EUROPEAN CENTRAL BANK

EUROSYSTEM

WORKING PAPER SERIES

NO 1535 / APRIL 2013

TRADE ADJUSTMENT IN THE EUROPEAN UNION

A STRUCTURAL ESTIMATION APPROACH

Vesna Corbo and Chiara Osbat

In 2013 all ECB publications feature a motif taken from the €5 banknote.

THE COMPETITIVENESS RESEARCH NETWORK

This paper presents research conducted within the Competitiveness Research Network (CompNet). The network is composed of economists from the European System of Central Banks (ESCB) - i.e. the 27 national central banks of the European Union (EU) and the European Central Bank – a number of international organisations (World Bank, OECD, EU Commission) universities and Think-tanks, as well as a number of non-European Central Banks (Argentina and Peru) and organisations (US International Trade Commission). The objective of CompNet is to develop a more consistent analytical framework for assessing competitiveness, one which allows for a better correspondence between determinants and outcomes.

The research is carried out in three workstreams: 1) Aggregate Measures of Competitiveness; 2) Firm Level; 3) Global Value Chains

CompNet is chaired by Filippo di Mauro (ECB). The three workstreams are headed respectively by Chiara Osbat (ECB), Jean-Charles Bricongne (Banque de France) and João Amador (Banco de Portugal). Christina Werner (ECB) is responsible for the CompNet Secretariat.

The refereeing process of this paper has been coordinated by Editorial Board of the ECB Working Paper Series, lead by Philipp Hartmann.

The paper is released in order to make the research of CompNet generally available, in preliminary form, to encourage comments and suggestions prior to final publication. The views expressed in the paper are the ones of the author(s) and do not necessarily reflect those of the ECB, the ESCB, and of other organisations associated with the Network.

Acknowledgements
The authors wish to thank Alexander Chudik, Nils Gottfries, Joakim Gullstrand, Jean Imbs, Filippo di Mauro, Alistair Dieppe, Andreas Westermark, seminar participants at the European Central Bank and Uppsala University, and the participants at the National Meeting of Swedish Economists 2011. Vesna Corbo thanks the ECB for the hospitality and financial support, as this project was initialized while Corbo was visiting the ECB. The opinions expressed in this paper are the authors' only and do not reflect those of the ECB or of the Eurosystem.

Vesna Corbo
UppsalaUniversity

Chiara Osbat
European Central Bank; e-mail: chiara.osbat@ecb.europa.eu

Address	Kaiserstrasse 29, 60311 Frankfurt am Main, Germany
Postal address	Postfach 16 03 19, 60066 Frankfurt am Main, Germany
Telephone	+49 69 1344 0
Internet	http://www.ecb.europa.eu
Fax	+49 69 1344 6000

ISSN	1725-2806 (online)
EU Catalogue No	QB-AR-13-032-EN-N (online)

This paper can be downloaded without charge from http://www.ecb.europa.eu or from the Social Science Research Network electronic library at http://ssrn.com/abstract_id=2246577.

Information on all of the papers published in the ECB Working Paper Series can be found on the ECB's website, http://www.ecb.europa.eu/pub/scientific/wps/date/html/index.en.html

Abstract

We estimate the elasticity of substitution of a country's imports, and that of its exports on the world market, for EU countries using sector level trade data. We present a new empirical strategy based on the identification scheme by Feenstra (1994), which enables the estimation of elasticities from data on exports. Moreover, our use of bootstrap methods allows us to obtain better elasticity measures, and to better characterize their accuracy. Our results show much heterogeneity in the estimates of the elasticity of substitution across industrial sectors. This, in turn, points to heterogeneity across countries, due to different production and trade structures. We obtain aggregate elasticities for the EU27 countries, with a mean of 3.5 for imports and 4.0 for exports, bringing us closer to traditional estimates and bridging the gap between the newer micro data estimates and the more traditional estimates found in the macroeconomic literature.

Keywords: Elasticity of substitution; Heterogeneity; Aggregation; Calibration of macroeconomic models

JEL classification: C23, F14, F47

Non-Technical Summary

The responsiveness of trade volumes to changes in relative prices, or elasticity of substitution, is a parameter of interest in international trade and policy and its measurement has been an intensely researched and hotly debated topic since the second world war, starting with articles such as Tinbergen (1946), Polak (1950) and Orcutt (1950). Despite this extensive literature, however, little consensus has been reached on the magnitudes of elasticities of substitution and trade elasticities. While the earlier time series literature mainly arrives at low, often insignificant values (an outcome dubbed "elasticity pessimism"), a newer branch that uses more disaggregated data and panel-based econometric methods tends to obtain considerably higher values (hence the "elasticity optimism" of e.g. Imbs and Méjean (2009)). This paper is motivated by this more recent literature and aims at estimating the elasticity of substitution for EU countries, at a rather disaggregated sectoral level.

We make three main contributions. First, we develop a new empirical strategy, in order to be able to identify the elasticity of exported goods on the world market. It is closely related to the identification scheme proposed by Feenstra (1994), which however is designed for estimation on import data only. Our new identification strategy enables us to take advantage of the export dimension of our data to obtain the elasticity of substitution of a declarant's exports – a measure which, albeit often neglected, is important for the calibration of open-economy macro models. Thus, we estimate the elasticity of substitution of a country's exports, without having access to the complete set of import data for each of the country's trading partners. Given that the quality and availability of data notably differs between countries, it is an advantage not to have to depend on as many sources as there are trading partners. Instead, we are able to use the more reliable and complete export data sets reported by each of our countries of interest.

Second, we make a methodological advancement in applying bootstrap methods to obtain less biased and more robust elasticity measures. As shown in Corbo and Osbat (2012), due to the Feenstra method involving a mapping from estimated reduced-form parameters into the structural parameters of interest via a non-linear mapping function, the obtained estimates are on average biased in large parts of the parameter space. Relying instead on the mode or median of the bootstrap distribution of our estimates considerably reduces this bias.

Our third contribution relates to the coverage and completeness of our results. Compared to many earlier studies, our database and the use of the bootstrap, which improves small-sample properties of the estimator, allow us to cover a much larger percentage of total trade and production. The data we use for estimation are from Eurostat's COMEXT database, which contains monthly observations on values and quantities of imports and exports reported by all European Union countries from and to up to 270 trading partners at a disaggregation level of up to 8 digits. For the majority of the countries in our sample, we span the period from 1995 to 2009, although for the newer EU members the time span is shorter. The level of disaggregation at which we estimate the sectoral elasticities is the 4-digit ISIC. The estimates are then properly aggregated into macro level estimates for each country, which are suitable for calibration purposes.

Our results show much heterogeneity in the estimates of the elasticity of substitution across industrial sectors, which in turn points to heterogeneity across countries, due to different production and trade structures. The aggregate elasticities that we obtain are lower than the ones presented in recent related studies, such as Imbs and Méjean (2009); we argue that the high estimates may partly be an artefact of the estimation method employed. Our aggregated import data elasticities have a mean of 3.5 and a median of 3.4, while the aggregated export data estimates have a mean of 4.0 and a median of 3.8. Accounting for the differences in the definitions of elasticities in micro data estimations and macroeconomic models, these values are in line with values used for calibration, albeit being in the high end. We are hence able to partly explain the large discrepancies in the estimates and bridge between the two strands of literature.

1 Introduction

The responsiveness of trade volumes to changes in relative prices is one of the central quantities of interest in international trade and policy. It is crucial, for example, in analyzing the responsiveness of trade to tariffs and the adjustment of external imbalances via changes in real exchange rates. It also relates to the concept of competitiveness, as a low elasticity may be indicative of low substitutability, possibly because of high technological content. The measurement of the elasticity of substitution has been an intensely researched and hotly debated topic since the Second World War, starting with articles such as Tinbergen (1946), Polak (1950) and Orcutt (1950). Despite the extent of the empirical literature since the 1940s, little consensus has been reached on the magnitudes of elasticities of substitution and trade elasticities. While the earlier time series literature mainly arrives at low, often insignificant values, a newer branch that uses more disaggregated data and panel-based econometric methods tends to obtain considerably higher values. Moreover, even though there is a fair amount of literature on the topic, estimates of the elasticity of substitution for most European countries are scarce and, even when they do exist, often outdated.

The aim of this paper is to provide reliable measures of aggregate elasticities of substitution for the EU countries. We do so for both the substitution of imported for domestically produced goods, and the substitution of a country's exports on the world market. Our methodology is motivated by the more recent literature; we perform our estimation at a rather disaggregated sectoral level following Feenstra (1994), with the intention of avoiding the endogeneity and aggregation biases present in demand equation estimations on aggregate data. In order to be able to identify the elasticity of exported goods on the world market, we develop a new empirical strategy. It is closely related to the identification scheme proposed by Feenstra (1994), which however is designed for estimation on import data only. As our data is such that a smaller set of countries act as declarants of both imports from, and exports to a much larger number of trade partners, we do not have the option of estimating the import data elasticities of all of the trading partners, and aggregating those into an elasticity of substitution of a declarant's exports. Instead, our identification strategy enables us to take advantage of the export dimension of our data to obtain these elasticities. We subsequently aggregate the sector level elasticities for both imports and exports into macro level elasticities of substitution for each EU27 country.

Thus, we estimate the elasticity of substitution of a country's exports, without having access to the complete set of import data for each of the country's trading partners.

Given that the quality and availability of data notably differs between countries, it is an advantage not to have to depend on as many sources as there are trading partners. Instead, we are able to use the more reliable and complete export data sets reported by each of our countries of interest. While often neglected in discussions, the obtained measure is important in the calibration of e.g. two-country models. While the calibration of the "domestic" country's elasticity of substitution is more frequently debated, little focus is placed on the other economy. There is, in theory, no reason to believe that a country's imports and exports are equally substitutable, and it is hence relevant to obtain reliable estimates of both measures. Here we provide the complete picture of the trade dynamics parameters relevant for calibration and policy analysis for each country under study, studying both the inflows and the outflows of goods. In other words, in the context of a two-country macroeconomic model, we estimate both the structural elasticity of substitution of the country under study, and the same elasticity of the foreign country, often assumed to represent the rest of the world.

We further make a methodological advancement in applying bootstrap methods to obtain less biased and more robust elasticity measures. As shown in Corbo and Osbat (2012), due to the Feenstra method involving a mapping from estimated reduced-form parameters into the structural parameters of interest via a non-linear mapping function, the obtained estimates are on average biased in large parts of the parameter space. Relying instead on the mode or median of the bootstrap distribution of our estimates considerably reduces this bias, which emerges in the actual mapping and is thus present even if the reduced-form parameters have been consistently estimated.

Compared to many earlier studies, our database and the use of the bootstrap allow us to cover a much larger percentage of total trade and production. The data we use for estimation are from Eurostat's COMEXT database, which contains monthly observations on values and quantities of imports and exports reported by all European Union countries from and to up to 270 trading partners at a highly disaggregated level. For the majority of the countries in our sample, we span the period from 1995 to 2009; for the newer members, however, the time span is shorter. The level of disaggregation at which we estimate the sectoral elasticities is the 4-digit ISIC, as this is the highest level of disaggregation at which we can obtain aggregation weights. We aggregate the sector-level estimates, using weighting schemes derived from theory and country-specific weights, obtaining a reliable macro level estimate for each country.

Our results show much heterogeneity in the estimates of the elasticity of substitution across industrial sectors, which in turn points to heterogeneity across countries, due to different production and trade structures. The aggregate elasticities that we obtain are lower than the ones presented in recent related studies, such as Imbs and Méjean (2009); we argue that the high estimates may partly be an artefact of the employed estimation method. Our aggregated import data elasticities have a mean of 3.5 and median of 3.4, while the aggregated export data estimates have a mean of 4.0 and a median of 3.8. Accounting for the differences in the definitions of elasticities in micro data estimations and macroeconomic models, these values are in line with values used for calibration, albeit in the high end. Thus we are able to partly explain the above discussed large discrepancies in the estimates and bridge the gap between the two strands of literature.

The rest of the paper is structured as follows. Section 2 contains a short overview of the literature related to our study. In Section 3, we briefly discuss the methodology developed in Feenstra (1994), which is also the one we use and extend in this paper. The extension to export data estimation is presented in Section 4. Section 5 discusses data and estimation methods applied. In Section 6, we present and discuss our aggregation methodology. Section 7 contains the results, starting with the microeconomic estimates and proceeding with the outcomes of the aggregation. The section is completed with a discussion of the calibration of macroeconomic models. Finally, we conclude in Section 8.

2 Overview of related literature

In the last two decades, literature employing panel data methods for the estimation of elasticities of international substitution and related concepts in international economics has emerged. It makes use of the increasingly available sources of micro level trade data to produce structural estimates of elasticities, at the good, sector or macroeconomic levels. These studies follow the methodology laid out in the seminal paper by Feenstra (1994), who introduces a way of structurally estimating the elasticities of substitution by explicitly modeling the supply side in addition to the demand side of the economy. Under the assumption of equal substitutability of all varieties independent of their origin, as in Armington (1969), the resulting system of equations can be estimated using only trade data. While older studies, dating as far back as the 1940s, struggled with the endogeneity problem present whenever quantities or volumes are regressed on prices, these new estimation methods are able to circumvent this by making use of multi-dimensional data. Furthermore, the use of micro level data significantly reduces any aggregation bias problems related to the use of aggregate time series.[1]

[1]An overview of the earlier literature employing macro time series for the estimation of trade elasticities and elasticities of substitution is beyond the scope of this paper. We refer to the summary in

Feenstra (1994) presents elasticities of substitution for a total of eight goods, ranging from 2.96 for typewriters to 42.9 for silver bullion.[2] Subsequent papers extend the exercise to comprise a much wider set of sectors. Broda and Weinstein (2006) estimate almost $30,000$ elasticities of substitution for US imports at the highest level of disaggregation, the unweighted median and mean of which are 3.1 and 12.6, respectively, over the period 1990-2001. In Broda, Greenfield, and Weinstein (2006), median estimates of the sectoral elasticities are presented country by country for more than 70 countries, 21 of which are EU members and can hence be compared to our estimates.

Broda and Weinstein (2006) further make the methodological contribution of extending the Feenstra method to deal with estimates that do not yield theory-consistent elasticities of substitution. Despite the large coverage of the Broda and Weinstein data, it is hard to interpret their estimates in a macroeconomic context, since they have not been properly aggregated. A method of performing theoretically consistent aggregation can instead be found in Imbs and Méjean (2009), who estimate and aggregate elasticities of substitution for 56 US sectors. This yields an aggregate estimate of around 7, which is also the value that the authors recommend for calibration of macroeconomic structural models. In their (2010) paper, Imbs and Méjean present estimates for a larger number of countries, among those 14 Euro Area (EA) countries. These are however aggregated into macro level trade elasticities, rather than elasticities of substitution, and thus do not have a clear theoretical interpretation.

Most of the above listed papers, as well as a large number of earlier studies, use US data in their estimations. For European countries the literature is more scarce, both when it comes to panel data and to time series data estimations. In fact, there are no studies, to the best of our knowledge, that provide aggregate elasticity estimates for all countries in the European Union. This lack of reliable estimates for calibration and policy evaluation is particularly pronounced for newer members of the European Union, for which the data is often scarce and spans over short periods of time. One recent study that does use data on all EU countries is Mohler and Seitz (2010). The authors do a similar exercise to the one in Broda and Weinstein (2006), with the goal of evaluating the gain from variety. To that end, they estimate a large number of disaggregated elasticities, with median values ranging from 3.4 for Greek imports to 4.9 for Romanian

McDaniel and Balistreri (2003), who discuss some of the main references in the field, among those Gallaway, McDaniel, and Rivera (2003). Other frequently cited references include Orcutt (1950), Houthakker and Magee (1969) and Marquez (1990).

[2] The elasticity of 2.96, for example, should be interpreted as the reaction of relative import quantities to a one percentage change in the relative price. In other words, if country v's price of typewriters imported by country c increases by 1% relative to the average price of typewriters consumed in country c, then the relative consumption of country-v typewriters in country c will decrease by 2.96%.

ones. The mean values lie between 7.9 and 290 for the different countries, but, just as in the case of Broda and Weinstein (2006), these estimates are hard to interpret in a macroeconomic context without proper aggregation.

3 Estimation using the Feenstra (1994) method

The main advantage of the Feenstra (1994) approach is that it explicitly models the supply side of the economy, in addition to the standard demand side specifications traditionally used when estimating elasticities of substitution. Identification is achieved exploiting the panel structure of the data; this way the simultaneity issues present whenever quantities are regressed on prices are avoided. We will not go into much detail on the Feenstra approach, as it has been presented extensively elsewhere in the literature. Here, we only present the underlying assumptions and the resulting estimation equation, which are needed for the understanding of the new empirical strategy that enables estimation using export data, and the discussion of the empirical results to come.

We begin by clarifying some concepts and definitions, which will be used throughout the paper, and commenting briefly on notation. The disaggregation level at which we compute and estimate the elasticities of substitution is the good, or equally, sector level. We use the terms *sector* and *good* interchangeably; in other words, the sector grouping is what determines the definition of a good. A good, in turn, comprises a number of varieties. By a *variety*, we here refer to a good originating from some country and meant to be consumed in some (possibly different) country. Hence, the definition of good is based on some product characteristics other than its origin and destination, and variety refers to products of a specific origin to some specific destination belonging to some category of goods. There will thus be as many varieties of each good as there are trading partners in that specific sector. The exact empirical definitions of goods and varieties will depend on the disaggregation level and availability of the data, and will be discussed further in Section 5 below. To make the formulas easily readable, we have set all indices to the first letter of what they are meant to index: we denote the country under study by c, the good (or sector) by g, the variety (or country producing good g) by v, and the time period by t. If a variable denotes the sum or aggregate across some of the indices, a dot will appear in the place of the index across which we are aggregating. In the case of parameters being assumed constant across some dimension, they will only be denoted by the indices across which they vary. Finally, the entire set of each lower-case letter will be denoted by its corresponding capital letter.

We start from the standard CES setting for the modeling of demand,

$$C_{cgvt} = \beta_{cgvt}^{\sigma_{cg}-1} \left(\frac{P_{cgvt}}{P_{cg.t}} \right)^{-\sigma_{cg}} C_{cg.t} \,, \tag{1}$$

and a simple supply structure, as given by

$$P_{cgvt} = \tau_{cgvt} \exp(\nu_{cgvt}) C_{cgvt}^{\omega_{vg}} \,. \tag{2}$$

Here, C_{cgvt} denotes the consumption of variety v of good g in country c at time t and $C_{cg.t}$ the total consumption of good g in country c at time t, and P_{cgvt} and $P_{cg.t}$ are the corresponding prices. The parameters β_{cgvt}, τ_{cgvt}, and ν_{cgvt} are a taste parameter, trade cost, and technology shock, respectively, all specific to each variety, good and country at each point in time. The parameter σ_{cg} denotes the elasticity of substitution of good g in country c, and is the main parameter of our interest, while $\omega_{vg} = \omega_{cg} \geq 0$, for all v, is the inverse of the price elasticity of supply of good g in country c, both assumed equal across varieties but allowed to differ between goods. Note that, not only are the elasticities of substitution assumed to be equal over imported varieties, but also between the imported and domestically produced varieties of good g. This is the Armington assumption, discussed above, upon which the identification strategy hinges.[3]

Denoting by $Z_{cgvt} \equiv P_{cgvt}C_{cgvt}$ the expenditures on variety v of good g in country c at time t, and substituting equation (1) into equation (2), we can take logs and difference the resulting system of equations to obtain

$$\begin{aligned}
\Delta \ln Z_{cgvt} &= (1-\sigma_{cg})\Delta \ln P_{cgvt} + \varepsilon_{cgvt} - (1-\sigma_{cg})\Delta \ln P_{cg.t} + \Delta \ln Z_{cg.t} && (3)\\
\Delta \ln P_{cgvt} &= \frac{\omega_{cg}}{1+\sigma_{cg}\omega_{cg}}\varepsilon_{cgvt} + \delta_{cgvt} \\
&\quad + \frac{\sigma_{cg}\omega_{cg}}{1+\sigma_{cg}\omega_{cg}}\Delta \ln P_{cg.t} + \frac{\omega_{cg}}{1+\sigma_{cg}\omega_{cg}}\Delta \ln C_{cg.t} \,, && (4)
\end{aligned}$$

[3]The Armington (1969) assumption was originally introduced to simplify the modeling of demand functions in the presence of a large number of products in a market. In addition to simplifying, it has the advantage of allowing us to measure the elasticity using trade data only, the quality and level of disaggregation of which far exceed those of other sources of data. The assumption has become rather standard in the empirical trade literature, a tradition that we will follow, and leave further discussions regarding its plausibility aside. As we are ultimately interested in the substitutability between domestically produced and imported goods, and will estimate the elasticities of substitution under the Armington assumption, we only note that, if anything, our estimates should be upward biased. This would be the case if there, contrary to the Armington assumption, was some home bias. As we argue that the elasticities of interest for calibration of macroeconomic models may have been overestimated in some of the previous micro-level studies, and that their true value should in fact be lower, our point is only reinforced in the presence of an upward bias stemming from the Armington assumption.

where

$$\varepsilon_{cgvt} \equiv (\sigma_{cg}-1)\Delta\ln\beta_{cgvt} \tag{5}$$

$$\delta_{cgvt} \equiv \frac{1}{1+\omega_{cg}\sigma_{cg}}\left[\Delta\ln\tau_{cgvt}+\Delta\nu_{cgvt}\right] \tag{6}$$

denote the idiosyncratic error terms in the demand and supply equations, respectively. The last two terms in each of the equations (3) and (4) are common across all varieties of good g imported by country c. By subtracting from each equation indexed v the same equation for some reference variety v_r, we can eliminate these terms, obtaining

$$\begin{aligned}
\breve{\varepsilon}_{cgvt} &\equiv \varepsilon_{cgvt}-\varepsilon_{cgv_rt} \qquad (7)\\
&= (\sigma_{cg}-1)\left[\Delta\ln P_{cgvt}-\Delta\ln P_{cgv_rt}\right]+\left[\Delta\ln Z_{cgvt}-\Delta\ln Z_{cgv_rt}\right]\\
\breve{\delta}_{cgvt} &\equiv \delta_{cgvt}-\delta_{cgv_rt} \qquad (8)\\
&= \frac{1+\omega_{cg}}{1+\omega_{cg}\sigma_{cg}}\left[\Delta\ln P_{cgvt}-\Delta\ln P_{cgv_rt}\right]\\
&\quad -\frac{\omega_{cg}}{1+\omega_{cg}\sigma_{cg}}\left[\Delta\ln Z_{cgvt}-\Delta\ln Z_{cgv_rt}\right].
\end{aligned}$$

In order to make identification of this system of equations possible, we need to make the assumption of no correlation between the error terms, i.e.

$$\mathrm{E}(\varepsilon_{cgvt}\delta_{cgvt})=0\,. \tag{9}$$

Given (9), we can multiply the above equations for the differenced error terms to obtain an expression for the i.i.d. variable u_{cgvt} which is a function of expenditure shares and prices only. Rearranging, we end up with the following estimation equation,

$$Y_{cgvt}=\theta_{1cg}X_{1cgvt}+\theta_{2cg}X_{2cgvt}+u_{cgvt}\,, \tag{10}$$

where we have defined:

$$Y_{cgvt}\equiv\left[\Delta\ln P_{cgvt}-\Delta\ln P_{cgv_rt}\right]^2\,, \tag{11}$$

$$X_{1cgvt}\equiv\left[\Delta\ln Z_{cgvt}-\Delta\ln Z_{cgv_rt}\right]^2\,, \tag{12}$$

$$X_{2cgvt}\equiv\left[\Delta\ln Z_{cgvt}-\Delta\ln Z_{cgv_rt}\right]\cdot\left[\Delta\ln P_{cgvt}-\Delta\ln P_{cgv_rt}\right]\,, \tag{13}$$

$$u_{cgvt}\equiv\frac{1+\omega_{cg}\sigma_{cg}}{(1+\omega_{cg})(\sigma_{cg}-1)}\,\breve{\varepsilon}_{cgvt}\breve{\delta}_{cgvt}\,, \tag{14}$$

$$\theta_{1cg} \equiv \frac{\omega_{cg}}{(1+\omega_{cg})(\sigma_{cg}-1)}\,, \tag{15}$$

and

$$\theta_{2cg} \equiv \frac{\omega_{cg}\sigma_{cg} - 2\omega_{cg} - 1}{(1+\omega_{cg})(\sigma_{cg}-1)}\,. \tag{16}$$

Note that the definitions of the estimation parameters θ_{1cg} and θ_{2cg} compose a system of two equations in two unknowns, namely the structural elasticity parameters σ_{cg} and ω_{cg}. Here we also emphasize, for future reference, that these expressions are non-linear in σ_{cg} and ω_{cg}. In fact, as will become clear from the parameter restrictions presented in Section 5.2, the function mapping θ_{1cg} and θ_{2cg} into σ_{cg} and ω_{cg} is not only non-linear but also discontinuous in parts of the parameter space.

To sum up, the variables we use in the estimation are the second moments of expenditures and prices, once they have been differenced over time and with respect to a reference variety in order to eliminate any level dependence or shocks common across varieties. Based on the estimates obtained from running a regression on those variables, we are able to induce the structural parameters of our interest.

4 Identification using a cross-section of exports from, instead of imports to, the country of interest

Due to the structure of our data, where a small cross-section of countries are acting as declarants and a much larger cross-section of countries are registered as trade partners, we do not have the liberty of estimating elasticities of substitution for each import partner and aggregating those into a measure of the export elasticity of substitution of our country of interest.[4] Even if we did have access to data from all trade partners, one may dispute whether the data quality and availability would be comparable among all the countries. Instead, we have at our disposal a large cross-section of importers for each declaring exporting country. To be able to make use of this data for estimation, however, we first need to develop an alternative identification strategy, since the Feenstra (1994) method only applies to a cross-section of exporters to a declaring importing country. In

[4]This is what Imbs and Méjean (2010) do, using the BACI data set where both imports to and exports from a large cross-section of countries is recorded. In the COMEXT data that we use for estimation, only the EU countries act as declarants reporting both their imports and exports, but we cannot observe the full trade of all of their trading partners. This drawback of our data set notwithstanding, the advantages compared to the BACI data are considerable; the COMEXT data is much more detailed and spans over a longer period of time. Moreover, we have to rely on fewer sources of data; this is an advantage given that the quality of data may be poorer in some countries than others, which in turn may imply a risk of biased estimates.

the present section, we present a way of doing so, drawing heavily on the ideas underlying Feenstra (1991, 1994).

Assume now that the preference structure in all countries is the same, i.e. a CES structure as in equation (1). We still allow for the prices and consumption shares to differ, assuming that the realization of shocks is specific to each country. Just as before, we define a variety as a good originating from some country and meant to be consumed in some possibly different country. Now, however, we will focus on the set of goods produced in and exported from the country of interest c to a cross-section of importing countries.[5] Our interest lies in finding the elasticity of substitution of the good g produced in country c on the world market, i.e. we are not interested in any one country v's elasticity of substitution of its imports of good g, but rather the general substitutability of all exports of good g from country c to the rest of the world. Imagining a structural small-open-economy or two-country model, where our country of interest is the small open economy or one of the modeled countries, respectively, and the rest of the world is modeled as *the other country*, the measure that we are after is *the other country's* elasticity of substitution between imported and domestically produced goods. This measure is relevant for the calibration of macroeconomic models; there is, however, very little information available about its magnitude for most, if not all countries.

4.1 The demand and supply equations

Since any country v potentially imports good g from lots of exporters, the variety coming from our exporter of interest c will be only one among many varieties consumed in v. We denote by i any single trading partner of country v, be it or not country c, and by I the full set of v's trading partners.

Just as in the import data derivations, the demand function for variety $i = c$ of good g in country v at time t is

$$C_{vgct} = \beta_{vgct}^{\sigma_{vg}-1} \left(\frac{P_{vgct}}{P_{vg.t}} \right)^{-\sigma_{vg}} C_{vg.t} \,, \tag{17}$$

[5] Note that, for the goods exported to different destinations from country c to qualify as different varieties, it is enough that the good is packaged differently, or that any other of its characteristics differs. Furthermore, since the disaggregation level we have in mind in our setting is the sector rather than the single firm, our definition of good encompasses a potentially large number of brands. The plausibility of our definition of variety is then further strengthened if a different composition of brands is exported to different destinations, or if the relative shares of the exported brands differ across countries of destination, which is likely to be the case in practice.

and the supply function is given by

$$P_{vgct} = \tau_{vgct} \exp(\nu_{vgct}) C_{vgct}^{\omega_{cg}} . \tag{18}$$

While the observed import prices were measured CIF (Cost, Insurance and Freight), on the export side we observe prices that are measured FOB (Free On Board). We therefore define the observed price[6]

$$\tilde{P}_{vgct} \equiv P_{vgct}^{FOB} = \frac{P_{vgct}}{\tau_{vgct}} \tag{19}$$

as the model price net of trade costs. Denoting the expenditures in an analogous way, so that $\tilde{Z}_{vgct} \equiv \tilde{P}_{vgct} C_{vgct} = Z_{vgct}/\tau_{vgct}$, taking logs and differencing, we can write the demand equation (17) as[7]

$$\Delta \ln \tilde{Z}_{vgct} = \varepsilon_{vgct} + (1 - \sigma_{vg}) \Delta \ln \tilde{P}_{vgct} + (\sigma_{vg} - 1) \Delta \ln P_{vg.t} + \Delta \ln Z_{vg.t} , \tag{20}$$

where Δ again denotes the first-order time difference, and where the error term of the demand equation is now defined as

$$\varepsilon_{vgct} \equiv (\sigma_{vg} - 1) \Delta \ln \beta_{vgct} - \sigma_{vg} \Delta \ln \tau_{vgct} . \tag{21}$$

Similarly, on the supply side, we have

$$\begin{aligned} \Delta \ln \tilde{P}_{vgct} \quad = \quad & \frac{\omega_{cg}}{1 + \omega_{cg}\sigma_{vg}} \varepsilon_{vgct} + \delta_{vgct} \\ & + \frac{\omega_{cg}\sigma_{vg}}{1 + \omega_{cg}\sigma_{vg}} \Delta \ln P_{vg.t} + \frac{\omega_{cg}}{1 + \omega_{cg}\sigma_{vg}} \Delta \ln C_{vg.t} , \end{aligned} \tag{22}$$

[6]FOB measures the cost of an import at the point of being loaded to a carrier for transport from the exporting to the importing country (Radelet and Sachs, 1998). Hence, the main carriage costs are not included in the measure. In the CIF measure, insurance, handling, and shipping costs are all included. To the extent our model provides a good description of the markets under study and our assumption on the trade cost being variety-specific is reasonable, the FOB/CIF distinction will not pose a problem for the accuracy of our estimates, as any trade cost will end up in the error term of the estimation equation.

[7]We previously worked with expenditure shares, instead of expenditures, as this is the variable used in Feenstra's original paper. On the import side, this distinction is irrelevant, as the denominator in the expenditure shares is eliminated in the differencing with respect to a reference variety. On the export side, we however observe expenditure shares as a function of the total export value of country c, and not the import value of country v. Writing the equations in terms of shares then yields a different estimation equation, as we are forced to do an additional differencing in order to eliminate the exporter-specific terms. The estimation results between the two methods differ only to the extent that the total export values from the exporter under study and from a reference exporter evolve differently over time. We have compared the results applying both of these methods, only to discover that there are nearly identical. We choose to present the estimation based on expenditures, as this method is simpler and more intuitive.

where the error term of the supply equation is now given by

$$\delta_{vgct} \equiv \frac{1}{1+\omega_{cg}\sigma_{vg}} \Delta\nu_{vgct} \,. \tag{23}$$

Comparing the equations (20) and (22) for country c to equations (3) and (4) for country v, we note two differences. The trade costs are now part of the error term of the demand equation, instead of the supply equation as was the case earlier. This is due to the export prices being measured FOB, and does not interfere with our identification strategy as long as we, just as before, assume that the trade costs are i.i.d. Moreover, as we have not yet imposed any restrictions on the elasticities, the supply and demand elasticities are still indexed differently. We deal with this in the following section. Note also that, to enable identification, we need not impose any further restrictions on the error terms. As in Section 3, we assume that the error terms in the demand and supply equations are uncorrelated, i.e. $E(\varepsilon_{vgct}\delta_{vgct}) = 0$.

4.2 Deriving the estimated equation

We now have a system consisting of a demand and a supply equation for each importing country, containing variables specific to country v but common to all its trading partners. Just as in the case of imports, we will make use of a reference variety to cancel out these terms. This will however require an additional dimension of data. The idea in Feenstra (1994) was to enable identification of in a time-series context unidentified parameters by using the panel dimension of the data. Here, the idea is to take this a step further and to use a three-dimensional panel in order to enable identification. We again introduce a reference exporter, but due to the structure of our data where, instead of the importer, the exporter acts as declarant, we need another cross-section dimension that will enable us to eliminate the unobserved importer-specific variables in our setting. This can be seen from equation (22), where the last two terms are specific to the importing country v, and hence would not disappear if we only used the cross-section of importers of variety c. This estimation strategy requires us to make slightly different identifying assumptions from the ones we made in Section 3. We discuss these assumptions next, before proceeding with the derivations.

Just as in the case of imports, we need to impose some restrictions on the elasticity parameters to enable identification. We again impose homogeneity across varieties, i.e. the elasticity of substitution and the supply elasticity are the same for all varieties of a good. We thus have that $\sigma_{vg} = \sigma^X_{cg}$, for all v and $\omega_{vg} = \omega^X_{cg}$, for all v, where the superscript X marks that the elasticity is specific to country c's exports. In both settings we assume that a variety is defined as a good originating from some country and meant to be consumed in some (possibly different) country. Note however that since, in the cases of estimation using import data and estimation using export data, the cross-sections of importers and exporters respectively span over very different geographical areas, it is not necessarily the case that the same assumptions will be equally plausible (or implausible, for that matter) in the two settings. In addition to the homogeneity-across-varieties assumption, using export instead of import data will force us to make an additional assumption, stemming from the introduction of a third panel dimension; we now assume that $\omega^X_{cg} = \omega^X_{c_rg}$ and $\sigma^X_{cg} = \sigma^X_{c_rg}$, where c_r is an optimally chosen exporter reference country. By *optimally chosen*, we here mean a country with similar characteristics and trading patterns as our country of interest. The additional assumption we are making is conceptually very close to the import assumptions made in Feenstra (1994), and it is not clear which set of assumptions is preferred to the other. For the import data estimations, on one hand we are assuming that the elasticity of substitution is the same *within a country* for imports of all origins. On the other hand, we assume that *all exporting countries* of a certain good have the same elasticity of supply. While this assumption is what enables identification and hence a necessary one, it is possible to think of examples when this assumption may be violated. For example, it is possible that two economies of very different size and endowments also differ with respect to their elasticities of supply, even if they are exporting the same good. For the estimations on export data, we are instead assuming that the elasticity of supply is the same for a good *within a country and as in one similar country of reference*. However, we now need to assume that the elasticity of substitution is the same *across importing countries*, an assumption perhaps as unlikely to be true as the supply elasticity being the same anywhere in the world. Given the available estimates of sectoral elasticities for a large number of countries, we can obtain a dispersion measure that can give us a clue about how problematic the latter assumption might be. Since the same is not possible for the supply elasticities, unfortunately, we can still say nothing regarding the assumptions' relative importance.

Note also that, even though we are primarily interested in the estimates of the elasticities of substitution rather than the supply elasticities, and it therefore might seem like a preferable option to impose restrictions on the supply elasticities rather than the elasticities of substitution, the only thing that will matter in the end is the severity of the restrictions, and not on which parameters they are imposed. Since we are not estimating any of the two elasticities directly, but obtaining estimates of the two coefficients θ_1 and

θ_2 which are combinations of the parameters of interest, in practice, any assumption that lowers the accuracy of these estimates will affect the elasticity of substitution equally, independently of which of the parameters it is theoretically restricting.

Moving on to the derivations of the estimated equation, we start by subtracting from equations (20) and (22) the same equations for the reference variety c_r.

$$\begin{aligned}
\breve{\varepsilon}_{vgct} &\equiv \varepsilon_{vgct} - \varepsilon_{vgc_rt} && (24)\\
&= (\sigma^X_{cg} - 1)\left[\Delta \ln \tilde{P}_{vgct} - \Delta \ln \tilde{P}_{vgc_rt}\right] + \left[\Delta \ln \tilde{Z}_{vgct} - \Delta \ln \tilde{Z}_{vgc_rt}\right]\\
\breve{\delta}_{vgct} &\equiv \delta_{vgct} - \delta_{vgc_rt} && (25)\\
&= \frac{1+\omega^X_{cg}}{1+\omega^X_{cg}\sigma^X_{cg}}\left[\Delta \ln \tilde{P}_{vgct} - \Delta \ln \tilde{P}_{vgc_rt}\right]\\
&\quad - \frac{\omega^X_{cg}}{1+\omega^X_{cg}\sigma^X_{cg}}\left[\Delta \ln \tilde{Z}_{vgct} - \Delta \ln \tilde{Z}_{vgc_rt}\right]
\end{aligned}$$

The terms, $\breve{\varepsilon}_{vgct}$ and $\breve{\delta}_{vgct}$ are independent by assumption. We can then again multiply the demand and supply equations to obtain an estimation equation where the residual term has an expected value of zero. A multiplication of (24) and (25) yields, just as before, an estimable equation of the following form:

$$Y^X_{cgvt} = \theta^X_{1cg} X^X_{1cgvt} + \theta^X_{2cg} X^X_{2cgvt} + u^X_{cgvt}\,. \tag{26}$$

The variables and coefficients are now defined as follows:

$$Y^X_{cgvt} \equiv \left[\Delta \ln \tilde{P}_{vgct} - \Delta \ln \tilde{P}_{vgc_rt}\right]^2\,, \tag{27}$$

$$X^X_{1cgvt} \equiv \left[\Delta \ln \tilde{Z}_{vgct} - \Delta \ln \tilde{Z}_{vgc_rt}\right]^2\,, \tag{28}$$

$$X_{2cgvt} \equiv \left[\Delta \ln \tilde{Z}_{vgct} - \Delta \ln \tilde{Z}_{vgc_rt}\right] \cdot \left[\Delta \ln \tilde{P}_{vgct} - \Delta \ln \tilde{P}_{vgc_rt}\right]\,, \tag{29}$$

$$u^X_{cgvt} \equiv \frac{1+\omega^X_{cg}\sigma^X_{cg}}{(1+\omega^X_{cg})(\sigma^X_{cg}-1)}\,\breve{\breve{\varepsilon}}_{vgct}\breve{\breve{\delta}}_{vgct}\,, \tag{30}$$

$$\theta^X_{1cg} \equiv \frac{\omega^X_{cg}}{(1+\omega^X_{cg})(\sigma^X_{cg}-1)}\,, \tag{31}$$

and

$$\theta^X_{2cg} \equiv \frac{\omega^X_{cg}\sigma^X_{cg} - 2\omega^X_{cg} - 1}{(1+\omega^X_{cg})(\sigma^X_{cg}-1)}\,. \tag{32}$$

The expressions for ω_{cg}^X and σ_{cg}^X are given by the same expressions as the ones for ω_{cg} and σ_{cg}, with θ_{1cg} and θ_{2cg} replaced by θ_{1cg}^X and θ_{2cg}^X, respectively.

5 Data and estimation

In this section, we briefly discuss the data we use for estimation, and the estimation methods applied.

5.1 Data

For all of our estimation purposes, we use the Eurostat COMEXT data. The database consists of trade data reported by all countries pertaining to the European Union, starting in year 1993 to and from up to 270 partners. The COMEXT database has the advantage of containing much richer and longer panels of European trade data than used in previous studies. It has the disadvantage that it reports detailed raw data, with frequently found outliers as well as commonly occurring missing values at the highest level of disaggregation. We face a trade off between high disaggregation and low data availability, and the opposite. The use of disaggregated data is desirable for two main reasons. First, we wish to avoid the aggregation biases present whenever macro level data is used for estimation, and homogeneity across sectors is hence implicitly imposed. As discussed in Section 2, it is an established fact that there exists non-negligible heterogeneity between different sectors of the economy, which, when neglected, tends to bias the estimates of the aggregate elasticity of substitution. This is what has motivated the use of micro level data even when the interest ultimately lies in obtaining macro level estimates. Second, we do not directly observe prices, but make use of data on values and quantities to compute unit values, which are then used as a proxy for true prices. The higher the level of aggregation, the more probable it is that the unit value is a biased measure, since we might be bunching together products with very different price levels and price movements.[8] On the other hand, the inference we make is based on asymptotic properties along the cross-section dimension, which implies that we need to have reasonably large samples of data for the exercise to make sense. Moreover, since

[8] We have used our data, to the extent it was possible, to assess the consequences of not having true price indices at the aggregation level we choose. Selecting a sector for which data availability was good enough to allow for the construction of time series at the most disaggregate level, we have constructed properly aggregated Fisher price indices and compared those to the unit values we use for estimation. The difference turned out to be very small, indicating that the use of unit values at the 4-digit level of disaggregation should not be a problem. Even though we cannot with certainty say whether this extends to all sectors, these indications do provide additional confidence in our estimation results.

our goal is to ultimately obtain properly aggregated macro level estimates of the elasticity of substitution, the exact choice of the disaggregation level of our estimation data is further governed by the availability of aggregation weights. For the aggregation exercise, in addition to trade data, we also need sector level data on internal production. This we obtain from the United Nations UNIDO database available only at the 4-digit ISIC level of disaggregation. So, despite having more disaggregated data that we could use for the estimation of sector level elasticities, we cannot reliably aggregate the estimates at a higher disaggregation level than the ISIC-4. Taking into account all of the above aspects, we choose not to use the highest level of disaggregation in our estimation exercise, but aggregate our data into 4-digit ISIC observations. This way, we substantially increase the size of the cross-section for each country-sector pair and make sure that we have proper weights for the aggregation of our elasticity estimates, while still keeping a high enough disaggregation level for the sectors to contain fairly homogeneous goods.

For the completion of our exercise, we need to do a mapping between the CN and ISIC nomenclatures, a quite tedious task since several revisions were done during the time span of our study and the lack of direct correspondence between the two nomenclatures. For the interested reader, we present some details on the mapping in Section B.2 in the appendix. Here, we only point out that we do this conversion before we perform the estimation, so that the only definition of sectors that we maintain throughout this paper is the ISIC-4.[9]

We split up the EU27 countries into two groups, depending on the time coverage of their data. For the countries for which we have longer time series, the "old" member states, we use yearly data starting in 1995. These countries are: Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, the Netherlands, Portugal, Spain, Sweden, and the UK. During the first part of our time span, Belgium and Luxembourg are reported as "BeLux", which we choose to maintain throughout in order to save on data; their results will therefore always overlap and will be reported as one. For the remaining countries, the "new" member states consisting of Bulgaria, Cyprus, the Czech Republic, Estonia, Hungary, Latvia, Lithuania, Malta, Poland, Romania, Slovakia, and Slovenia, we use yearly data starting in 2005. In all cases the sample ends after 2009, since the inclusion of 2010 would have implied the need to deal with yet another major CN revision, with the potential consequence of being forced to drop more sectors from our analysis.

[9]Note that for the aggregation of the export data estimates, we make use of yet another database: the CEPII BACI data. We do this because we need values of the complete trade taking place in each sector in each of the trading-partner countries, not only the EU ones. The BACI data is also grouped into 4-digit ISIC sectors.

The ISIC nomenclature Rev. 3 and 3.1 contains a total of 127 sectors at the 4-digit level. In the synchronization of the CN nomenclature, which we need to do because our data spans over several revisions, we lose 8 ISIC-4 codes that could not be traced throughout our time period of interest. Using the remaining 119 codes, we look for any sectors that have very low relative importance in all countries. With the final 106 sectors for which we perform estimations, we end up covering nearly 100% of production in most countries (the lowest coverage is obtained for Italy, where the sectors we include cover 98.14% of internal production), since any sector that has importance in any country is included for all.[10] Before aggregating our estimation data, we do a first round of outlier cleaning at the highest level of disaggregation – the 8-digit CN level, which encompasses some $14,000$ sectors. The details on the cleaning of the data are discussed in Section B.3 in the appendix. Having cleaned out the most extreme observations, we move on to aggregating the data into ISIC-4 sector, yearly time series.

Finally, when constructing the actual variables used for estimation, we need to choose reference varieties for each country-sector pair, to be used for the differencing of the price and expenditure shares data. As the reference variety v_r, used to eliminate terms common to all import varieties, we select a country with complete time coverage, i.e. a country that appears as a trading partner in all years of interest. If no such country exists, we choose the one with the greatest coverage. If, on the contrary, there are several candidates with complete time coverage, we select the one for which the mean traded value is the largest, as the quality of data may be better at high values, as discussed in Mohler (2009). Mohler further shows that the elasticity estimates are not very sensitive to the choice of reference variety, as long as the chosen variety is among the ones with the highest values. Our sensitivity checks confirm his findings: the elasticity estimates change only marginally with the change in reference variety, which is why we conclude that the impact of the reference variety choice on our results is rather small. As the exporter reference country, c_r, we select the country with the largest overlap of trading partners with our country of interest, using this measure as a proxy of similar trading patterns. Taking into account the possible heterogeneity in trade patterns across sectors within a country, we redo this selection process for each country-sector pair, to ensure that the reference varieties are always optimally chosen.

[10] In practice, when selecting what sectors to include, we make sure that we keep at least 85% of the domestic production in all countries, and that we do not lose any sectors that make up 1% or more of the domestic production in any country.

5.2 Estimation

For each of the 27 EU countries and for each sector separately, we estimate the elasticity of substitution for imports and exports, using (10) and (26) as estimating equations. If estimated by OLS, they will not yield consistent estimates as prices and expenditure shares are correlated with the demand and supply shocks, and hence also with the error terms in the estimation equations. To deal with this problem, Feenstra (1994) suggests including country-specific fixed effects as instruments, a method we also pursue. Our estimator yields consistent estimates of θ_{1cg} and θ_{2cg}, but not all such estimates are *theory-consistent*, due to the restrictions imposed by the structural model on σ_{cg}.

Defining

$$\rho_{cg} \equiv \frac{\omega_{cg}(\sigma_{cg}-1)}{1+\omega_{cg}\sigma_{cg}}\,, \tag{33}$$

the parameter restrictions required to obtain theory-consistent estimates can be summarized as shown below. For a derivation of the theory-implied restrictions, see Section A.1 in the appendix.

Given that $\hat{\theta}_{1cg} > 0$:

if $\hat{\theta}_{2cg} > 0$, then

$$\hat{\rho}_{cg} = \frac{1}{2} + \left[\frac{1}{4} - \frac{1}{4+\left(\hat{\theta}^2_{2cg}/\hat{\theta}_{1cg}\right)}\right]^{\frac{1}{2}}; \tag{34}$$

if $\hat{\theta}_{2cg} < 0$, then

$$\hat{\rho}_{cg} = \frac{1}{2} - \left[\frac{1}{4} - \frac{1}{4+\left(\hat{\theta}^2_{2cg}/\hat{\theta}_{1cg}\right)}\right]^{\frac{1}{2}}. \tag{35}$$

Furthermore, it is possible that a negative value of $\hat{\theta}_{1cg}$ yields a theory-consistent estimate if[11]

$$\hat{\theta}_{1cg} > -\frac{\hat{\theta}^2_{2cg}}{4}\,, \tag{36}$$

and in all cases

$$\hat{\sigma}_{cg} = 1 + \frac{2\hat{\rho}_{cg}-1}{\hat{\theta}_{2cg}(1-\hat{\rho}_{cg})}\,. \tag{37}$$

Finally, if $\hat{\theta}_{2cg} \to 0$, then $\hat{\rho}_{cg} \to \frac{1}{2}$ and $\hat{\sigma}_{cg} \to 1 + \hat{\theta}_{1cg}^{-\frac{1}{2}}$.

[11] See Feenstra (1994) for derivations. This case does occur from time to time in our estimations. In Feenstra (1994), it occurs for one out of the eight sectors for which theory-consistent estimates are obtained.

We report estimation results obtained from the 2SLS estimation, as in the original Feenstra (1994) paper, as well as results obtained using a bootstrapping procedure over the residuals from the original estimation. Note that we do not use the bootstrap as an alternative to the IV estimation, as the identification of the reduced-form parameters is done in the exact same way as in Feenstra (1994). Instead, it works as an addition to the original estimation procedure, alleviating the bias caused by the non-linearity of the function mapping the reduced-form parameters into structural ones. Specifically, our bootstrapping procedure works as follows. For each country-sector pair, we start by estimating the reduced-form parameters $\hat{\theta}_{1cg}$ and $\hat{\theta}_{2cg}$ (or $\hat{\theta}^X_{1cg}$ and $\hat{\theta}^X_{2cg}$, in the case of exports data estimation), using the original 2SLS estimator of Feenstra (1994). We bootstrap the residuals from the original estimation equation, given by (10) for import data and (26) for export data, applying a wild bootstrapping procedure.[12] We then use the new vector of residuals to reconstruct the dependent variable. Our reconstruction looks as follows:

$$\breve{Y}_{cgvt} = \hat{\theta}_{1cg} X_{1cgvt} + \hat{\theta}_{2cg} X_{2cgvt} + \hat{u}_{cgvt} \breve{\varepsilon}_t \,, \tag{38}$$

where $\breve{Y}_{cgvt}$ denotes the reconstructed dependent variable, $\hat{u}_{cgvt}$ is the residual from the original regression, and $\breve{\varepsilon}_t$ is white noise following a distribution such that $E(\breve{\varepsilon}_t) = 0$ and $E(\breve{\varepsilon}_t^2) = 1$. Finally, we estimate equation (38), saving the estimates if $\hat{\theta}_{1cg}$ and $\hat{\theta}_{1cg}$ fulfill the above outlined theoretical restrictions. The bootstrapping procedure is then repeated until we obtain $5,000$ pairs of theory-consistent estimates, keeping track of the total number of required draws as an indication of the plausibility of the imposed modeling assumptions. If the total number of draws is very large, we take it as an indication of the model providing a poor description of data for the country-sector pair in question. Note, however, that for high elasticity estimates, located close to the theory-inconsistent region of the parameter space, it is likely that we sometimes end up with theory-inconsistent results. For high estimates, we would thus expect the total number of draws to sometimes exceed $5,000$, even if the modeling assumptions are suitable.

The bootstrapping procedure provides us with a complete distribution of the estimated parameters and the mapped elasticity of substitution for each country-sector pair, allowing us to detect and characterize potential problems with the estimation. Given that we are often dealing with skewed distributions, we focus on robust statistics, i.e. the mode, median and the interquartile range, rather than the means and variances of the distributions. The bootstrapping procedure, and the output it generates, are explained

[12] We apply the wild bootstrap since it has been shown to work better than, for example, the paired bootstrap in the presence of heteroskedasticity, and equally well in its absence. We follow the wild bootstrapping procedure laid out in Mammen (1993).

in further detail in Corbo and Osbat (2012). Applying Monte Carlo techniques, we also show that the mode and the median of the bootstrap distributions are the least biased measures of the elasticity of substitution, when compared to the bootstrap mean and the 2SLS estimate. The advantage of the robust bootstrap measures is especially pronounced for mid-range estimates, where we find ourselves most of the time and where the mode slightly outperforms the median. For low estimates, all of the considered measures perform satisfactorily. For very high estimates, instead, none of them do – estimates in this region do not frequently occur.

The results from the original Feenstra (1994) estimation method are always included for comparison purposes. As we will see in Section 7, the point estimates obtained using the two methods are very close whenever we find a low elasticity, which we should also expect given that our bootstrapping procedure yields unbiased estimates. However, for some sectors where the elasticity estimates are high, notable differences emerge. As we discuss in more detail in Corbo and Osbat (2012), the problems emerge from the shape of the function mapping the estimated coefficients θ_{1cg} and θ_{2cg} (θ^X_{1cg} and θ^X_{2cg}), in case of import (export) data estimation, into the parameter of interest σ_{cg} (σ^X_{cg}). It is worth emphasizing again that the aforementioned bias does not stem from the estimation of the reduced-form parameters. It emerges in the actual mapping, and is present even under the assumption that the reduced-form parameters are consistently estimated. Given the bias-reducing properties of the bootstrap median and mode, it is natural that our preferred results are the bootstrapped ones. There are however additional advantages to the procedure. First, the bootstrapping offers higher transparency into our estimation method; given that it generates entire distributions of the parameters, we can detect and explain possible problems with the estimates. Moreover, it offers a natural way of measuring the accuracy of our estimates, as robust measures of dispersion can easily be obtained from the distributions. As the mapping function is discontinuous, previously used approximation methods for the computation of measures of dispersion are unsuitable in some regions of the relevant parameter space. With the bootstrap in place, we avoid having to rely on approximation methods. Second, it allows us to keep a larger share of the sectors without having to do a grid search based on rather arbitrary restrictions, as has been common in the literature following Broda and Weinstein (2006). As we have found that it is the sectors that tend to generate high estimates of the elasticity of substitution that are the least robust, increasing the robustness of these estimates turns out to be of importance for the magnitude and sensitivity of the aggregate elasticity estimate.

6 Aggregation of sector level elasticities

The ultimate goal of our exercise is to obtain reliable aggregate estimates of the elasticity of substitution between imports and domestically produced goods on one hand, and the elasticity of substitution of a country's exports on the world market on the other. To do so, we need to aggregate the sector level estimates obtained using the above discussed estimation procedures. Before moving on to the technical details on our methods of aggregation, we first briefly discuss potential problems emerging from the conceptual differences in the definitions of the elasticities that we measure and the elasticities that we are really interested in measuring.

The sector level elasticity estimates are based on shifts in relative prices and relative market shares for each sector. For aggregate elasticities, however, we will need to assume the occurrence of some aggregate shock, such as a shock to the exchange rate. If the simple supply function we assume in equations (2) and (18) is a good approximation locally, but poorer globally, we may not be capturing exactly the magnitude of the aggregate elasticity of substitution. While it seems plausible to assume that we can easily substitute from one supplier of our imports to another (or several others) in response to a relative price increase, as assumed in our estimation, to substitute away from all of our importers to our domestic production may prove more difficult. Given a substantial increase in the relative price of all imports to the price of the domestically produced variety, it may not always be possible to substitute to the same extent as in the case of a price increase in one imported variety, as total consumption may exceed the total domestic production capacity in some sectors. Whenever this is the case, the elasticity of substitution following an overall price increase in imports will be lower than the elasticity of substitution following a relative price increase in some imported variety obtained in our estimation exercise. Our derived aggregate estimates can then be thought of as representing the upper bound of the actual elasticities, that coincides with the true value if substitution is possible to the same extent following an overall price shock as following a relative price shock to import prices. Note, however, that this reasoning applies only to an increase in the relative price of imports to the domestically produced variety. In the case of a decrease in the price of imports, such production capacity restrictions are not likely to bind unless the importing country consumes a large fraction of the world's total production of some goods.

Sofar, the discussion has concerned the aggregation of sector level elasticities of substitution for imported goods. For the exported goods, the aggregate elasticity of substitution that we are interested in measuring quantifies the extent to which our exported goods can be substituted on the world market. The shock we have in mind, when performing the aggregation, is one that changes the price of exports of some good originating from the country under study, c relative to all other varieties of that good on the world market. This could, for example, be a shock that shifts the production costs in country c, or a change in its exchange rate that hits equally for all currencies. Just as in the case of imports, we may be overestimating the aggregate elasticity in case production capacity is not enough to satisfy demand. In theory, this may happen in response to a large decrease in the price of exported varieties from country c relative to all other varieties on the world market. Such a shock would imply substitution away from all other varieties and towards the ones originating from c, creating a risk for supply not being able to match demand. Given a shock that instead increases the price of exported varieties from country c, this reasoning does not apply.[13]

Our sector level estimates, as well as our aggregation exercise, rely on the assumption that the simple model we assume here provides a reasonably accurate description of the reality. However, other trade theories have been brought to light over the last decade, where much of the focus lies on firm heterogeneity and the importance of the extensive margin of trade. Chaney (2008) presents a model where firms are assumed to be heterogeneous in productivity, and there are fixed as well as variable trade costs. This renders a demand function similar to that we assume, but with an elasticity with respect to cost changes that depends on the heterogeneity of firms, and is larger than the elasticity for each individual firm.[14] Furthermore, Chaney finds that the elasticity of substitution makes the extensive margin less sensitive to changes in costs. Under certain conditions, and given our identification scheme, the fixed costs, such as the trade barriers assumed in Chaney (2008), will not affect our estimates. The conditions needed are that the fixed costs are either constant over time, in which case they cancel out when we difference over time, importer-specific, in which case they cancel out when we difference with respect to the reference variety, or variety- and time-specific, in which case they end up in the error term. If we believe that the heterogeneous firm model provides a better description of the world, and to the extent that the listed conditions on the fixed costs are fulfilled, then we are in fact estimating the degree of heterogeneity of firms and thus overestimating the elasticity of substitution. Hence, we need again worry about our elasticity estimates being too high, rather than too low. Moreover, the time-differencing

[13]In the extreme case of world consumption exceeding the production capacity of the world minus that of country c, we may wrongly estimate the aggregated elasticity even following an increase in the price of exported varieties from country c. However, for the goods and the disaggregation level of interest for our estimation, this is not likely to occur.

[14]This is the measure denoted by γ in Chaney (2008).

involved in our method, implies that we are not capturing all of the extensive margin of trade, as an observation is included only if there is registered trade between the two countries in the preceding period. As we are not using firm-level data, but letting a variety encompass all of the brands originating from a certain country, the effect of firm entry will be captured as long as there is already an established trade relationship between the importer and the exporter of interest. Any new trade relationships between countries, however, will not be. Based on the model in Chaney (2008), by not fully capturing the extensive margin, if anything, we are again overestimating the reaction of trade to cost changes – the more so the higher the true elasticity is. This supports our earlier conclusion that our elasticity estimates represent the upper bound of the true elasticity we wish to measure.

In summary, the results of our aggregation exercise are likely to be more accurate for smaller than for larger shocks. Furthermore, the possible bias stemming from the elasticities we are estimating being conceptually different from the ones that we are interested in measuring is asymmetric and will only be present in the cases where production capacity restrictions bind. This may in theory occur following overall increases in the prices of imported goods relative to domestically produced ones, and decreases in the price of exported varieties from the country under study relative to all other varieties on the world market. In both of these cases, the aggregate elasticities based on our sector level elasticity estimates will be upward biased. Since we cannot accurately measure the production capacities, however, we cannot with certainty tell the size of this bias nor the frequency with which it occurs. Even if firm heterogeneity not captured by our model plays an important role, we expect our measures to be too high rather than too low.

We finally note that the experiment we have in mind for our aggregation exercise, i.e. a one percent increase in the price of all imported (or exported) goods, may not be realistic for sectors for which the elasticity is very high. These sectors are characterized by very strong competition, in which case such cost increases are unlikely to be passed through to consumer prices. In the extreme case of perfect competition, any increase in the price of a variety would imply that it would stop being traded, as in this setting all firms are price takers. Our model assumes that there is at least some degree of monopolistic power in price setting, and so this extreme case does not apply. However, for sectors with very high elasticity estimates, we should still worry more about overestimating rather than underestimating the elasticities. As discussed in Section 5.2, all of our considered measures suffer from biases in the case of a high elasticity of substitution. As we demonstrate in Corbo and Osbat (2012) though, the 2SLS estimate displays an upward bias, while the bootstrap estimates are instead biased downward. This further favors the bootstrap measures as the more reliable estimates to be used for aggregation.

6.1 Aggregation of import data estimates

For the aggregation of import data estimates, we follow the methodology laid out in Imbs and Méjean (2009). They define the aggregate elasticity of substitution between bundles of domestic and foreign goods as

$$\sigma_c \quad \equiv \quad 1 + \frac{\partial \ln \sum_g \sum_{v \neq c} P_{cgv} C_{cgv} - \partial \ln \sum_g P_{cgc} C_{cgc}}{\partial \ln \psi_c}, \tag{39}$$

where we have denoted by ψ_c is a shock to the international relative price of country c's domestic goods, uniform across all trading partners v. Note that the variables are now not indexed by t, since the elasticity of substitution is assumed constant over time and we will eventually use time-averaged data for calibration of the aggregation weights. The elasticity of substitution obtained from value data is simply one plus the elasticity that would have been obtained on volumes data. The assumed properties of the shock allow us to interpret the estimate as capturing the substitutability between bundles of domestic goods and goods from *the rest of the world*. In other words, this collapses the multilateral model into a two-country setting, similar to the standard setting assumed in structural macroeconomic models. We next rewrite equation (39), so as to obtain an expression in terms of the sector level elasticities and weights that we can measure. In the interest of brevity, we have placed the derivations in Section A.2 in the appendix.

Inserting demand equation (1), and assuming that consumption of different goods is aggregated according to a Cobb-Douglas function, we can rewrite equation (39) as follows:

$$\begin{aligned} \sigma_c - 1 \quad = \quad & \sum_g \sum_{v \neq c} m_{cgv}(1-\sigma_{cg}) \frac{\partial \ln P_{cgv}}{\partial \ln \psi_c} - \sum_g m_{cgc}(1-\sigma_{cg}) \frac{\partial \ln P_{cgc}}{\partial \ln \psi_c} \\ & - \sum_g (m_{cg.} - m_{cgc})(1-\sigma_{cg}) \frac{\partial \ln P_{cg.}}{\partial \ln \psi_c} + \sum_g (m_{cg.} - m_{cgc}) \frac{\partial \ln P_{c..}}{\partial \ln \psi_c}, \end{aligned} \tag{40}$$

where the weights

$$
\begin{aligned}
m_{cgv} &\equiv \frac{P_{cgv}C_{cgv}}{\sum_g \sum_{v\neq c} P_{cgv}C_{cgv}} & (41)\\
m_{cgc} &\equiv \frac{P_{cgc}C_{cgc}}{\sum_g P_{cgc}C_{cgc}} & (42)\\
m_{cg.} &\equiv \sum_{v\neq c} m_{cgv} & (43)
\end{aligned}
$$

denote the expenditure share of good-g imports from country v in country c's total imports; the share of domestic expenditures on good g in total domestic expenditures in country c; and the expenditure share of good-g imports in country c's total imports, respectively. As we are estimating the response of expenditures to price changes in the long run, it is assumed that the domestic price will react fully, while the foreign prices will not to react at all to domestic cost changes in country c. The final aggregation equation for the import data elasticity estimates can then be written as

$$
\begin{aligned}
\sigma_c = & \sum_g m_{cgc}\sigma_{cg} + \sum_g (m_{cg.} - m_{cgc})(\sigma_{cg} - 1)(1 - w_{cg.}^M) \\
& + \sum_g (m_{cg.} - m_{cgc})\alpha_{cg}(1 - w_{cg}^M)\,, \qquad (44)
\end{aligned}
$$

where we have defined the share of imports in total expenditures on good g in country c as

$$
w_{cg}^M \equiv \frac{\sum_{v\neq c} P_{cgv}C_{cgv}}{\sum_{v\in V} P_{cgv}C_{cgv}} = 1 - \frac{P_{cgc}C_{cgc}}{P_{cg.}C_{cg.}}\,, \qquad (45)
$$

and where

$$
\alpha_{cg} \equiv \frac{P_{cg.}C_{cg.}}{P_{c..}C_{c..}} \qquad (46)
$$

is the expenditure share of good g in country c's total consumption.

6.2 Aggregation of export data estimates

In this section, we develop a way of aggregating the sector level elasticity estimates obtained from export data. We think of the aggregate elasticity of substitution of a country's exports as the substitutability of its basket of exported goods on the world market. In other words, we want to measure the decrease in export volumes following a one percent increase in the price of all exported goods. In a two country macroeconomic model, this measure corresponds to the foreign country's elasticity of substitution.

We define the aggregate elasticity of substitution for country c's exports as

$$\sigma_c^X \equiv 1 - \frac{\partial \ln \sum_g \sum_{v \neq c} P_{cgv} C_{cgv} - \partial \ln \sum_g \sum_{v \neq c} \sum_{i \neq c} P_{vgi} C_{vgi}}{\partial \ln \psi_c}, \tag{47}$$

where ψ_c is again a shock to the price of all country c's domestically produced goods, including its exports, and hence an import price shock for all trading partners v. Our interest lies in the extent to which the variety c of some good g will be substituted away on the world markets in response to a price change, and hence we sum up the response in the relative consumption of variety c to the consumption of all domestically produced and imported goods in all countries that are importing good g from country c. Henceforth, when referring to the world minus our country of interest c, we shall use the term *rest of the world*. We proceed with the derivations of the aggregation equation for exports in the same way as we did in the case of imports, with the goal of obtaining an expression in terms of the sector level elasticities and measurable weights. The more detailed derivations can be found in Section A.3 in the appendix.

Under the same assumptions as in the case of imports, inserting country v's demand equations in (47) and rearranging, we have

$$\begin{aligned}
\sigma_c^X - 1 = & - \sum_g \sum_{v \neq c} x_{vgc} (1 - \sigma_{vg}) \frac{\partial \ln P_{vgc}}{\partial \ln \psi_c} \\
& + \sum_g \sum_{v \neq c} \sum_{i \neq c} x_{vgi}^W (1 - \sigma_{vg}) \frac{\partial \ln P_{vgi}}{\partial \ln \psi_c} \\
& + \sum_g \sum_{v \neq c} (x_{vgc} - x_{vg}^W)(1 - \sigma_{vg}) \frac{\partial \ln P_{vg.}}{\partial \ln \psi_c} \\
& - \sum_g \sum_{v \neq c} (x_{vgc} - x_{vg}^W) \frac{\partial \ln P_{v..}}{\partial \ln \psi_c},
\end{aligned} \tag{48}$$

where

$$x_{vgc} \equiv \frac{P_{vgc} C_{vgc}}{\sum_g \sum_{v \neq c} P_{vgc} C_{vgc}} \tag{49}$$

$$x_{vgi}^W \equiv \frac{P_{vgi} C_{vgi}}{\sum_g \sum_{v \neq c} \sum_{i \neq c} P_{vgi} C_{vgi}} \tag{50}$$

$$x_{vg}^W \equiv \sum_{i \neq c} x_{vgi}^W \tag{51}$$

are the share of good-g exports to country v out of c's total exports; the share of country i's good-g exports to country v in rest of the world's total output; and non-

c consumption of good g in country v as a share in rest of the world's total output, respectively. Note that the export aggregation equation (48) has the same components as the import aggregation equation (40). The first two terms capture the, adequately weighted, responses of the individual prices, the third term the response of the sector price index, and the last term the response of the aggregate consumption price index, in the importing countries. Since we have assumed that $\sigma_{vg} = \sigma^X_{cg}$, for all v, letting

$$x_{cg} \equiv \sum_{v \neq c} x_{vgc} \tag{52}$$

and

$$\alpha_{vg} \equiv \frac{P_{vg.}C_{vg.}}{P_{v..}C_{v..}}\,, \tag{53}$$

we can write the final equation for the aggregate elasticity of substitution for country c's as

$$\begin{aligned} \sigma^X_c \;=\; & \sum_g x_{cg}\sigma^X_{cg} + \sum_g \sum_{v \neq c} (x_{vgc} - x^W_{vg})(1 - \sigma^X_{cg}) w^X_{vg} \\ & - \sum_g \sum_{v \neq c} (x_{vgc} - x^W_{vg}) \alpha_{vg} w^X_{vg}\,, \end{aligned} \tag{54}$$

where

$$w^X_{vg} \equiv \frac{P_{vgc}C_{vgc}}{P_{vg.}C_{vg.}}\,, \tag{55}$$

is the share of country-c exports in country v's total consumption of good g.

Details on how we compute the weights used for the aggregation of export as well as import data estimates are shown in Section A.4 in the appendix.

7 Results

7.1 Estimation results

In total, including all of the 27 countries for which we perform our estimations, we obtain more than 2,500 elasticity estimates for each of the considered elasticity measures. Given the obvious difficulty in presenting all of these results, we choose to present the complete set of sector level results for one country only. We chose Germany as it is the largest of the EU economies. Detailed results on the remaining 26 countries, for which only aggregate estimates are reported in Section 7.3 below, are available from the authors upon request.

Table 1: Sector level results for the German economy

Sector	#Partn.	2SLS		Bootstrap					
		Est.	StDev	#Draws	Mean	25%	Median	75%	Mode
1511	68	7.3	18.6	5027	11.4	5.2	6.6	9.0	5.6
1512	103	10.7	31.4	5000	8.2	6.8	7.8	9.3	7.2
1513	118	4.3	4.8	5000	4.2	3.8	4.2	4.6	4.1
1514	91	2.8	1.5	5000	2.7	2.5	2.7	2.9	2.7
1520	48	2.2	1.7	5000	2.3	1.9	2.2	2.6	2.0
1531	77	3.2	3.2	5000	3.3	3.1	3.3	3.6	3.3
1532	46	2.8	1.6	5000	2.8	2.6	2.8	3.1	2.8
1533	49	3.3	2.3	5000	3.2	2.9	3.2	3.5	3.1
1541	69	2.4	1.7	5001	2.6	2.1	2.4	2.9	2.3
1542	46	3.0	1.1	5000	3.0	2.7	2.9	3.2	2.9
1543	84	4.0	4.3	5000	4.0	3.4	3.9	4.4	3.7
1549	96	13.0	53.4	5107	17.3	8.3	10.5	14.3	8.7
1551	81	3.0	2.9	5000	2.8	2.5	2.8	3.1	2.6
1552	73	3.9	3.6	5000	3.6	3.3	3.6	3.9	3.6
1553	61	6.7	11.9	5000	6.8	5.4	6.3	7.2	6.4
1554	75	5.1	1.1	5000	4.7	3.8	4.4	5.2	4.0
1600	52	12.5	50.9	5000	9.5	5.9	8.1	11.3	7.2
1711	107	4.6	5.7	5000	4.5	4.0	4.4	4.9	4.2
1721	109	7.8	18.3	5004	7.9	5.5	6.7	8.5	5.9
1722	86	3.3	3.0	5000	3.2	2.9	3.2	3.5	3.1
1729	75	4.2	4.5	5000	3.8	3.4	3.7	4.2	3.6
1730	103	4.9	5.7	5000	4.5	4.0	4.4	4.8	4.3
1810	147	3.4	2.6	5000	3.4	3.1	3.3	3.6	3.3
1911	75	15.1	72.5	5021	13.2	7.7	9.7	13.3	8.4
1912	97	19.6	202.5	5183	19.2	5.5	8.1	13.3	6.9
1920	91	5.9	7.2	5000	5.9	5.0	5.7	6.6	5.4
2010	93	3.3	2.6	5000	3.2	3.0	3.2	3.5	3.1
2021	68	1.6	0.0	5000	1.6	1.5	1.5	1.7	1.5
2022	80	3.6	4.2	5000	3.5	3.1	3.5	3.9	3.4
2023	91	2.4	2.0	5000	2.3	2.2	2.3	2.5	2.3
2029	122	3.1	2.3	5000	3.0	2.8	3.0	3.2	2.9
2101	72	2.8	1.8	5000	2.7	2.5	2.6	2.8	2.6
2102	85	3.8	4.4	5000	3.6	3.2	3.5	3.9	3.4
2109	82	2.6	1.8	5000	2.6	2.4	2.6	2.7	2.5
2211	98	4.1	4.9	5000	4.0	3.6	3.9	4.3	3.9
2212	71	3.8	3.4	5000	3.7	3.2	3.6	4.0	3.6
2221	93	4.2	5.8	5000	3.6	3.2	3.5	3.9	3.5
2222	67	3.9	5.0	5000	3.6	3.2	3.5	4.0	3.3
2310	19	4.4	0.7	5000	4.5	3.5	4.6	5.2	4.8
2320	40	2.4	0.1	5003	2.4	2.2	2.3	2.5	2.3
2330	8	-	-	-	-	-	-	-	-
2411	105	2.2	1.6	5000	2.3	2.0	2.2	2.4	2.1
2412	41	2.4	1.3	5000	2.4	2.2	2.3	2.5	2.3
2413	78	5.5	7.5	5000	5.6	4.6	5.5	6.5	5.5
2422	63	3.3	3.9	5000	3.2	3.0	3.2	3.5	3.2
2423	75	3.5	5.1	5000	3.4	2.7	3.1	3.8	2.8
2424	92	3.3	2.3	5000	3.2	3.0	3.1	3.4	3.1
2429	103	2.6	0.2	5000	2.6	2.3	2.5	2.8	2.5
2430	63	4.1	3.4	5000	4.1	3.5	3.9	4.5	3.8
2511	74	3.5	4.1	5000	3.4	2.9	3.3	3.8	3.2
2519	79	3.7	3.7	5000	3.5	3.2	3.4	3.8	3.3
2520	122	5.6	10.8	5000	5.2	4.3	4.9	5.7	4.6
2610	98	3.5	4.9	5000	3.4	2.8	3.2	3.7	3.0
2691	100	3.4	2.9	5000	3.1	2.9	3.1	3.3	3.1
2692	57	3.6	1.8	5000	3.5	3.3	3.5	3.8	3.5
2693	55	2.9	3.2	5000	3.0	2.5	2.9	3.4	2.7

Continued on next page

Sector	#Partn.	2SLS		Bootstrap					
		Est.	StDev	#Draws	Mean	25%	Median	75%	Mode
2694	38	3.2	1.3	5000	3.1	3.0	3.1	3.2	3.1
2695	54	2.2	1.8	5000	2.1	2.0	2.1	2.3	2.1
2696	80	2.0	0.9	5000	2.0	1.9	2.0	2.0	2.0
2699	75	2.3	1.7	5000	2.3	2.0	2.2	2.4	2.1
2710	87	3.3	3.3	5000	3.0	2.8	3.0	3.2	2.9
2720	102	3.5	3.1	5000	3.7	3.0	3.6	4.2	3.2
2811	80	2.8	2.2	5000	2.8	2.5	2.7	3.0	2.7
2812	62	4.5	7.6	5000	4.2	3.5	3.9	4.6	3.7
2813	38	3.6	3.1	5000	3.3	3.1	3.3	3.6	3.2
2893	109	3.4	4.4	5000	3.4	2.9	3.2	3.7	3.1
2899	122	3.2	2.4	5000	3.2	2.9	3.2	3.4	3.1
2911	84	30.5	432.5	5656	36.4	7.7	10.8	18.5	8.6
2912	113	2.5	2.2	5000	2.5	2.3	2.5	2.7	2.4
2913	99	10.4	46.4	5011	10.1	6.4	8.1	10.7	7.0
2915	73	3.5	0.3	5000	3.5	3.1	3.4	3.8	3.3
2919	106	7.8	23.8	5029	9.7	5.3	6.7	9.0	5.8
2921	64	3.1	3.2	5000	3.0	2.7	2.9	3.2	2.8
2922	76	4.0	5.9	5000	4.2	3.2	3.7	4.5	3.5
2923	54	3.6	4.2	5000	3.4	3.0	3.3	3.8	3.2
2924	88	3.2	2.9	5000	3.1	2.8	3.0	3.3	3.0
2925	75	5.0	16.7	5593	31.8	3.6	4.6	5.4	4.6
2926	81	4.0	7.3	5000	3.8	3.3	3.7	4.2	3.5
2927	48	2.6	2.3	5000	2.4	2.2	2.4	2.6	2.4
2929	104	5.2	17.7	5001	5.0	3.5	4.2	5.3	3.7
2930	72	8.9	32.1	5002	8.5	5.9	7.2	9.3	6.6
3000	131	3.7	7.2	5000	3.2	2.9	3.2	3.5	3.1
3110	117	1.8	0.1	5000	1.8	1.7	1.8	1.9	1.8
3120	100	6.3	1.0	5017	9.0	4.4	6.2	8.7	4.7
3130	71	3.1	2.2	5000	3.2	2.8	3.1	3.5	3.0
3140	39	2.2	3.1	5000	2.2	2.0	2.2	2.4	2.2
3150	83	8.3	34.8	5042	10.4	4.9	6.0	9.5	5.4
3190	77	2.4	2.9	5000	2.4	2.1	2.3	2.6	2.3
3210	71	20.1	15.5	6254	25.0	5.7	8.3	13.7	6.2
3220	99	2.9	2.5	5000	2.8	2.6	2.8	2.9	2.8
3230	84	2.3	2.1	5000	2.2	2.0	2.2	2.4	2.2
3311	93	4.4	0.7	5000	4.3	3.3	4.1	4.9	3.9
3312	115	11.4	85.3	5335	17.4	5.2	6.9	11.0	5.4
3313	62	9.1	41.6	5014	7.6	4.7	5.7	7.6	5.0
3320	91	2.2	2.2	5000	2.1	1.9	2.0	2.2	2.0
3410	102	44.8	823.2	6125	36.7	7.9	10.3	15.3	8.4
3420	66	4.7	8.1	5000	5.8	3.9	4.8	6.2	4.1
3430	123	2.4	2.7	5000	2.4	2.2	2.4	2.5	2.3
3511	44	3.4	2.8	5000	3.3	3.0	3.3	3.6	3.2
3520	52	3.7	0.6	5000	3.5	3.0	3.3	3.9	3.2
3530	86	16.6	80.1	5207	22.0	8.6	11.2	17.7	9.2
3610	127	5.7	11.0	5000	5.1	4.3	4.9	5.7	4.6
3691	136	2.0	1.6	5000	2.0	1.8	1.9	2.1	1.9
3693	69	5.3	8.7	5000	4.8	3.8	4.5	5.4	4.0
3694	88	3.3	6.1	5000	3.1	2.4	2.8	3.4	2.6
3699	116	5.4	9.2	5000	5.2	4.5	5.0	5.7	4.7

Table 1 contains the detailed results for the estimation performed on German import data. The first column contains the sector codes according to the ISIC-4 classification, while the second contains the number of trading partners that we observe in each sector. In other words, column 2 contains the size of the cross-section in our estimations. Note that, in case the number of trading partners is below 10, no results are reported as we refrain from doing any estimations when the sample size is too small. Columns 3 and 4 in the table contain the original 2SLS estimate, and the variance computed using the delta method, as in Feenstra (1994). The rest of the table is devoted to the bootstrap results, with column 5 displaying the number of draws needed to generate $5,000$ theory-consistent estimates, column 6 the mean of the bootstrap distribution, columns 7-9 the interquartile range including the median, and column 10 the mode of the distribution. As is obvious from the table, the mode is generally lower, and occasionally even much lower than the mean of the distribution, indicating the existence of skewed distributions. While some sectors display a close-to-normal distribution, a large fraction of the bootstrap samples suffer from the presence of outliers. These outliers, albeit few, can take on very extreme values. The mean, therefore, does not yield a good representation of the distribution. Much more robust, and for the sample more representative measures are the median and the mode.

As already mentioned, based on Monte Carlo simulations, the results of which are discussed in more detail in Corbo and Osbat (2012), we find that the mode is the measure that yields the lowest bias in the relevant parts of the parameter space when compared to the 2SLS estimate, the bootstrap mean and the bootstrap median.[15] In what follows, we will refer to the mode of the bootstrap distribution when discussing the results, unless otherwise mentioned.

As discussed in Section 5.2, for most of the sectors, the 2SLS estimates are very close, if not identical to the bootstrap mode and median estimates, especially so if the elasticities are low. For any sector for which the estimation is unproblematic and the elasticity falls within a well-behaved region of the mapping function, we would also expect this to be the case. However, for higher elasticity estimates, located near the explosive part of the mapping function, we are generally obtaining lower bootstrapped elasticity estimates. The shape of the mapping function, which displays a discontinuity close to the region where the elasticities are very high, is also what causes the sectors where the 2SLS displays a high variance to generally require a larger number of bootstrap draws. The mode of the distributions is again largely unaffected by this, and allows us to obtain elasticity estimates even in the cases where 2SLS produces no relevant results.

Finally, note that it is generally the case that, the higher the elasticity estimate, the higher its variance. This feature is expected given the form of the mapping function used to convert the estimates of θ_1 and θ_2 into elasticity estimates, as the function gets

[15] For very high elasticity values, the mode is also somewhat biased, but it still outperforms the 2SLS and the bootstrap mean. The median performs slightly better than the mode in this region, and nearly as well in the rest of the parameter space. As very high elasticity estimates are not frequently occurring, however, we select the mode as the measure for our baseline set of results.

considerably steeper when the elasticity values increase. As discussed in Section 5.2, applying the bootstrapping procedure, we obtain full distributions of both the estimated parameters θ_1 and θ_2, and the structural parameter of interest σ. We often observe non-normal σ-distributions, with a few and sometimes very extreme outliers. We therefore conclude that the variance of the estimates is a bad measure of dispersion; indeed, from column 4 in Table 1 we see that the variances are sometimes quite high. The interquartile range, however is not sensitive to the irregularities in the distributions, and thus offers a better picture of the accuracy of our estimates.

7.2 Sectoral heterogeneity

We found a rather high level of heterogeneity at the sector level. Comparing the dispersion of the estimates across sectors within a country to the dispersion within a sector across countries, we see that the cross-sector heterogeneity exceeds the cross-country one. This suggests that compositional effects may be the main driver behind the differences in the aggregate country level estimates. The mean (median) standard deviation of the sector level import data estimates within a country is 4.15 (3.52), while the same measure across countries within a sector is 3.27 (2.43). For the export data estimates, the corresponding figures are 2.57 (2.27) and 1.81 (1.49), respectively.

We next look at the ranking of different sectors, in terms of their estimated elasticity, across countries. This we do in order to assess whether there are some sectors that display consistently higher or lower estimates than others. Having ranked the sectors according to the elasticity estimates for each country separately, we perform a simple t-test for each pair of sectors to check whether their average rankings are significantly different from each other. For the import data estimates, we find that the 26 sectors ranked in the middle do not have significantly different ranks, while the higher and lower ranked sectors do. For the export data estimates, it is only the 12 sectors in the middle that are not significantly different. The sector with the on average highest rank across both import and export data estimates is sector *3530: Aircraft and spacecraft*. Looking at the frequency with which a certain sector appears among the ones that require more than $5,000$ bootstrap draws in order to obtain $5,000$ theory-consistent estimates, we also see that sector *3530* is the most frequently occurring one. The second most frequently occurring sector, among the ones that require a larger number of draws, is *3312: Measuring/testing/navigating appliances, etc.*. This sector has the fourth highest average ranking across countries and across both import and export data estimates. Hence, these sectors seem to be characterized by a high substitutability and thereby also problematic to obtain estimates for. On the other extreme, there are a fairly large number of sectors that never require more than $5,000$ bootstrap draws: 48 for the import data estimates, 65 for the export data estimates, and 35 when looking across both.

We then instead look at whether the sector elasticity estimates are in fact so similar, that their difference across countries is negligible. Figures 1 and 2 display box plots of all of the bootstrap distributions obtained from import data for two selected sectors, offering a fairly representative picture of the sectoral distributions. It is clear from the figures that, even though the sector elasticities are often close to each other, they are estimated precisely enough for their inter-quartile ranges not to overlap for most countries. Compared to Figure 1, Figure 2 shows a sector where the elasticity estimates are generally higher. It is clear from the figure that, as we have pointed out earlier, the estimates that are very high are also imprecise. Nevertheless, these estimates differ from a group of much lower ones, even though they might not be significantly different from each other. Based on the analysis of the entire set of sectors, we conclude that there are clear patterns in the sector level estimates, with some sectors displaying consistently higher and some consistently lower elasticities across countries. Still, the elasticity estimates for most sectors are not equal across countries, which indicates that there is a case for using country-specific sector level estimates in the aggregation, despite the fact that the cross-sector heterogeneity exceeds the cross-country one.

Finally, comparing our median sector level estimates to the ones presented in Broda et al. (2006), we find that, in most cases, our 2SLS estimates are of the same order of magnitude. There are exceptions, however; for Cyprus, Greece and the UK, Broda et al. (2006) obtain notably lower estimates than we do, while, for Romania, their estimate is much higher than ours. This, perhaps, is not surprising, due to the differences in the data sets employed and time periods of study. When we compare our median sector level estimates to Mohler and Seitz (2010), which are based on the same database as ours although covering a somewhat different time period, we find that our estimates in most cases are in line with theirs. As Mohler and Seitz run their estimations at a higher order of disaggregation than we do, we may expect their estimates to be somewhat higher, as it is usually found that more disaggregated goods are more substitutable than less disaggregated ones. They are indeed higher for nine countries, but they are also lower for some. However, the median elasticity is a rather blunt measure for comparison. The same argument could also be applied to the Broda et al. (2006) estimates, although this specific discrepancy between the results may be obscured by the differences in the data sets employed in their estimations and in ours. We note here also that we can only do these comparisons for the import data estimates. Corresponding estimates based on

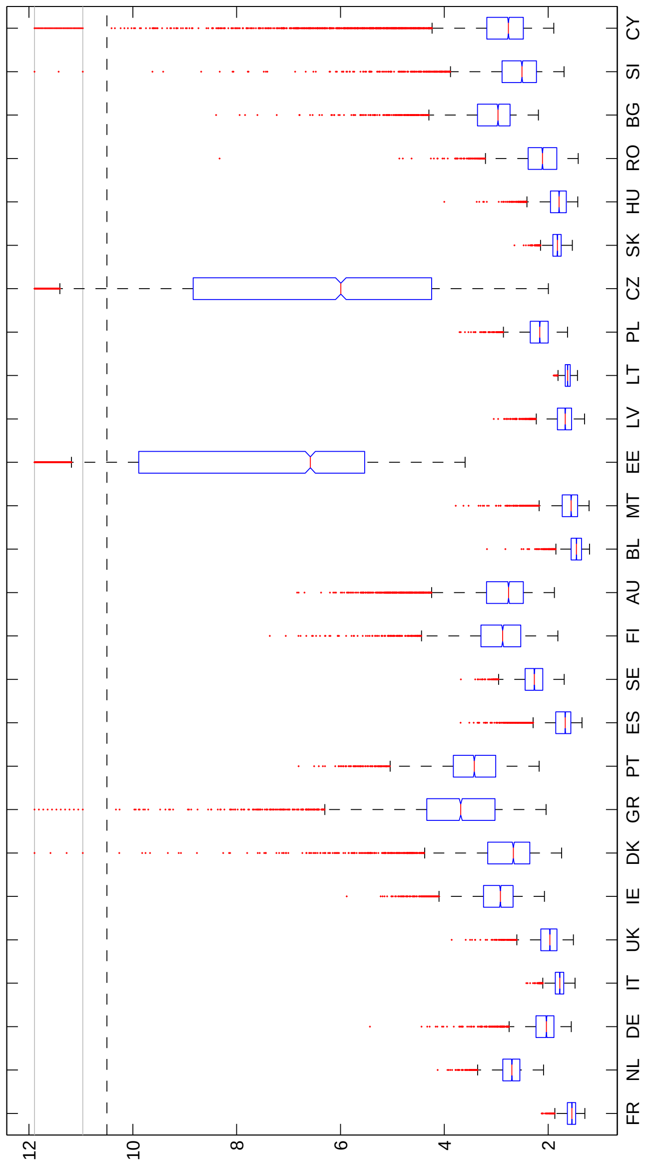

Figure 1: Box plots of import data estimates for sector *3320: Optical instruments & photographic equipment*

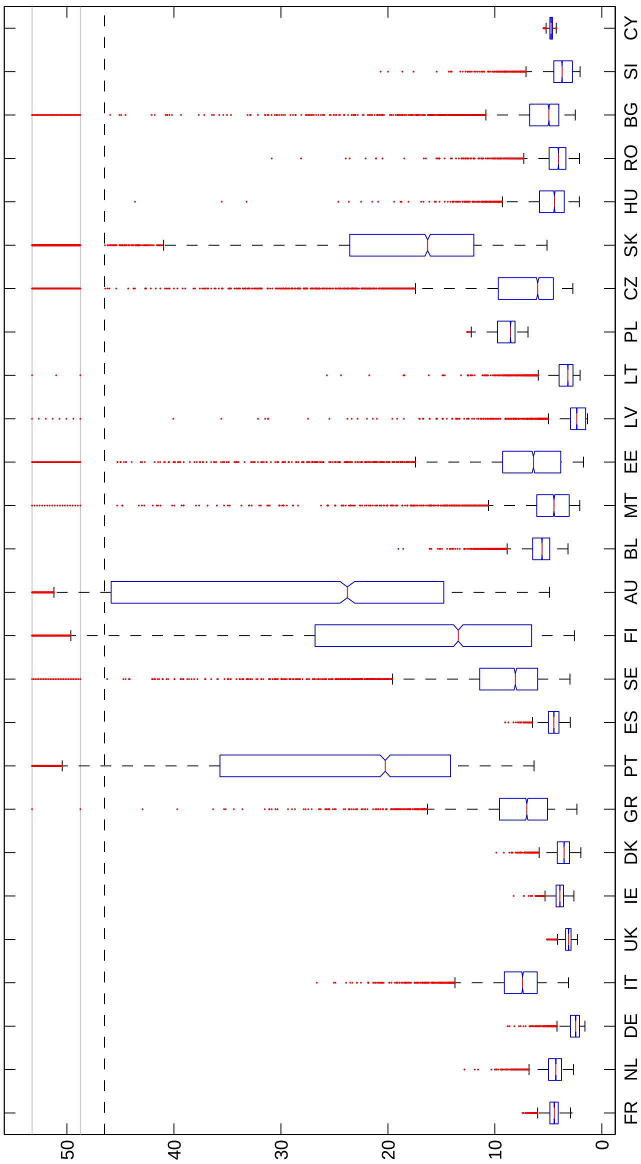

Figure 2: Box plots of import data estimates for sector *1541: Bakery products*

export data have, to the best of our knowledge, not been presented earlier.

7.3 Aggregation results

Applying the aggregation procedure described in Section 6 on our estimates of the sector level elasticities, we obtain the aggregate elasticities of substitution listed in Tables 2 and 3. We display the aggregation results obtained using the 2SLS estimates as well as the results obtained using the bootstrapped elasticities. To facilitate comparison between the two methods, in addition to the results based on the complete set of bootstrapped estimates, we also show aggregates including only the sectors for which we have been able to obtain theory-consistent estimates using 2SLS.

As is evident from the tables, the bootstrap median and mode are generally lower than the 2SLS estimates. The aggregates based on the means of the bootstrap distributions are displayed only to give some indication of the skewness of the underlying distributions; as the means are not a representative measure of these distributions, we will not comment further on the mean-based results. For Cypriot imports, and for Estonian exports, the 2SLS aggregate is very high. For Cyprus, it is almost entirely due to the extreme elasticity estimate in sector *2692: Refractory ceramic products*; the 2SLS estimate of the elasticity is almost $50,000$. For the Estonian exports, the high aggregate estimate is explained by the sector *3610: Furniture*, for which the 2SLS estimate equals 133. It should be noted, however, that the standard deviations of these high elasticity estimates are extremely high as well, rendering the elasticities for these sectors insignificant. This clearly shows the main problem one encounters when trying to obtain aggregate estimates: very high elasticity estimates in single sectors will pull up the aggregate considerably, severely altering the implications for macroeconomic models. It is therefore crucial that these estimates are reliable if they are to be included in the aggregate. At the same time, in the present setting, simply excluding the insignificant estimates does not provide a good solution either, as we are dealing with highly heteroskedastic measures. It will almost certainly be the case that the high estimates are the insignificant ones, due to the shape of the mapping function. However, as already discussed, the high estimates seem, to some extent, to be specific to the estimation method as they are lowered when we bootstrap the estimates. Moving on to the aggregation based on all sectors for which data was available, we see that the import elasticities are all unchanged for Germany and Ireland, indicating that no new sectors have been added. For most of the remaining countries, the estimates when all sectors are included are only slightly higher. That the aggregate estimates increase when we include all sectors is

Table 2: Aggregated elasticities of substitution of imports

Country	2SLS	Bootstrap (comparison)			Bootstrap (All sectors)		
		Mean	Median	Mode	Mean	Median	Mode
Austria	5.1	5.4	3.6	3.3	24.5	5.7	4.5
BeLux	3.7	5.9	3.2	3.0	7.2	3.5	3.1
Bulgaria	13.6	12.8	4.8	3.8	13.7	4.9	3.8
Cyprus	89.2	17.4	4.3	3.8	25.1	5.6	4.4
Czech Republic	4.3	5.4	3.4	3.2	13.3	3.8	3.4
Denmark	4.0	4.2	3.1	2.9	8.6	3.9	3.3
Estonia	9.0	9.8	3.9	3.5	29.3	6.5	4.8
Finland	3.6	4.2	3.2	3.0	10.9	4.0	3.5
France	3.8	5.3	3.3	3.1	13.2	4.3	3.7
Germany	8.7	8.2	4.2	3.7	8.2	4.2	3.7
Greece	3.4	28.3	2.9	2.8	28.5	3.1	2.9
Hungary	6.3	6.5	3.2	2.9	9.6	3.8	3.3
Ireland	11.9	5.0	2.9	2.7	5.0	2.9	2.7
Italy	3.5	3.5	3.2	3.1	4.6	3.4	3.2
Latvia	5.6	7.1	3.0	2.7	8.1	3.1	2.8
Lithuania	11.5	4.8	2.7	2.5	8.0	3.1	2.7
Malta	13.2	7.7	2.9	2.7	8.7	2.9	2.6
Netherlands	5.9	9.2	4.0	3.4	9.5	4.1	3.5
Poland	5.0	9.2	3.8	3.4	13.4	4.5	3.7
Portugal	6.0	4.5	3.1	2.9	6.8	3.6	3.3
Romania	3.5	4.7	2.9	2.7	6.0	3.1	2.8
Slovakia	5.8	5.9	3.9	3.6	8.5	4.1	3.7
Slovenia	5.5	7.4	3.6	3.3	13.1	4.8	4.0
Spain	3.8	3.8	3.4	3.2	5.5	3.8	3.4
Sweden	8.3	14.1	4.7	4.1	15.8	5.0	4.2
UK	3.3	3.2	2.9	2.7	4.4	3.1	2.9

expected, since, as discussed earlier, the function mapping the estimated coefficients into elasticities tends to be problematic for high, rather than low, elasticity values. For most countries, however, the increase is marginal, as the set of sectors included in the two aggregations is largely overlapping. The exceptions are Austria, Cyprus, Estonia, and Slovenia, where the increase is more substantial. For Austria, almost all of the increase is explained by one additional sector, *3410: Motor vehicles*, for which the 2SLS estimate was theory-inconsistent and hence not included in the aggregation. Similarly, for Estonia and Slovenia, the responsible sector is *3610: Furniture* and *2930: Domes-*

Table 3: Aggregated elasticities of substitution of exports

Country	**2SLS**	**Bootstrap (comparison)**			**Bootstrap (All sectors)**		
		Mean	Median	Mode	Mean	Median	Mode
Austria	6.8	5.3	4.0	3.8	5.3	4.0	3.8
BeLux	3.6	3.4	3.1	3.0	3.4	3.1	3.0
Bulgaria	5.1	5.6	3.3	3.1	7.4	3.4	3.0
Cyprus	17.2	8.0	5.2	4.5	22.0	6.4	5.2
Czech Republic	5.4	6.4	4.0	3.6	9.5	4.3	3.8
Denmark	4.1	4.2	3.6	3.4	4.2	3.6	3.4
Estonia	19.8	13.7	7.3	6.2	19.2	7.9	6.5
Finland	4.3	3.7	3.4	3.3	4.6	3.5	3.4
France	12.9	8.7	4.4	3.8	8.9	4.4	3.8
Germany	4.5	3.8	3.4	3.3	13.3	5.3	4.3
Greece	5.9	17.9	4.1	3.7	22.0	4.6	4.1
Hungary	6.0	7.6	4.5	4.2	7.7	4.5	4.2
Ireland	3.7	3.4	3.3	3.2	3.4	3.3	3.2
Italy	3.9	3.8	3.2	3.1	4.7	3.4	3.2
Latvia	6.2	6.7	4.6	4.3	19.3	6.0	5.1
Lithuania	7.5	11.0	5.4	4.7	14.0	6.0	4.9
Malta	4.7	8.9	4.0	3.9	34.4	5.0	4.4
Netherlands	5.7	3.9	3.5	3.4	5.3	3.7	3.5
Poland	7.2	8.2	5.3	4.7	9.0	5.3	4.7
Portugal	5.7	5.0	4.0	3.8	5.4	4.1	3.9
Romania	4.1	5.8	3.4	3.2	6.4	3.4	3.2
Slovakia	4.9	6.2	4.0	3.8	7.8	4.3	3.9
Slovenia	5.2	7.9	4.2	3.8	10.8	4.4	3.9
Spain	3.5	4.2	3.0	2.9	6.7	3.5	3.2
Sweden	4.6	5.6	4.0	3.9	25.5	5.2	4.5
UK	6.7	6.0	3.3	3.0	6.0	3.3	3.0

tic appliances n.e.c., respectively, while several sectors are contributing in the case of Cyprus. On the exports side, we observe unchanged elasticities for Austria, Belgium and Luxembourg, Denmark, Ireland, and the UK. Again, the increases for the rest of the countries are marginal, except for a few cases. For Germany and Latvia, the changes are more pronounced, which can be attributed to single sectors in both cases; the sectors are *3530: Aircraft and spacecraft* and *1512: Processing/preserving of fish*, for Germany and Latvia, respectively.

Table 4: Aggregated bootstrapped elasticities of substitution of imports, using country-specific weights together with country-specific or all German elasticities

Country	**Mean**		**Median**		**Mode**	
	Country-spec.	All DE	Country-spec.	All DE	Country-spec.	All DE
Austria	24.5	5.8	5.7	3.5	4.5	3.2
BeLux	7.2	3.4	3.5	2.8	3.1	2.6
Bulgaria	13.7	3.7	4.9	3.0	3.8	2.8
Cyprus	25.1	3.5	5.6	3.0	4.4	2.8
Czech Republic	13.3	7.1	3.8	3.8	3.4	3.5
Denmark	8.6	3.8	3.9	3.0	3.3	2.8
Estonia	29.3	3.8	6.5	3.0	4.8	2.8
Finland	10.9	4.0	4.0	3.1	3.5	2.9
France	13.2	8.5	4.3	4.4	3.7	3.9
Germany	8.2	8.2	4.2	4.2	3.7	3.7
Greece	28.5	3.9	3.1	3.0	2.9	2.8
Hungary	9.6	6.1	3.8	3.4	3.3	3.0
Ireland	5.0	3.8	2.9	2.9	2.7	2.7
Italy	4.6	7.1	3.4	4.1	3.2	3.7
Latvia	8.1	3.5	3.1	2.9	2.8	2.7
Lithuania	8.0	3.1	3.1	2.6	2.7	2.4
Malta	8.7	3.2	2.9	2.5	2.6	2.4
Netherlands	9.5	3.8	4.1	2.9	3.5	2.7
Poland	13.4	5.1	4.5	3.3	3.7	3.1
Portugal	6.8	4.1	3.6	3.1	3.3	2.9
Romania	6.0	5.4	3.1	3.3	2.8	3.0
Slovakia	8.5	4.5	4.1	3.4	3.7	3.2
Slovenia	13.1	4.0	4.8	3.1	4.0	2.9
Spain	5.5	7.2	3.8	3.9	3.4	3.6
Sweden	15.8	7.7	5.0	4.1	4.2	3.7
UK	4.4	7.3	3.1	3.9	2.9	3.5

The lower aggregate elasticities obtained using the bootstrapped estimates in Tables 2 and 3 are much closer to the values we observe in the macroeconomic literature than are the estimates obtained using 2SLS. In fact, albeit being in the higher end, the mean and median values of our aggregate bootstrap mode elasticities – equalling 3.5 and 3.4 for imports when all sectors are included – fall within the ballpark of values used for calibration. For exports, the same values are somewhat higher – with the mean equalling 4.0 and the median 3.8 – but they are still below the values that are reported in the earlier cited studies using disaggregate trade data. Hence, according to our estimates,

Table 5: Aggregated bootstrapped elasticities of substitution of exports, using country-specific weights together with country-specific or all German elasticities

Country	Mean		Median		Mode	
	Country-spec.	All DE	Country-spec.	All DE	Country-spec.	All DE
Austria	5.3	4.4	4.0	3.5	3.8	3.4
BeLux	3.4	6.3	3.1	3.7	3.0	3.4
Bulgaria	7.4	3.9	3.4	3.6	3.0	3.5
Cyprus	22.0	3.2	6.4	3.0	5.2	2.9
Czech Republic	9.5	4.3	4.3	3.4	3.8	3.3
Denmark	4.2	6.6	3.6	4.0	3.4	3.6
Estonia	19.2	3.8	7.9	3.4	6.5	3.4
Finland	4.6	3.5	3.5	3.2	3.4	3.2
France	8.9	34.8	4.4	9.6	3.8	6.7
Germany	13.3	13.3	5.3	5.3	4.3	4.3
Greece	22.0	8.5	4.6	4.3	4.1	3.7
Hungary	7.7	3.5	4.5	3.2	4.2	3.1
Ireland	3.4	5.1	3.3	3.3	3.2	3.1
Italy	4.7	7.4	3.4	4.0	3.2	3.6
Latvia	19.3	3.9	6.0	3.3	5.1	3.2
Lithuania	14.0	3.6	6.0	3.5	4.9	3.4
Malta	34.4	14.5	5.0	5.1	4.4	4.5
Netherlands	5.3	3.6	3.7	3.2	3.5	3.1
Poland	9.0	5.0	5.3	3.7	4.7	3.5
Portugal	5.4	3.4	4.1	3.2	3.9	3.2
Romania	6.4	5.8	3.4	3.8	3.2	3.5
Slovakia	7.8	3.6	4.3	3.2	3.9	3.1
Slovenia	10.8	3.8	4.4	3.1	3.9	3.0
Spain	6.7	8.7	3.5	4.3	3.2	3.8
Sweden	25.5	7.2	5.2	4.1	4.5	3.7
UK	6.0	16.0	3.3	5.8	3.0	4.5

some of the discrepancy between the values observed in the macro and trade literature can be attributed to methodology.

To shed further light on the cross-country differences in the estimates, Tables 4 and 5 compare the aggregate elasticities obtained using country-specific weights and estimates, and the ones obtained allowing only the weights to differ. This comparison allows us to disentangle the differences stemming from the sectoral composition of a country from the differences in the actual elasticities. We use Germany as the benchmark, and repeatedly aggregate its sector level elasticity estimates using the weights of each country. Comparing the individual countries' elasticities, aggregated this way, to the German elasticity aggregate, we are able to tell whether their imports or exports are oriented more towards high-elastic or low-elastic sectors than the German one. We note that nearly all countries are importing more low-elasticity goods than Germany; if we believe that technological content is negatively correlated with elasticity, this would imply that Germany imports relatively little high-technology goods. Even on the export side, however, we observe that nearly all of the countries are actually oriented *more* towards low-elasticity sectors than Germany. Judging by these results, hence, the trade orientation of these countries actually acts in their favor compared to German trade orientation. To the extent Germany is performing relatively better on the export market, our results indicate that this should be attributed to higher competitiveness within sectors, rather than the trade orientation itself. Comparing next, for each country, the two aggregates based on country-specific and on German elasticities, we see that the export data aggregates in nearly all cases are lowered when we use German elasticities for aggregation. There are some exceptions, however, most notably France. Looking closer at the sector level results, we note that sector *3530: Aircraft and spacecraft* is mainly responsible for the large increase. The elasticity estimate for this particular sector is significantly higher for Germany – 31.5 compared to 8.6 for France – while, at the same time, it has a high weight in the French aggregation – approximately 12%. The same applies to the UK, where sector *3530* has a weight of 5% and the country-specific estimate is 8.0. These results stress the importance for the aggregate of not overestimating the sector level elasticities, in particular for high-elasticity sectors. The change in the aggregate elasticity for France, for example, is considerable, and would have sizable effects in a calibration context, and yet it mainly comes from the change in the estimate of one single sector. For imports, most of the aggregates display small changes. There are, however, some exceptions. The aggregate estimate for Estonia decreases notably when German aggregation weights are used. This is mostly due to the German elasticity being considerably lower in the two sectors *3610: Furniture* and *1520: Dairy products*, even though several other sectors contribute to the change as well. For Cyprus, the decrease in the aggregate is also pronounced. This is due to a higher price sensitivity compared to Germany in a number of sectors, out of which *1511: Processing/preserving of meat* accounts for the largest part of the change. For the UK, on the other hand, the aggregate is notably increased. Here, most of the difference is explained by the sector *3410: Motor vehicles*, for which German imports are considerably more elastic than UK ones. Unsurprisingly, we observe a convergence in the aggregate elasticities across countries when we use only

German sector-level elasticity aggregates in the aggregation. However, for exports the elasticities are only marginally closer to each other when looking at the mode estimates and even slightly more dispersed when looking at the medians. The standard deviation of the aggregate estimates for the bootstrap mode (median) goes from 0.83 (1.17) and 0.75 (1.34). For imports, the corresponding numbers are 0.59 (0.93) to 0.43 (0.52). Hence, on the export side, most of the cross-country heterogeneity can be attributed to differences in sectoral composition across the different countries. On the import side, however, this result is not as strong; for imports, it seems to be the case that the differences in the actual elasticities across countries play a crucial role, in addition to the compositional effect.

7.4 Calibration of macroeconomic models

The results of our estimation exercise, once properly aggregated, can be used for calibration of macroeconomic models. As discussed briefly in the introduction, there is an ongoing debate about the true magnitude of the elasticities of substitution, with a ridge between the trade and the macro literature. On one side, there are the low elasticity estimates traditionally obtained from time series data that have also been extensively used for calibration of macroeconomic models. On the other, there are the more recent elasticity estimates obtained with more modern econometric techniques, which are considerably higher. We will devote this section to an attempt to bridge between these two literatures. We have already shown that there may be caveats to the estimation, and that using more robust methods leads to lower elasticity aggregates, which fall in between these two extremes. Here, we discuss the conceptual differences between the definitions of the elasticities in macroeconomic models and the empirical literature.

Elasticities in macroeconomic models are usually aggregated in a different order than we assume in estimation, and in the aggregation of empirical estimates. In the macroeconomic literature, multiple sectors are often not explicitly modeled, but it is implicitly assumed that all goods/sectors are first aggregated into a domestically produced and an imported basket, and that these, in turn, are aggregated into total consumption using some elasticity of substitution. In practice, this is not what the world looks like and we can never directly estimate this elasticity from the data. In reality, trade takes place within sectors, and the imports of each separate good are substituted for domestically produced varieties, or vice versa. Non-tradables are sometimes modeled to account for domestically produced goods that are not traded internationally, but there are also goods that are imported and that cannot be substituted for domestically produced ones, certainly not in the short or medium run. This could be due to lack of natural resources, differing technologies, climate, etc. Our estimates are based on actual trade transactions using the Armington assumption, so they are conditional on the hypothesis that a good is substitutable. In other words, we look at *realized* trade, hence our estimates might possibly overstate the true aggregate elasticity, even though they may be true for goods for which there is a domestically produced substitute. Specifically, what needs to be fulfilled for the elasticity estimate for a certain good to be accurate, is that the good is produced domestically, and that it can be produced in large enough quantities, as was discussed in the beginning of Section 6. Hence, to the extent the goods that are imported but not domestically produced are important, we may be overstating the true aggregate elasticity of substitution of a country's imports.

Our estimates of the aggregate elasticity of substitution of imports are centered at 3.5, with a minimum of 2.6 and a maximum of 4.8. Thus, for many European countries, an elasticity of 3 or 3.5 seems suitable for calibration. The aggregate elasticities for exports are centered at 3.9, with a minimum of 3.0 and a maximum of 6.5. The import and export elasticity estimates are, hence, generally fairly close to each other. Elasticities do vary across countries, however, and this needs to be taken into account when country-specific results are called for. Moreover, there are non-negligible differences in the import and export estimates for some countries, in particular Latvia, Lithuania, Malta, and Estonia. This suggests that a symmetric calibration of the elasticity of substitution of a country's imports and that of its export, although common, may occasionally be of very poor accuracy. Finally, the results we obtain support to some extent the critique that the elasticities of substitution in macroeconomic models are often assigned too low values in calibration, following calibration tradition based on outdated estimates which are likely to suffer from endogeneity and aggregation biases. Our suggested values for calibration are however not nearly as high as the ones suggested by, for example, Imbs and Méjean (2009). In the light of the arguments presented in Section 6 and in the present section, we further claim that there may be reason to further adjust these estimates downward, because of the conceptual differences. We cannot, however, quantify the importance nor the magnitude of these adjustments, why we view the presented aggregate elasticities as the best attainable estimates.

8 Conclusions

We present and apply a new empirical strategy, based on the ideas of Feenstra (1994), which enables the estimation of elasticities of substitution of exports on disaggregated sector level data. We estimate and subsequently aggregate sector level elasticities of substitution for the EU27 countries, of both each country's imports and that of its exports on the world market. Using the bootstrap, we are also able to shed light on some potential problems with the Feenstra (1994) method, while at the same time offering a possible solution. We find a high degree of sectoral heterogeneity in our estimates, which partly also explains the differences in the trade reaction to relative prices across countries. Our elasticity estimates are on the low side of the micro-data estimates reported in the literature. We argue that this is due to the bootstrap alleviating the bias caused by the non-linear mapping from reduced-form to structural parameters, as the original estimator tends to produce elasticities that are, on average, upward biased. Taking into account the conceptual differences in the definitions of the relevant elasticities in the macroeconomic and the microeconomic studies, we further argue that our results are close to the more traditional values of the elasticity of substitution, in the ballpark of what is normally used for calibration.

It is generally assumed that the elasticity of substitution is a deep parameter, which does not vary over time. That is also what we assume in our estimations. Equipped with the rich data set we have used for the estimation exercise in this paper, we could potentially have a say regarding the plausibility of this assumption, in particular for the older euro area members. However, this is beyond the scope of this paper. For now, we only note that one may need to bear in mind that the introduction of the euro took place during the time period of our study, even if it is not clear if and how it may have affected our results. The role played by the euro for the trade dynamics of the EU countries is an interesting venue for future research.

References

ARMINGTON, P. (1969): "A Theory of Demand for Products Distinguished by Place of Production," *Staff Papers-International Monetary Fund*, 159–178.

BRODA, C., J. GREENFIELD, AND D. WEINSTEIN (2006): "From groundnuts to globalization: A structural estimate of trade and growth," *NBER Working Paper Series*, 12512.

BRODA, C. AND D. WEINSTEIN (2006): "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 121, 541–585.

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

CORBO, V. AND C. OSBAT (2012): "Optimism bias? The elasticity puzzle in international economics revisited," ECB Working Paper, forthcoming.

FEENSTRA, R. (1991): "New goods and index numbers: US import prices," *NBER Working Paper No. 3610.*

——— (1994): "New product varieties and the measurement of international prices," *The American Economic Review*, 84, 157–177.

GALLAWAY, M., C. MCDANIEL, AND S. RIVERA (2003): "Short-run and long-run industry-level estimates of US Armington elasticities," *North American Journal of Economics and Finance*, 14, 49–68.

HOUTHAKKER, H. AND S. MAGEE (1969): "Income and price elasticities in world trade," *The Review of Economics and Statistics*, 51, 111–125.

IMBS, J. AND I. MÉJEAN (2009): "Elasticity optimism," *IMF Working Paper Series.*

——— (2010): "Trade Elasticities: A Final Report for the European Commission," *European Economy-Economic Papers.*

MAMMEN, E. (1993): "Bootstrap and Wild Bootstrap for High Dimensional Linear Models," *The Annals of Statistics*, 21, pp. 255–285.

MARQUEZ, J. (1990): "Bilateral trade elasticities," *The Review of Economics and Statistics*, 72, 70–77.

McDaniel, C. and E. Balistreri (2003): "A review of Armington trade substitution elasticities," *Economie internationale*, 301–313.

Mohler, L. (2009): "On the Sensitivity of Estimated Elasticities of Substitution," *Working Paper, University of Basel.*

Mohler, L. and M. Seitz (2010): "The gains from variety in the European Union," *Munich Discussion Paper.*

Orcutt, G. (1950): "Measurement of price elasticities in international trade," *The Review of Economics and Statistics*, 32, 117–132.

Pierce, J. and P. Schott (2009): "Concording US harmonized system categories over time," Tech. rep., National Bureau of Economic Research.

Polak, J. (1950): "Note on the measurement of elasticity of substitution in international trade," *The Review of Economics and Statistics*, 32, 16–20.

Radelet, S. and J. Sachs (1998): "Shipping costs, manufactured exports, and economic growth," in *American Economic Association Meetings, Harvard University, mimeo.*

Tinbergen, J. (1946): "Some measurements of elasticities of substitution," *The Review of Economics and Statistics*, 28, 109–116.

A Technical appendix

A.1 Mapping estimated coefficients into elasticities: the theoretical restrictions on the parameters

Combining expressions (15) and (16) in the main text, we can obtain an expression for σ_{cg} in terms of θ_{1cg} and θ_{2cg}. In order to derive the restrictions on the parameters, Feenstra (1994) rewrites the above expressions in terms of the parameter ρ_{cg}, defined as in equation (33) in the main text, instead of ω_{cg}. This simplifies things since both an upper and a lower restriction can be obtained for ρ_{cg}, while ω_{cg} is only bounded below by zero. Noting that ρ_{cg} is increasing in ω_{cg}, we obtain the lower bound of $\rho_{cg} = 0$ by letting $\omega_{cg} = 0$. On the other extreme, we have

$$\begin{aligned} \rho_{cg}\big|_{\omega_{cg}\to\infty} &= \left.\frac{1}{\frac{1}{\rho_{cg}}}\right|_{\omega_{cg}\to\infty} = \left.\frac{1}{\frac{1+\omega_{cg}\sigma_{cg}}{\omega_{cg}(\sigma_{cg}-1)}}\right|_{\omega_{cg}\to\infty} \\ &= \frac{1}{0+\frac{\sigma_{cg}}{\sigma_{cg}-1}} = \frac{\sigma_{cg}-1}{\sigma_{cg}} < 1\,. \end{aligned} \tag{A.1}$$

Hence, it holds that $0 \leq \rho_{cg} < \frac{\sigma_{cg}-1}{\sigma_{cg}} < 1$. Rewriting the expressions for $\hat{\theta}_{1cg}$ and $\hat{\theta}_{2cg}$ in terms of $\hat{\rho}_{cg}$, next, yields

$$\hat{\theta}_{1cg} = \frac{\hat{\rho}_{cg}}{(\hat{\sigma}_{cg}-1)^2(1-\hat{\rho}_{cg})} \tag{A.2}$$

and

$$\hat{\theta}_{2cg} = \frac{2\hat{\rho}_{cg}-1}{(\hat{\sigma}_{cg}-1)(1-\hat{\rho}_{cg})}\,. \tag{A.3}$$

Dividing the square of (A.3) by (A.2), and solving the resulting equation, results in an expression for ρ_{cg} given by

$$\hat{\rho}_{cg} = \frac{1}{2} \pm \left[\frac{1}{4} - \frac{1}{4+\left(\hat{\theta}^2_{2cg}/\hat{\theta}_{1cg}\right)}\right]^{\frac{1}{2}}\,. \tag{A.4}$$

Inserting (A.4) into (A.3), we have

$$\hat{\sigma}_{cg} = 1 + \frac{2\hat{\rho}_{cg}-1}{\hat{\theta}_{2cg}(1-\hat{\rho}_{cg})}\,. \tag{A.5}$$

In order to insure that $\hat{\sigma}_{cg} > 1$, we must choose a value of $\hat{\rho}_{cg} > \frac{1}{2}$ when $\hat{\theta}_{2cg} > 0$, and a value of $\hat{\rho}_{cg} < \frac{1}{2}$ when $\hat{\theta}_{2cg} < 0$. Finally, knowing that $\hat{\rho}_{cg}$ and $(1-\hat{\rho}_{cg})$ must both be positive, it is clear from equation (A.2) that, as long as $\hat{\theta}_{1cg} > 0$, the restrictions on $\hat{\sigma}_{cg}$ and $\hat{\rho}_{cg}$ are always fulfilled. As noted int he main text, in some cases of highly negative $\hat{\theta}_{2cg}$, even a negative value of $\hat{\theta}_{1cg}$ could yield a theory-consistent $\hat{\sigma}_{cg}$.

A.2 Deriving the aggregation equation for imports

Inserting demand equation (1) in the main text along with its counterpart for the domestically produced variety c into equation (39) in the main text, yields

$$\begin{aligned}\sigma_c &= 1 + \frac{\partial \ln \sum_g \sum_{v \neq c} P_{cgv} \beta_{cgv}^{\sigma_{cg}-1} \left(\frac{P_{cgv}}{P_{cg.}}\right)^{-\sigma_{cg}} C_{cg}}{\partial \ln \psi_c} \\ &\quad - \frac{\partial \ln \sum_g P_{cgc} \beta_{cgc}^{\sigma_{cg}-1} \left(\frac{P_{cgc}}{P_{cg.}}\right)^{-\sigma_{cg}} C_{cg.}}{\partial \ln \psi_c} .\end{aligned} \tag{A.6}$$

While the aggregator of imported and domestic goods within sectors is assumed CES, we assume the aggregator between sectors to be given by the following Cobb-Douglas function,

$$C_{c..} = \prod_{g \in G} \frac{C_{cg.}^{\alpha_{cg}}}{\alpha_{cg}^{\alpha_{cg}}} , \tag{A.7}$$

where α_{cg} is the expenditure share of good g in country c's total consumption. This restricts the elasticity of substitution between sectors to one, yielding a sequence of demand function for $C_{cg.}$ in terms of $C_{c..}$ given by[16]

$$C_{cg.} = \alpha_{cg} \frac{P_{c..}}{P_{cg.}} C_{c..} . \tag{A.8}$$

Under the assumption that

$$\frac{\partial \ln C_{c..}}{\partial \ln \psi_c} = 0 , \tag{A.9}$$

we can rewrite equation (A.6) to obtain an expression for $\sigma_c - 1$ as given by equation (40) in the main text. Since it is the long-run elasticities that are estimated, it is assumed that the domestic price will react fully to the change in costs, i.e. the representative domestic producer modifies her price. Foreign produces, on the other hand, are assumed

[16] In Imbs and Méjean (2010), the aggregator between sectors is assumed to take on the more general CES form, allowing the elasticity of substitution between sectors to assume values different from unity. They display results assuming that the elasticity of substitution between sectors equals 1, 2 and 0.5. However, since we do not know with certainty what this elasticity should be, we choose to hold on to the Cobb-Douglas modeling choice for simplicity.

not to respond at all to domestic costs changes in country c. Hence,

$$\frac{\partial \ln P_{cgc}}{\partial \ln \psi_c} = 1 \qquad \text{and} \qquad \frac{\partial \ln P_{cgv}}{\partial \ln \psi_c} = 0\,, \forall v \neq c\,. \tag{A.10}$$

It then holds that

$$\frac{\partial \ln P_{cg.}}{\partial \ln \psi_c} = \beta_{cgc}^{\sigma_{cg}-1} \left(\frac{P_{cgc}}{P_{cg.}}\right)^{1-\sigma_{cg}}, \tag{A.11}$$

and

$$\frac{\partial \ln P_{c..}}{\partial \ln \psi_c} = \sum_{g \in G_c} \alpha_{cg} \beta_{cgc}^{\sigma_{cg}-1} \left(\frac{P_{cgc}}{P_{cg.}}\right)^{1-\sigma_{cg}}. \tag{A.12}$$

Using the assumed demand equations, and noting that $\sum_g n_{cgc} = 1$, we can finally rewrite equation (40) as[17]

$$\begin{aligned} \sigma_c = & \sum_g m_{cgc}\sigma_{cg} + \sum_g (m_{cg.} - m_{cgc})(\sigma_{cg} - 1)(1 - w_{cg.}^M) \\ & + \sum_g (m_{cg.} - m_{cgc})\alpha_{cg}(1 - w_{cg}^M)\,, \end{aligned} \tag{A.13}$$

which is the aggregation equation we use in practice, presented as equation (44) in the main text.

[17]Note here that we have made a somewhat different assumption about how consumption is affected by the domestic cost shock than in Imbs and Méjean (2009). We have assumed that real total consumption remains unchanged in response to a shock in domestic costs, while still allowing for a non-zero response in the nominal consumption $P_{c..}C_{c..}$ which is represented by the last term in (44). Imbs and Méjean (2009), on the other hand, assume in their derivations that it is the nominal consumption that remains constant, which explains the difference in our final equations. It is not obvious which of the assumptions that is optimal. In the short run, on one hand, it is reasonable to assume that the consumers' have a fixed budget constraint, given that financial markets are not complete internationally. In this case, assuming that the consumption expenditures are constant seems reasonable. However, since we are measuring the long-run response in data, we may not expect the consumers' budget to be fixed in nominal terms, due to inflation, wage increases etc. Even if real consumption is clearly not constant in the long run either, we believe that this assumption is at least as close to the truth as the assumption of constant nominal consumption expenditures.

A.3 Deriving the aggregation equation for exports

Inserting country v's demand equations into equation (47) in the main text, yields

$$\sigma_c^X \equiv 1 - \frac{\partial \ln \sum_g \sum_{v \neq c} P_{vgc} \beta_{vgc}^{\sigma_{vg}-1} \left(\frac{P_{vgc}}{P_{vg.}}\right)^{-\sigma_{vg}} C_{vg.}}{\partial \ln \psi_c} - \frac{\partial \ln \sum_g \sum_{v \neq c} \sum_{i \neq c} P_{vgi} \beta_{vgi}^{\sigma_{vg}-1} \left(\frac{P_{vgi}}{P_{vg.}}\right)^{-\sigma_{vg}} C_{vg.}}{\partial \ln \psi_c}. \tag{A.14}$$

Assuming again, as in the case of imports, that the real consumption remains unchanged in response to the shock, i.e. that

$$\partial \ln C_{v..} / \partial \ln \psi_c = 0 \,, \; \forall v \,, \tag{A.15}$$

we can rewrite (A.14) to obtain equation (48) in the main text. As the assumed shock is the same as in the case of imports, i.e. a shock to the price of all goods produced in country c, it holds that

$$\frac{\partial \ln P_{vgc}}{\partial \ln \psi_c} = 1 \qquad \text{and} \qquad \frac{\partial \ln P_{vgi}}{\partial \ln \psi_c} = 0 \,, \; \forall i \neq c \,. \tag{A.16}$$

Noting that

$$\frac{\partial \ln P_{vg.}}{\partial \ln \psi_c} = \beta_{vgc}^{\sigma_{vg}-1} \left(\frac{P_{vgc}}{P_{vg.}}\right)^{1-\sigma_{vg}} , \tag{A.17}$$

$$\frac{\partial \ln P_{v..}}{\partial \ln \psi_c} = \sum_{g \in G_v} \alpha_{vg} \beta_{vgc}^{\sigma_{vg}-1} \left(\frac{P_{vgc}}{P_{vg.}}\right)^{1-\sigma_{vg}} , \tag{A.18}$$

and that $\sum_g \sum_{v \neq c} x_{vgc} = 1$, and combining with the demand equations for country v, we obtain

$$\sigma_c^X = \sum_g \sum_{v \neq c} x_{vgc} \sigma_{vg} + \sum_g \sum_{v \neq c} (x_{vgc} - x_{vg}^W)(1 - \sigma_{vg}) \frac{P_{vgc} C_{vgc}}{P_{vg.} C_{vg.}} - \sum_g \sum_{v \neq c} (x_{vgc} - x_{vg}^W) \frac{P_{vg.} C_{vg.}}{P_{v..} C_{v..}} \frac{P_{vgc} C_{vgc}}{P_{vg.} C_{vg.}} \,. \tag{A.19}$$

Remembering that we have assumed $\sigma_{vg} = \sigma_{cg}^X$, for all v, the above expression becomes

$$
\begin{aligned}
\sigma_c^X &= \sum_g x_{cg}\sigma_{cg}^X + \sum_g \sum_{v \neq c} (x_{vgc} - x_{vg}^W)(1 - \sigma_{cg}^X) \frac{P_{vgc}C_{vgc}}{P_{vg.}C_{vg.}} \\
&\quad - \sum_g \sum_{v \neq c} (x_{vgc} - x_{vg}^W) \frac{P_{vg.}C_{vg.}}{P_{v..}C_{v..}} \frac{P_{vgc}C_{vgc}}{P_{vg.}C_{vg.}}, \qquad \text{(A.20)}
\end{aligned}
$$

where

$$
x_{cg} \equiv \sum_{v \neq c} x_{vgc} \qquad \text{(A.21)}
$$

is the share of good-g exports in the total exports of country c. Using the same notation as in the case of import aggregation, we denote by

$$
\alpha_{vg} \equiv \frac{P_{vg.}C_{vg.}}{P_{v..}C_{v..}} \qquad \text{(A.22)}
$$

the expenditure share of good g in country v's total consumption. Defining the share of country-c exports in country v's total consumption of good g as shown in equation (55) in the main text, the final equation for the aggregate elasticity of substitution for country c's exports is given by equation (54) in the main text.

A.4 The weights used for aggregation of the sector level elasticity estimates

In order to be able to aggregate the import data elasticities in a model-consistent way, we need data on consumption shares, in addition to the imports data on bilateral trade flows. Specifically, we need to calibrate the following weights for the aggregation of import data elasticities: m_{cgc}, m_{cg}, w_{cg}^M and α_{cg}. Given the definitions of w_{cg}^M and α_{cg}, we can write m_{cgc} and m_{cg} as follows:

$$
\begin{aligned}
m_{cgc} &\equiv \frac{P_{cgc}C_{cgc}}{\sum_g P_{cgc}C_{cgc}} = \frac{\frac{P_{cgc}C_{cgc}}{P_{c..}C_{c..}}}{\frac{\sum_g P_{cgc}C_{cgc}}{P_{c..}C_{c..}}} = \frac{\frac{P_{cg.}C_{cg.}}{P_{c..}C_{c..}}\frac{P_{cgc}C_{cgc}}{P_{cg.}C_{cg.}}}{\sum_g \frac{P_{cg.}C_{cg.}}{P_{c..}C_{c..}}\frac{P_{cgc}C_{cgc}}{P_{cg.}C_{cg.}}} \\
&= \frac{\alpha_{cg}(1 - w_{cg}^M)}{\sum_g \alpha_{cg}(1 - w_{cg}^M)} \qquad \text{(A.23)}
\end{aligned}
$$

$$
\begin{aligned}
m_{cg} &\equiv \frac{\sum_{v\neq c} P_{cgv} C_{cgv}}{\sum_g \sum_{v\neq c} P_{cgv} C_{cgv}} = \frac{\frac{\sum_{v\neq c} P_{cgv} C_{cgv}}{P_{c..}C_{c..}}}{\frac{\sum_g \sum_{v\neq c} P_{cgv} C_{cgv}}{P_{c..}C_{c..}}} = \frac{\frac{P_{cg.}C_{cg.}}{P_{c..}C_{c..}} \frac{\sum_{v\neq c} P_{cgv} C_{cgv}}{P_{cg.}C_{cg.}}}{\sum_g \frac{P_{cg.}C_{cg.}}{P_{c..}C_{c..}} \frac{\sum_{v\neq c} P_{cgv} C_{cgv}}{P_{cg.}C_{cg.}}} \\
&= \frac{\alpha_{cg} w_{cg}^M}{\sum_g \alpha_{cg} w_{cg}^M} \,. \qquad \text{(A.24)}
\end{aligned}
$$

Hence, what we need for aggregation is only the import shares w_{cg}^M and the expenditure shares α_{cg} for each country c. These we obtain from our estimation data in combination with the United Nations' UNIDO data on internal production.

For a model-consistent aggregation of the export-data elasticity estimates, we need to calibrate the following weights: x_{cg}, x_{vgc}, x_{vg}^W, w_{vg}^X and α_{vg}. The first two, i.e. the share of sector g in country c's total exports and the variety share of each variety v in country c's total exports, we can obtain from the bilateral trade data that we use for estimation. The rest of the weights, however, need to be calibrated using other data sources. Note first that we can rewrite the weight x_{vg}^W as

$$
\begin{aligned}
x_{vg}^W &\equiv \sum_{i\neq c} x_{vgi}^W \\
&= \frac{\sum_{i\neq c} P_{vgi} C_{vgi}}{\sum_g \sum_{v\neq c} \sum_{i\neq c} P_{vgi} C_{vgi}} \\
&= \frac{P_{vg.}C_{vg.} - P_{vgc}C_{vgc}}{\sum_g \sum_{v\neq c} (P_{vg.}C_{vg.} - P_{vgc}C_{vgc})} \\
&= \frac{P_{vg.}C_{vg.}(1 - w_{vg}^X)}{\sum_g \sum_{v\neq c} P_{vg.}C_{vg.}\left(1 - w_{vg}^X\right)} \,. \qquad \text{(A.25)}
\end{aligned}
$$

Since we can calibrate $P_{vgc}C_{vgc}$ for all v and g from our export data, all we need to collect are the total sector expenditures for each country $v \in V_{cg}$, i.e. $P_{vg.}C_{vg.}$ for each country for which data is included in the estimation, all expressed in the same currency. Having access to these, we can then construct the α_{vg} directly, and the w_{vg}^X and x_{vg}^W using the sector expenditures together with the export data.

Since what we have access to is production data on one hand, and trade data on the other, we calculate the weights w_{vg}^X as

$$
w_{vg}^X = \frac{X_{vgc}}{Y_{vg} - X_{vg} + M_{vg}} \,, \qquad \text{(A.26)}
$$

where X_{vgc} denotes the value of country c's exports of good g to country v, $Y_{vg} - X_{vg}$ denotes the value of the domestic consumption of good g in country v, i.e. the value of total production less exports, and M_{vg} denotes the value of country v's total imports of good g. Hence, the term in the numerator correspond to the theoretical term $P_{vgc}C_{vgc}$ and the one in the denominator to the theoretical term $P_{vg.}C_{vg.}$. The weights α_{vg} we calculate as

$$\alpha_{vg} = \frac{Y_{vg} - X_{vg} + M_{vg}}{\sum_g (Y_{vg} - X_{vg} + M_{vg})}\,, \tag{A.27}$$

where the term in the denominator now corresponds to $P_{v..}C_{v..}$. Finally, we compute x^W_{vg} as

$$x^W_{vg} = \frac{(Y_{vg} - X_{vg} + M_{vg})(1 - w^X_{vg})}{\sum_g \sum_{v \neq c}(Y_{vg} - X_{vg} + M_{vg})\left(1 - w^X_{vg}\right)}\,, \tag{A.28}$$

while x_{vgc} and x_{vg}, being computing using export data only, do not require any further comment.[18]

[18] It could of course be the case that some countries are importing a certain good g from some other exporter $i \notin V_{cg}$. Since our estimation data only provide total trade for the EA countries, and not for all of their trading partners, we need yet another source of data for the calibration of our aggregation weights. The CEPII BACI data serve this purpose, as they contain the total trade in each sector for all of the more than 200 listed countries.

B Data appendix

B.1 The Eurostat COMEXT data

Our database is based on Eurostat's COMEXT, which contains monthly observations on values and quantities of imports and exports reported by all European Union countries from and to up to 270 trading partners. The full database is available at a disaggregation level of 8 digits in the *Combined nomenclature* (CN), based on the Harmonized System, which however goes only up to 6 digits (HS-6).

For each country-partner-sector triplet at each point in time, COMEXT provides information on the value of each monthly transaction in ECU-EUR, the quantity in $1,000$ Kg and, if available, the corresponding Special Units (which vary by sector and can be items, liters, meters, etc). The reporting of quantities is not always consistent, so that often values without a corresponding quantity are observed. These end up as missing values in our sample, because one of our main variables is the unit value of imports, which we obtain by dividing the value by the quantity. The unit values at the 8-digit level of disaggregation are the variable we use to clean the data from outliers, which are an endemic feature of this database.

B.2 Mapping between nomenclatures (CN→HS→ISIC)

Since we wish to perform the outlier cleaning at the highest level of disaggregation, as discussed further in Section B.3 below, we start from the 8-digit raw data and aggregate that to suit our needs. Our goal is to ultimately aggregate the sector level estimates, and, since the weights available to us are classified according to the *International Standard Industrial Classification* (ISIC) and up to 4 digits of disaggregation, we need these CN-8 sectors to be ultimately aggregated into ISIC-4 sectors. As both these nomenclatures can be linked to the *Harmonized System* (HS), specifically at the 6-digit level, we first map the HS-6 codes into ISIC-4 ones and then the CN-8 codes into the resulting groups of HS-6 codes corresponding to each ISIC-4 category.[19]

The HS nomenclature has been subject to three revisions during the period of our study, in the years of 1996, 2002 and 2007. Major changes were also made in 2010, which explains the choice of our ending point. For any revision, we risk having to drop some data due to there being no clear mapping between the different nomenclatures over our entire sample. We begin by transforming all HS revisions to the 2002 one, and then

[19] At the 8-digit level, the CN contains more than $15,000$ codes, while the HS comprises some $4,700$ codes at the 6-digit level. The number of ISIC codes at the 4-digit level is 127 (for manufacturing).

group the 2002 HS-6 codes into their corresponding ISIC-4 codes.[20] In this process we loose nine ISIC-4 codes in total, due to not being able to track their content throughout the entire time period.[21] The mapping of HS6 codes into ISIC-4 ones provides us with a list of HS6 codes to be aggregated into each and every one of the ISIC-4 categories. We thereafter move to grouping the CN-8 codes into ISIC-4.

In addition to the major revisions of the HS, which also affected the CN since the first six digits generally coincide with the HS-6 classification, there have been revisions of the 6-digit CN codes during the years 1993-95 and 1999. In addition to these, a large number of changes to the 8-digit codes have been made each year; as long as the first six digits remained unchanged, however, we do not need to keep track of these changes.[22] We make sure that the changes made to the CN-6 codes, in addition to the HS-6 changes discussed above, do not imply any shifts between ISIC-4 groupings. Finally, for each year in our sample, we aggregate the list of CN-8 codes pertaining to each of the 119 kept ISIC-4 codes, thereby obtaining time series of bilateral trade, sector by sector, for each of the 27 countries of our interest.

B.3 Outlier cleaning

We use a cross-sectional benchmark to identify outliers in the data. For each sector, we take all observations on import unit values of all 27 declarants from all partners; the outliers are those observations that lie "too far" from the median of this cross-section. Note that this procedure only makes sense at a very disaggregated level, when the unit values refer to goods that are as similar as it gets when dealing with trade data. As a metric of distance we use the absolute deviation from the median (mad), which is much more robust to outliers than the standard deviation around the mean. We aim

[20]Our concordance of the HS codes is similar to that presented in Pierce and Schott (2009), with a few relevant differences. To economize on data, we choose the 2002 revision as benchmark, instead of the first year of our sample as the overlap is larger between two consecutive revisions, than between revisions that are further apart. Moreover, as we merge the HS-6 codes into ISIC-4, we want to make sure that we don't throw away any HS-6 codes that have stayed within the same ISIC-4 category, even though we may not be able to obtain a time series for each specific HS-6 code. Finally, we note that the CN nomenclature differs from the HS one at higher levels of disaggregation, why we cannot rely on algorithms constructed for US data, but instead need to construct CN-specific concordances.

[21]These codes are *1712: Finishing of textiles*, *2230: Reproduction of recorded media*, *2421: Pesticides and other agro-chemical products*, *2731: Casting of iron and steel*, *2732: Casting of non-ferrous metals*, *2891: Metal forging/pressing/stamping/roll-forming*, *2892: Treatment and coating of metals*, *3710: Recycling of metal waste and scrap*, and *3720: Recycling of non-metal waste and scrap*.

[22]The amount of changes made at the 8-digit level is another reason for why we choose to work with somewhat more aggregated data. We would otherwise need to throw away all of the codes which have been split or merged on several occasions. In addition to this, we observe re-usage of some codes without any overlapping content. This makes it very tedious, and for many sectors even impossible, to create reliable time series on a large scale basis.

at eliminating a small percentage of observations, so that we progressively increase the number of *mad* around the median if we see that the procedure tends to eliminate too many observations. We start off with a distance of $(2 * 1.4785 * mad)$, but if more than 3 percent of observations are classified as outliers, we increase the cutoff from $(2 * 1.4785 * mad)$ to $(3 * 1.4785 * mad)$ and so on. We alternate between raising and lowering the cutoff by fractions of $(\text{cutoff} * 1.4785 * mad)$ for at most 100 times. If after 100 runs the percentage of classified outliers is still very high, we accept this as a sign of high variability and accept the algorithm's decision, keeping track of the percentage of outliers for every sector-year pair. We find that these extreme cases tend to occur in sector-year pairs with very few observations, which will in any case drop out of our analysis due to a low number of bilateral transactions.