turkey	24	-7.5	-5.8	-29.0	-3.3
Taiwan	20	-6.4	-5.2	-29.0	-2.7
United Kingdom	26	-6.3	-5.2	-13.1	-2.5
United States	27	-5.9	-5.0	-29.0	-3.0

Notes: The table reports summary statistics on the estimated elasticities, $\hat{\varepsilon}_j^s$, by importing country.

Table 3: Summary statistics on estimated elasticities, by sector

Sector	Count	Mean	Median	Min	(Country)	Max	(Country)
Food	28	-7.2	-6.1	-15.0	(Greece)	-3.8	(Finland)
Beverage	21	-6.6	-5.5	-29.0	(Malaysia)	-2.3	(Hungary)
Tobacco	3	-3.2	-2.8	-4.8	(USA)	-2.0	(Germany)
Textile	27	-11.3	-7.4	-29.0	(Australia, Chile, Guatla, Cyprus)	-3.5	(Taiwan)
Wearing Apparel	17	-13.9	-10.5	-29.0	(Australia, Cyprus, Sweden, Taiwan)	-4.6	(Korea)
Leather products	19	-8.9	-6.3	-29.0	(Greece)	-3.8	(Malaysia)
Footwear	22	-10.0	-6.9	-29.0	(Cyprus)	-3.0	(Korea)
Wood products	21	-6.0	-4.9	-23.2	(Australia)	-2.4	(Italy)
Furniture	19	-6.4	-3.6	-29.0	(USA)	-1.5	(Hungary)
Paper products	26	-4.3	-4.0	-8.0	(Portugal)	-1.8	(Chile)
Printing & Publishing	26	-6.4	-4.2	-29.0	(Malaysia)	-2.2	(Chile)
Industrial chemicals	21	-6.0	-5.0	-12.7	(Guatemala)	-4.1	(USA)
Other chemicals	21	-5.9	-5.9	-8.2	(Chile)	-2.7	(Finland)
Petroleum	12	-7.9	-4.6	-26.3	(Spain)	-2.5	(UK)
Rubber products	27	-7.4	-4.8	-29.0	(Indonesia)	-3.5	(France)
Plastic products	27	-5.6	-4.3	-29.0	(Indonesia)	-2.9	(Italy)
Potteries	18	-5.5	-3.8	-29.0	(Sweden)	-1.9	(Australia)
Glass products	26	-6.2	-4.4	-29.0	(Indonesia)	-2.4	(Chile)
Other mineral products	26	-4.0	-3.8	-7.1	(Taiwan)	-2.1	(Chile)
Iron and steel	22	-6.2	-5.2	-29.0	(Chile)	-3.3	(France)
Non-ferrous metal	19	-6.2	-5.3	-13.1	(UK)	-3.0	(Portugal)
Fabricated metal pdcts	26	-7.7	-5.6	-26.6	(Indonesia)	-3.5	(France)
Machineries	23	-9.0	-6.7	-29.0	(Cyprus)	-4.9	(USA)
Electrical apparatus	27	-10.5	-8.6	-29.0	(Cyprus, Hungary, Chile)	-4.7	(Germany)
Transport equipment	25	-11.4	-8.0	-29.0	(China, Turkey, Canada)	-4.1	(Spain)
Measuring equipment	17	-13.5	-10.9	-29.0	(Sweden, Guatla)	-3.4	(Finland)
Other manufacturing	20	-8.0	-6.2	-29.0	(Portugal)	-3.1	(Norway)

Notes: The table reports summary statistics on the estimated elasticities, $\hat{\varepsilon}_j^s$, by ISIC-rev2 industry. The countries displaying the minimum and maximum elasticities in each sector are displayed under parentheses.

Table 4: Aggregate welfare gains, domestic expenditure shares, and trade elasticities

Country	Domestic share (λ_{jj}) (1)	Elasticity (ε_j) (2)		Welfare ($d\ln W_j^{MS}$) (3)		Aggregate Elast (ε_j^A) (4)	
Australia	0.690	-6.466	(.810)	-0.057	(.007)	-1.944	(.220)
Austria	0.522	-4.831	(.238)	-0.135	(.007)	-2.199	(.192)
Canada	0.577	-8.018	(.630)	-0.069	(.005)	-1.442	(.106)
Chile	0.647	-8.999	(.609)	-0.048	(.003)	-3.578	(.855)
China	0.810	-6.920	(.433)	-0.030	(.002)	-2.972	(.247)
Cyprus	0.532	-9.989	(.648)	-0.063	(.004)	-3.089	(.419)
Finland	0.644	-5.040	(.187)	-0.087	(.003)	-2.368	(.307)
France	0.701	-4.624	(.118)	-0.077	(.002)	-2.158	(.158)
Germany	0.691	-4.312	(.095)	-0.086	(.002)	-2.112	(.170)
Greece	0.487	-9.206	(1.509)	-0.078	(.013)	-4.825	(.898)
Guatemala	0.585	-4.486	(.428)	-0.120	(.011)	-3.950	(1.046)
Hong Kong	0.167	-4.932	(.262)	-0.363	(.019)	-2.081	(.193)
Hungary	0.604	-6.381	(.319)	-0.079	(.004)	-3.442	(.618)
India	0.898	-4.755	(.198)	-0.023	(.001)	-2.859	(.344)
Indonesia	0.596	-4.401	(.456)	-0.118	(.012)	-2.093	(.238)
Italy	0.698	-5.314	(.156)	-0.068	(.002)	-2.990	(.278)
Japan	0.935	-5.420	(.164)	-0.012	(.000)	-3.030	(.138)
Korea	0.777	-5.113	(.239)	-0.049	(.002)	-4.895	(.805)
Malaysia	0.382	-3.418	(.375)	-0.282	(.031)	-2.516	(.239)
Norway	0.615	-5.363	(.317)	-0.091	(.005)	-3.165	(.417)
Portugal	0.608	-6.417	(.610)	-0.078	(.007)	-3.197	(.345)
Slovakia	0.716	-7.781	(.938)	-0.043	(.005)	-3.879	(.711)
Spain	0.730	-5.035	(.214)	-0.062	(.003)	-2.634	(.233)
Sweden	0.545	-6.044	(.260)	-0.100	(.004)	-2.314	(.240)
Taiwan	0.707	-5.243	(.315)	-0.066	(.004)	-1.988	(.209)
Turkey	0.771	-6.491	(.315)	-0.040	(.002)	-3.351	(.216)
United Kingdom	0.662	-5.389	(.188)	-0.076	(.003)	-3.617	(.369)
United states	0.837	-4.907	(.160)	-0.036	(.001)	-1.463	(.072)

Notes: The table reports the calibrated aggregate share of domestic goods in consumption (Column (1)), the aggregate elasticity of imports computed using equation (8) (column (2)), the welfare impact of moving to autarky implied by this elasticity (column (3)) and the aggregate elasticity estimated from aggregate data (column (4)). Standard errors in parentheses, obtained using the Delta method described in appendix. All estimates are significant at the one percent level.

Figure 3: Sources of heterogeneity in one-sector elasticities

(a) Elasticity gap with the US

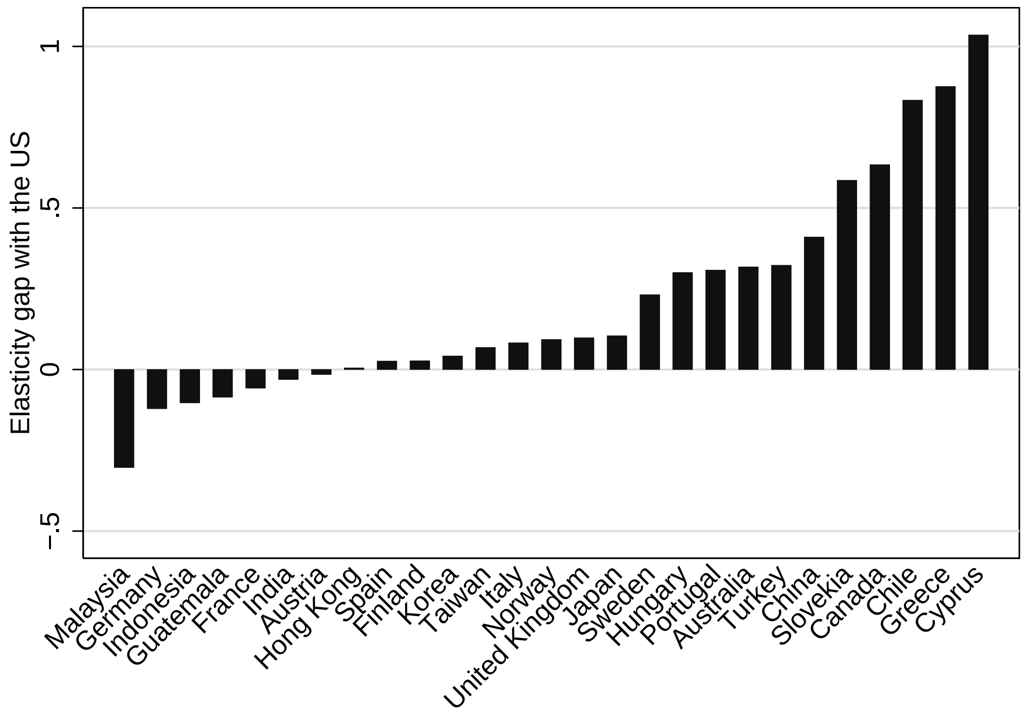

(b) Decomposition of the elasticity gap

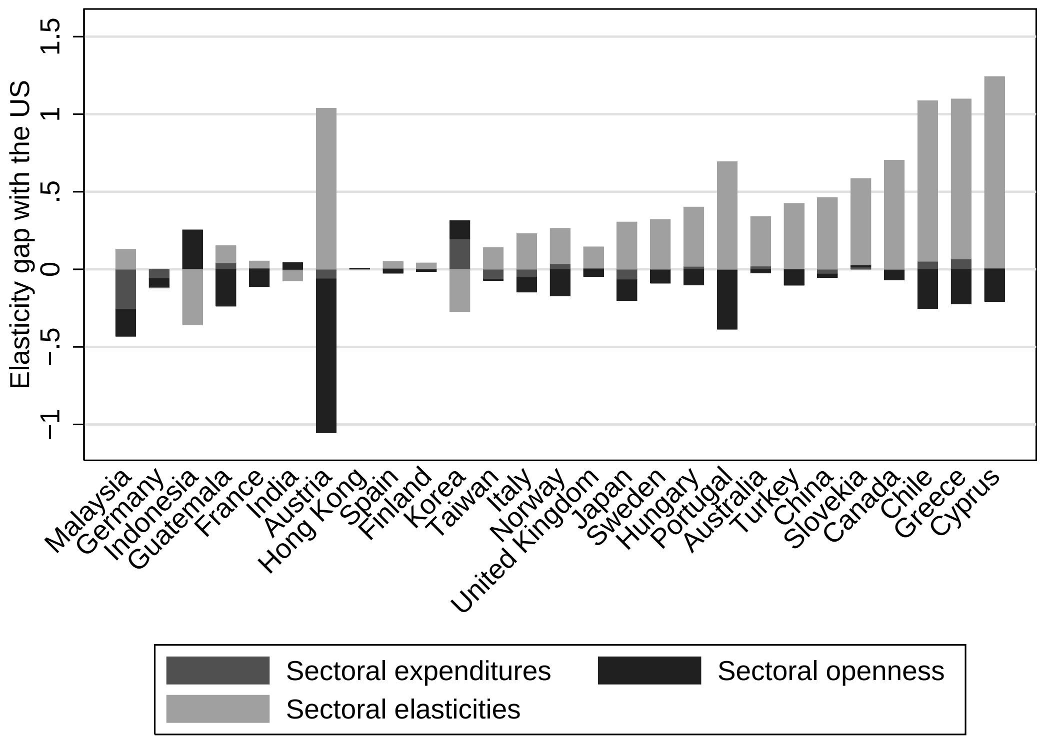

Note: Panel (a) depicts each country's estimated aggregate elasticity, in deviation with respect to the US ($\hat{\varepsilon}_j - \hat{\varepsilon}_{US}$). Panel (b) then reports the decomposition in equation (13). Medium gray corresponds to the first term B_j, black corresponds to the second term L_j, and light grey corresponds to the third term E_j.