Elasticity Optimism*

Jean Imbs †   Isabelle Méjean‡

May 2010

Abstract

In most macroeconomic models, the substitutability between domestic and foreign goods is calibrated using aggregated data. This imposes homogeneous elasticities across goods, and the calibration is only valid under this assumption. If elasticities are heterogeneous, the aggregate substitutability is a weighted average of good-specific elasticities, which in general cannot be inferred from aggregated data. We identify structurally the substitutability in US goods using multilateral trade data. We impose homogeneity, and find an aggregate elasticity similar in value to conventional macroeconomic estimates. It is more than twice larger with sectoral heterogeneity. We discuss the implications in various areas of international economics.

Keywords: Elasticity of Substitution, Trade Elasticities, Heterogeneity, Calibration.

JEL Classification: F41, F32, F21.

*We thank the Fondation Banque de France and the National Center of Competence in Research "Financial Valuation and Risk Management" for financial support. The National Centers of Competence in Research (NCCR) are a research instrument of the Swiss National Science Foundation. This paper was begun while Imbs was visiting the Hong Kong Institute for Monetary Research, whose hospitality is gratefully acknowledged, and while Méjean was at CEPII. We thank audiences at the Paris School of Economics, the European Central Bank, the Stockholm School of Economics, Ente Einaudi, the NBER 2008 Summer Institute, UCLA, UC Berkeley, UC Santa Cruz, the San Francisco Fed, HEC Paris, Insead, Sciences-Po and Cambridge, as well as George Alessandria, Nicolas Coeurdacier, Jonathan Eaton and Alessandro Rebucci for fruitful exchanges.

†Paris School of Economics, HEC Lausanne, Swiss Finance Institute and CEPR. Corresponding author: HEC Lausanne, Extranef 235, Lausanne Switzerland 1015. +41 21 692 3484, jimbs@unil.ch, www.hec.unil.ch/jimbs

‡IMF, Ecole Polytechnique, CREST and CEPR, isabelle.mejean@polytechnique.edu, www.isabellemejean.com

1 Introduction

Most calibrated models in international macroeconomics assume a representative agent, and a unique final good sector in all countries. There is a single, constant elasticitiy of substitution between domestic and foreign goods, and it is typically calibrated using aggregate data. That calibration choice is important. Depending on the value assigned to the parameter, the predictions of virtually any calibration exercise with an international dimension change quantitatively, sometimes dramatically.

Unsurprisingly, the value of the elasticity draws from decades of empirical work. But little consensus has emerged from the effort, except for two broad conclusions. First, elasticity estimates inferred from aggregate data are barely positive. Second, there are enormous differences between goods. Long time ago, Orcutt (1950) referred to an "elasticity pessimism", which he related to the low observed volatilities in aggregate quantities, and the high volatility of international relative prices. Here we show inferring the parameter from aggregate data actually imposes homogeneity. The constraint can create a conventional heterogeneity bias, which we show is the main reason why aggregate elasticity estimates are close to zero in aggregate US data. We propose a correction of the parameter that accounts for heterogeneity at the microeconomic level.[1]

In virtually all macroeconomic models, the elasticity of substitution between domestic and foreign bundles of goods is calibrated on the basis of estimates that arise from aggregate data. But an elasticity obtained from aggregated data implicitly imposes homogeneity across sectors, because the estimation is in fact performed on quantities and price indices that have been aggregated up to the country level. With aggregate data, the estimation can only hope to identify *the* elasticity of substitution between the domestic and foreign bundles of goods produced in different sectors. Cross-sector aggregation is performed first, and the only international substitutability that can prevail is between aggregate bundles of goods. In the words of Kreinin (1967): "It is a major shortcoming of most estimates that they are concerned with demand for manufactured products as a whole. Elasticity estimates for individual commodities are rare. This lack imposes a severe constraint on attempts to quantify the effects of policy measures on the volume of trade and economic welfare. Of necessity most such studies apply a common elasticity

[1]The parameter we propose to correct is the *single* elasticity of substitution in a one-sector open economy model. Naturally it is also possible to account for sectoral heterogeneity in preferences introducing multiple sectors in the theory. Such an approach, which we view as complementary, is sometimes followed, but often at the cost of tractability, especially in an open economy context.

figure to all commodity groups." (p.510).

Identification issues notwithstanding, the aggregate elasticity of substitution σ can be inferred from the observed response of imported aggregates to appropriately identified shifts in price indices. But the inference is *only* valid under the restriction that the sector-specific elasticities of substitution, σ_k, are equal across sectors. It is undisputable that the price elasticity of aggregate imports is close to zero in the data; but jumping to the conclusion that σ is, as well, is only warranted under the restriction that $\sigma_k = \sigma$ for all sectors k.[2] Given the overwhelming evidence that estimates of σ_k vary by sector, we argue this is an implausible assumption.

With heterogeneous elasticities of substitution across sectors and CES preferences, we show the aggregate elasticity is given by an appropriately weighted average of sector-specific elasticities σ_k, say $\sum_k \omega_k \sigma_k$, with weights ω_k given by theory. The weighted average should be understood as the representative agent's elasticity of substitution, computed in a way that accounts for the sector-level heterogeneity in σ_k. Whenever the estimates of σ based on aggregate data are significantly different from $\sum_k \omega_k \sigma_k$, the homogeneity constraint implicit in aggregate data effectively creates a conventional heterogeneity bias.[3] Then, the only option available to pin down an aggregate substitutability that accounts for heterogeneity is to estimate σ_k from disaggregated data, and infer from theory the weights used in aggregation. Using aggregate data to estimate the price elasticity of aggregate imports, as is customary in macroeconomics, will only identify an aggregate substitutability that imposes homogeneity.

Now it is in fact possible the estimates of σ obtained from aggregated data happen to pinpoint precisely the adequately weighted average of σ_k. Aggregate data then imply estimates of σ that happen to account for sectoral heterogeneity. Sector-level elasticities are still constrained to homogeneity, but at a level that happens to be the one implied by the theory with heterogeneity, i.e., $\sigma = \sum_k \omega_k \sigma_k$. The heterogeneity bias happens to be so small it cannot be detected. The question is therefore empirical: Do aggregate data imply an estimate of σ that happens to concord with the value of $\sum_k \omega_k \sigma_k$ implied by sectoral data? We show the answer for US data is negative. There is a heterogeneity bias, so that aggregating the *data* is fundamentally different from aggregating elasticity

[2] We show this formally for CES preferences in Appendix A

[3] A heterogeneity bias arises when homogeneity is forced on coefficient estimates that are in reality heterogeneous, as aggregating the data imposes. Then aggregating estimates is different from aggregating the data. The magnitude of the bias is an empirical matter. See for instance Pesaran and Smith (1995).

estimates, as we do in this paper.

We develop a model of CES preferences telling us how to properly aggregate microeconomic elasticities. Households have CES preferences over a continuum of goods, or sectors. Consumption in each sector is in turn a CES aggregate of different varieties, produced in different countries. Crucially, the elasticity of substitution between varieties is allowed to be sector-specific. We construct a measure of aggregate substitutability consistent with a representative agent choosing between country-level aggregates of domestic and foreign quantities. We show there is potentially a discrepancy between conventional macroeconomic estimates of the elasticity of substitution, imposing equal elasticities across goods, and aggregate estimates allowing for heterogeneity.

A recent literature has argued the observed response of traded quantities to price changes is in fact unrelated with the elasticity of substitution given in preferences. This happens in Chaney (2008), Eaton and Kortum (2002), or Arkolakis, Costinot and Rodriguez-Clare (2009), when endogenous supply decisions on the part of heterogeneous firms are introduced. Trade elasticities then depend directly on the distribution of firm heterogeneity, a supply parameter. For instance, Dekle, Eaton and Kortum (2008) examine the consequences on global supply of a resolution of global imbalances in a framework where the international substitutability of goods is entirely irrelevant because of firm entry decisions.

The interpretation of the heterogeneity bias we document is best understood in the context of a model with CES preferences and a representative firm. There, the bias has implications for the calibration of preferences. Our results are not meant to vindicate such a modeling approach. Within the confines of a CES demand system with a representative firm, we point to the possibility of a heterogeneity bias in venerable estimates of an important parameter. As most *calibrated* models in international macroeconomics continue to impose a representative firm and CES preferences, we argue the implications of a heterogeneity bias are of practical relevance. That said, we show in section A of a Technical Appendix that our interpretation of trade elasticities in terms of a preference parameter carries through in a slightly generalized version of Chaney (2008). We find the response of traded quantities continues to reflect the substitution elasticity in preferences, even with endogenous supply decisions and firm entry. This obtains provided the substitutabilities between firms and countries are allowed to be different. Such generalization suggests it is *possible* to construct a conventional model with firm heterogeneity where trade elasticities and CES continue to be directly related. Both then suffer from a heterogeneity bias.[4]

Our second goal is to estimate disaggregated elasticities. To do so, we borrow from a methodology introduced by Feenstra (1994) and recently implemented by Broda and Weinstein (2006). In the context of our model, the approach can be used to identify the parameter of interest. Demand at a sectoral level is given by a CES aggregator of domestic and foreign goods varieties. Under an Armington assumption, the substitutability of domestic and foreign varieties is the same as the substitutability between two foreign varieties. As a consequence, we are able to identify sectoral elasticities of substitution using the observed cross-country variation in the trade flows towards a given importer. In comparison with conventional approaches, this provides estimates that are structural, and do not fall victim to the endogeneity concerns that plague any regression of (relative) quantities on (relative) prices.

Armed with an econometric methodology to estimate σ_k and a theory to aggregate them, we conduct two experiments. First, we constrain the estimates to equality across sectors. We expect the results to be in line with the macroeconomic literature that has used aggregate data to infer the aggregate elasticity of substitution. This puts our theory to the test of its ability to reproduce conventional macroeconomic estimates under conventional macroeconomic assumptions. Second, we relax the homogeneity constraint, and obtain estimates of $\sum_k \omega_k \sigma_k$. If a significant difference exists with constrained estimates, a heterogeneity bias prevails in the data. Then, macroeconomic data cannot be used to obtain a substitutability measure that reflects the sectoral heterogeneity in σ_k.

In macroeconomic applications, calibration exercises typically favor values of the parameter that are inferred from aggregate data, and chosen as "plausible mid-points" to the wide range of estimates the literature has uncovered. For instance, Obstfeld and Rogoff (2005) use a value of 2; Backus, Kehoe and Kydland (1994) [BKK] use 1.5, as do Chari, Kehoe and McGrattan (2002), and it is set between 0.6 and 2 in Coeurdacier, Kollmann and Martin (2007). When all elasticities are forced to be equal across sectors, our approach generates aggregate estimates between 2.5 and 3 for the US. The range is

[4] Our results also suggest a heterogeneity bias will affect the parameter governing the distribution of producing firms in the original version of Chaney (2008). If firm distributions are heterogeneous across sectors, an estimate based on the whole universe of firms at the aggregate level will suffer from a heterogeneity bias.

within the ballpark of the calibrated values used in the macroeconomic literature, and within the range of estimates obtained from aggregate data.[5] With heterogeneity however, aggregate elasticity estimates more than double, with values in the US up to 6 or 7. Such significant difference demonstrates a heterogeneity bias is at play in aggregate US data. It also tells us an aggregate elasticity of substitution that accounts for sectoral heterogeneity is closer to 7 than to 2.

We argue σ is a preference parameter, whose calibration ought to be motivated by adequate structural estimates. The alternative is to consider σ a free parameter, whose calibration is determined by the model's ability to match particular moments of interest in the data. Most prominently, σ is often chosen in such a way that the model matches the response of aggregate imports to shocks in relative prices. Clearly, the corrected estimate we propose will not perform well along this margin. This should not be surprising. The paper's point is precisely to argue heterogeneity drives a wedge between estimates of the price elasticity of imports and an elasticity of substitution that accounts for sectoral heterogeneity. Such calibration choice will by construction assume away a heterogeneity bias, which we argue is problematic given the evidence on sectoral heterogeneity in σ_k. In general structural estimates ought to be preferred for the calibration of preference parameters, provided the underlying model is general enough. We argue CES preferences continue to be the standard in open economy macroeconomics.

The comparison between constrained and unconstrained estimates focuses squarely on the importance of heterogeneity. We implement the same estimator, on the same data set, with the sole difference that we impose or not a homogeneity constraint. Potential alternative explanations for a discrepancy between the elasticities inferred from microeconomic or macroeconomic data are effectively held constant. When we perform our estimation on an aggregated version of our data, we find estimates for the elasticity of substitution near one. This is close to the results obtained with a homogeneity constraint, and close to conventional macroeconomic estimates. But the comparison is less valuable, since aggregating the data obscures the ceteris paribus nature of our main experiment. For instance, aggregation substantially reduces the dimensionality of the data relevant to identification.

Does heterogeneity matter economically? We discuss some illustrations in several areas

[5] See Goldstein and Kahn (1985).

of international macroeconomics. Whether the elasticitiy of substitution is in fact 2 or 7 does make a quantitative difference, and calibrated equilibrium responses often change sizably. In some instances, the correction is also relevant qualitatively. We close with an illustration of the heterogeneity bias constructed within the conventional international real business cycle framework due to BKK. We simulate a multi-sector version of the model, with heterogeneous sector-level elasticities σ_k and calibrated weights ω_k. We then ask from the conventional one-sector version of BKK what calibrated value of the elasticity of substitution is best able to reproduce the data implied by a multi-sector model. We find the match is best for an elasticity given by $\sum_k \omega_k \sigma_k$, i.e. a calibration that accounts for the heterogeneity bias we uncover.

In what follows, we first present the model used to guide the aggregation of industry specific elasticities. Section 3 discusses the identification of sector specific parameters, their aggregation and the data involved. We also describe our adaptation of Feenstra's approach. Section 4 reports our results, and document their relevance in recent standard models of international trade in goods and assets. We also ascertain their robustness. Section 5 concludes.

2 Trade Elasticities: Practice ahead of Theory

We open with a summary of the empirical literature concerned with estimating trade elasticities to infer substitutability. We discuss common practice, and how end estimates tend to be heterogeneous in disaggregated data. We then lay out the theory we use to map out conventional elasticity estimates with the utility parameter of interest in macroeconomic calibration exercises. The model tells us how to compute aggregate elasticities in ways that may or may not allow for heterogeneity at the good's or sector's level. It is constructed to also accommodate firm entry decisions.

2.1 Practice

Estimating import price elasticities is an old business in economics. Historically, estimates arising from sectoral or aggregate data have virtually always been the object of separate studies. The dichotomy is natural, since microeconomic trade elasticities typically consider the effect of a shift in sectoral price, holding constant the price of close substitutes, whereas macroeconomic elasticities consider shocks to all relative prices. In what follows, we are careful to consider the response of aggregate quantities to the same, macroeconomic, shock, allowing or not for heterogeneity in the response of quantities to a given shock in prices. Whether we impose homogeneity or not, we do so on the same data, characterized by an identical variance-covariance matrix, and on the basis of the same estimator. The comparison brings the focus squarely on the importance of heterogeneity in elasticity estimates.

Goldstein and Kahn (1985) survey cross-country evidence about trade elasticities, estimated on the basis of aggregated data. Their meta-analysis summarizes the results implied by aggregate data for a variety of countries, time periods and econometric techniques. They report estimates for import price elasticities rarely above 2 in absolute value. Such small estimates have largely remained unchanged since, even though the econometrics involved have become considerably more sophisticated. Thus, Marquez (1990) implements a frequency domain estimator, Gagnon (2003) instruments import prices using the real exchange rate, and Hooper, Johnson and Marquez (1998) use co-integration techniques. In most instances, estimates of the import price elasticity arising from aggregated data are found to be weakly negative, not always significant, and rarely larger than 2 in absolute value. The conventional inference in international macroeconomics is to deduce that the elasticity of substitution between aggregate bundles of domestic and foreign goods is close to zero as well. For instance, Backus, Kehoe and Kydland (1994) note that "there is some uncertainty about what value [of the elasticity of substitution] is indicated by the data. The most reliable studies seem to indicate that for the United States the elasticity is between 1 and 2" (p.91).

Studies that provide disaggregated elasticity estimates have been on the rise over the past decades, because of improved data availability. An early contribution is Kreinin (1967), with data on three sectors for ten industrial countries. Import elasticities at the sector level display considerable heterogeneity, ranging from -1.1 to -5.0 for instance in Kreinin (1967). Heterogeneity is largely confirmed in more recent contributions that merge disaggregated information on trade restrictions, traded quantities, and prices. For instance, Blonigen and Wilson (1999) document substitution elasticities between 0 and 3.52 in 146 US sectors. Microeconomic studies also open the door to instrumentation strategies where changes in international relative prices can be ascribed to events, such as trade liberalizations, whose magnitude and timing are arguably exogenous to each market's circumstances. Thus, using NAFTA and detailed information on good specific tariff changes, Head and Ries (2001) find disaggregated substitution elasticities between

8 and 12, while Romalis (2007) documents estimates between 4 and 13. In the words of Anderson and Van Wincoop (2004), "overall the literature leads us to conclude that the elasticity is likely to be in the range of 5 to 10" in disaggregated data (p.716).

Even the relatively infrequent papers that do estimate both macro and micro elasticities carefully avoid drawing a link between them. For instance, Houthakker and Magee (1969) dedicate a section to elasticities estimated on aggregate traded quantities and prices. They find import price elasticities close to zero for fifteen developed economies, and reject heterogeneity in estimates across countries. In a separate section, they estimate import elasticities for five US commodity classes broken down by degree of processing. They find substantial heterogeneity, with values ranging from 0 to -4.05. But they never explicitly or formally relate heterogeneity at the sectoral level and aggregate estimates. In footnote 8, they remark that "the price elasticity for total imports is considerably below a weighted average of the five separate elasticities; this is due to the fact that crude foods and crude materials, which have a low elasticity, are more variable than the other classes and consequently have a relatively greater influence on the relevant sums of squares and cross-products." The intuition resembles a heterogeneity bias.

We do not propose to dispute the argument that micro-based elasticity estimates can be different objects than the values obtained from macroeconomic data. For instance, Ruhl (2005) argues cross-sectional elasticity estimates are naturally higher, for they embed long run phenomena, such as firm dynamics and the associated adjustments in the quantities produced. Time series data, in contrast, tend to focus on high or medium frequency developments, and may overlook most entry or exit decisions. But in practice, disaggregated datasets tend to be cross-sectional, whereas aggregate ones have information over time. The apparent importance of aggregation may therefore be an artefact of data availability and differences in econometric methodologies. The parameters estimated in micro- and macro-economic studies can in fact be fundamentally different, since in practice they may not capture the extensive margin to the same extent.

In what follows, we are careful to accommodate this possibility. Both our corrected and conventional elasticity estimates arise from the same dataset. Both are affected by a putative extensive margin to exactly the same extent. In fact, Feenstra (1994) discusses how his empirical approach can accommodate time-varying number of firms in each exporting economy, an argument we clarify when we describe the estimation. In short, our homogeneous and heterogeneous estimations are conceptually similar, and they use the same dimension in the same data. Ruhl's (2005) argument cannot explain the discrepancies we identify in structural estimates of the elasticity of substitution. In fact, the experiment we propose is ceteris paribus in a general sense. We estimate homogeneous and heterogeneous elasticities on the same data set, and within the same methodology. Our econometric approach holds constant all other potential explanations for a difference between macro and micro elasticity estimates. If a discrepancy subsists, it will have to stem from heterogeneity. We zoom in onto the heterogeneity question, and point to its potential importance in macroeconomics.

Macroeconomists have in fact recently recognized such potential importance. When calibrating the elasticity of substitution, an increasing number of international macroeconomists eschew estimates of the parameter that were obtained from aggregate data. They refer instead to microeconomic studies. For instance, Obstfeld and Rogoff (2000) choose a value of 6, arguing it is consistent with disaggregated estimates. In their theory of the international diffusion of technology shocks, Corsetti, Dedola and Deluc (2008) consider a value of 8, once again "based on the estimates in the trade literature" (p.460). Coeurdacier (2009) chooses 5, as "the lower bound of estimates from the trade literature" (p.88), to study the impact of trade costs on aggregate portfolio choice. All these models are macroeconomic in nature, and all build upon the CES preferences that we assume here as well. Such a recent trend in calibration choices draws an implicit link between the heterogeneous disaggregated estimates and their macroeconomic counterpart. In this paper, we formalize the link.

2.2 Theory

2.2.1 Preferences

Consumption in economy j is an aggregate of imperfectly substitutable sectors $k = 1, ..., K$. Utility is given by $C_j = \left[\sum_{k\in K} (\alpha_{kj} C_{kj})^{\frac{\gamma-1}{\gamma}}\right]^{\frac{\gamma}{\gamma-1}}$ where α_{kj} denotes an exogenous preference parameter and γ measures the substitutability between sectors. Consumption in each sector is derived from a continuum of varieties of good k, that may be imported or not, as in $C_{kj} = \left[\sum_{i\in I} (\beta_{kij}\ C_{kij})^{\frac{\sigma_k-1}{\sigma_k}}\right]^{\frac{\sigma_k}{\sigma_k-1}}$, where $i \in I$ indexes varieties/producing countries (including j, the domestic variety). Crucially, the elasticity of substitution σ_k is specific to each industry, and assumed identical across all varieties, imported or not. β_{kij} lets preferences vary exogenously across varieties, reflecting for instance differences in quality or home bias in consumption. The sectors that verify $\beta_{kij} = 0$ for all $i \neq j$ are effectively non-traded.

This structure of demand is classic in international economics. The key assumption for our purposes is equal substitutability between two varieties, no matter their origin. Introducing the assumption is largely what opened the door to the New Trade literature, pioneered by Krugman (1980), and laid the foundation for the more recent models of trade with heterogeneous firms, starting with Melitz (2003).

The representative maximizing agent chooses her consumption allocation on the basis of Cost, Insurance, Freight prices, labeled in local currencies.[6] Utility maximization implies that demand for variety i in each sector k is given by

$$C_{kij} = \beta_{kij}^{\sigma_k - 1} \left(\frac{P_{kij}}{P_{kj}}\right)^{-\sigma_k} \alpha_k^{\gamma-1} \left(\frac{P_j}{P_{kj}}\right)^{-\gamma} C_j \tag{1}$$

with: P_{kij} the local currency price of variety i of good k, $P_{kj} = \left[\sum_{i \in I} \left(\frac{P_{kij}}{\beta_{kij}}\right)^{1-\sigma_k}\right]^{\frac{1}{1-\sigma_k}}$ and $P_j = \left[\sum_{k \in K} \left(\frac{P_{kj}}{\alpha_{kj}}\right)^{1-\gamma}\right]^{\frac{1}{1-\gamma}}$.

We now ask our model how the estimated response of *aggregate* quantities to changes in *aggregate* international relative prices is affected depending on whether heterogeneity in σ_k is permitted. For this to be a meaningful experiment in a model with multilateral trade at the industry level, we consider disturbances to international relative prices of a specific kind. First, we focus on changes in all relative prices, across all sectors k. This means reallocation of demand across industries is solely driven by the heterogeneous response of sectoral quantities to a uniform price shock. It is relative quantities whose responses may be heterogeneous, which in turn may obscure aggregate estimates.

Second, we focus on uniform shocks to the international price of domestic goods, across all exporters i in $I \neq j$. This assumes away reallocation of demand across source exporting economies, with relative prices changing identically in all markets. We do this for practical reasons, so that the multilateral dimension of the model collapses into a two-country version, and we can interpret our estimate as capturing the substitutability between composite goods in the domestic economy and in the rest of the world.[7] A

[6]Without loss of generality, we could introduce an additional price wedge, reflecting distribution costs that presumably affect both domestic and foreign varieties. This would merely add some notation, but no further insight. In the empirics, the price of each variety is measured Free on Board, i.e. net of both retail and transportation costs.

[7]The second assumption is made for convenience. The intuition remains the same if we focus on a change in (all) relative prices between the domestic economy and a specific exporter i. The data needed

natural candidate is a domestic shock to relative production costs, driven for instance by a productivity disturbance. It will change the international price of domestic goods by an identical amount across all sectors k and exporters $i \neq j$. As a short-hand, we label the shock a "domestic wage" disturbance w_j. An increase in w_j represents a real appreciation driven by a positive shift in relative domestic costs and ultimately prices across all sectors.

Consider the definition of an aggregate elasticity of substitution σ between bundles of domestic and foreign goods. By definition,

$$\sigma = 1 + \frac{\partial \ln \sum_k \sum_{i \neq j} P_{kij} C_{kij} - \partial \ln \sum_k P_{kjj} C_{kjj}}{\partial \ln w_j}$$

The elasticity of substitution captures the relative response of demand for domestic or foreign bundles of goods. Demand is expressed in nominal terms because virtually all trade data are expressed in value, especially at a disaggregated level. Since the driving force to the shift in relative prices is aggregate, the difference between the elasticity of substitution arising from volume or value data is simply 1.

Using equation (1) and its counterpart for the domestic variety, simple algebra implies

$$\begin{aligned} \sigma - 1 \quad = \quad & \sum_k \sum_{i \neq j} n_{kij} (1 - \sigma_k) \frac{\partial \ln P_{kij}}{\partial \ln w_j} - \sum_k n_{kjj} (1 - \sigma_k) \frac{\partial \ln P_{kjj}}{\partial \ln w_j} \\ & + \sum_k \left(n_{kj} - n_{kjj} \right) (\sigma_k - \gamma) \frac{\partial \ln P_{kj}}{\partial \ln w_j} \end{aligned} \tag{2}$$

with $n_{kij} \equiv \frac{P_{kij}\ C_{kij}}{\sum_{k \in K} \sum_{i \neq j} P_{kij}\ C_{kij}}$, $n_{kj} \equiv \frac{P_{kjj}\ C_{kjj}}{\sum_{k \in K} P_{kjj}\ C_{kjj}}$ and $n_{kj} = \sum_{i \neq j} n_{kij}$.

The aggregate elasticity of substitution has two components. The first two terms correspond to appropriately weighted averages of the responses of country-specific prices to the cost shock, at the variety level i or j. The third term captures the responses of the sector price index to the cost shock.

2.2.2 Representative Firm

Since we are interested in long run estimates, in each sector k the representative domestic producer modifies her price according to the change in costs, $\frac{\partial \ln P_{kjj}}{\partial \ln w_j} = 1$. The representative foreign producer, in contrast, does not respond to changes in domestic costs, and $\frac{\partial \ln P_{kij}}{\partial \ln w_j} = 0$.[8] As a result, the response of the sector price index is given by

to perform aggregation are just slightly different, and identification becomes more complicated since it relies on the cross-section of all exporters to a given destination.

[8]The assumption of CES preferences is crucial here. With linear quadratic preferences, for instance, or more sophisticated market structure, price responses embed both demand elasticities and changes in

$\frac{\partial \ln P_{kj}}{\partial \ln w_j} = \frac{\partial \ln P_{kj}}{\partial \ln P_{kjj}} \frac{\partial \ln P_{kjj}}{\partial \ln w_j} = \frac{P_{kjj}\ C_{kjj}}{P_{kj}\ C_{kj}} = 1 - w^M_{kj}$ where $w^M_{kj} = \frac{\sum_{i \neq j} P_{kij}\ C_{kij}}{P_{kj}\ C_{kj}}$ is the share of imports in total expenditures on good k. Equation (2) simplifies into

$$\sigma_{NoFirm} = \sum_k n_{kjj}\ \sigma_k + \sum_k \left(n_{kj} - n_{kjj}\right)\left(1 - w^M_{kj}\right)\left(\sigma_k - \gamma\right) \tag{3}$$

The aggregate elasticity σ_{NoFirm} contains two terms. First, a weighted average of industry-specific elasticities, with weights corresponding to the importance of sector k in overall domestic expenditures. The second term reflects the response of the industry specific price index P_k. Since by assumption the relative price of good k changes identically across all source economies $i \neq j$, the composition of the ideal price index in sector k changes significantly in response to the shock considered. We label σ_{NoFirm} a total elasticity, i.e. one that takes the response of price indices into account. In contrast, a partial elasticity assumes aggregate price responses away.[9] We note the second term in equation (3) is likely small. By definition of n_{kjj}, $\sum_k \left(n_{kj} - n_{kjj}\right)\left(1 - w^M_{kj}\right)\left(\sigma_k - \gamma\right) < \sum_k n_{kj}\sigma_k \left(1 - w^M_{kj}\right)$. The difference between partial and total elasticities is bounded above. Relative to partial elasticity, the upper bound contains an extra multiplicative term smaller than 1, since $w^M_{kj} < 1$ for all k. Partial and total elasticites are therefore likely to differ by small amounts. In fact, if sector allocations of expenditures are similar for domestic and foreign goods, $n_{kj} \simeq n_{kjj}$, and partial and total elasticities are virtually identical.

Focus for now on the partial elasticity and suppose $\frac{\partial \ln P_{kj}}{\partial \ln w_j} = 0$. The aggregate substitutability between foreign and domestic varieties is given by a weighted average of each industry's corresponding preference parameter. Strictly speaking, this weighted average is the direct equivalent of the aggregate elasticity of substitution obtained from an aggregation of *estimates*, rather than an aggregation of the *data*. It captures the direct effect of a shock to w_j on P_{kjj}, and the resulting immediate change in C_{kjj} visible from equation (1). In macroeconomic data, traded quantities are summed up across sectors and exporting countries. The econometrician asks how these synthetic aggregates respond to changes in international prices. In other words, she estimates one single value for σ, which may well differ from $\sum_k n_{kjj}\ \sigma_k$. If it does, a heterogeneity bias prevails in the data, and disaggregated data are necessary to calculate an aggregate elasticity accounting

markups.

[9]The second term in equation (3) exists because of our focus on an aggregate shock in relative prices. If instead the shock considered were microeconomic in nature and focused on a specific exporter i -a change in tariff- then under standard atomistic assumptions and large I, the second term in equation (3) would disappear, and we would be left with a partial elasticity.

for heterogeneity.[10] In fact, equation (3) tells us the direction and magnitude of a heterogeneity bias depend on the cross-sector correlation between expenditures shares n_{kjj} and goods' substitutability σ_k. If substitutable varieties tend to form a large share of domestic expenditures, the unconstrained elasticity σ_{NoFirm} will take larger values than its constrained counterpart, $\bar{\sigma}_{NoFirm}$.

Sectoral price indices also respond to macroeconomic shocks. Consider a positive shock to domestic costs, driving up domestic prices and therefore P_{kj}, to an extent that increases with the share of domestic varieties, $1 - w_{kj}^M$. The response of the price index affects the relative prices of both domestic and imported varieties. It acts to dilute the initial upwards change in $\frac{P_{kjj}}{P_{kj}}$, which dampens measured total elasticity. But it also drives a negative response in $\frac{P_{kij}}{P_{kj}}$, with opposite end effects on measured total elasticity.

Estimates of the elasticity of substitution obtained from aggregate data provide potentially biased values of a preference parameter that accounts for heterogeneity in σ_k. The bias arises because differences in sector-specific substitutability are assumed away. This is true of partial and total elasticity, and we later report results for both concepts. In what follows, we obtain structural estimates of σ_k (and its counterpart constrained to homogeneity), and infer the corresponding, theory-implied (partial or total) values for σ_{NoFirm} and $\bar{\sigma}_{NoFirm}$. We will show that the implied values for $\bar{\sigma}_{NoFirm}$ are in line with classic results from the macroeconomic empirical literature, whereas the implied values for σ_{NoFirm} are not.

2.2.3 Firm Heterogeneity

In a conventional model of international trade assuming CES preferences and a representative firm, the heterogeneity bias affects the elasticity of substitution σ. How does such interpretation depend on the structure of the supply side of the economy? Recent papers by Eaton and Kortum (2002), Chaney (2008) or Arkolakis, Costinot and Rodriguez-Clare (2009) have argued traded quantities disconnect from preferences when firm heterogeneity is introduced. In Section A of the Technical Appendix, we present a slightly generalized version of the model in Chaney (2008).

Firms are heterogeneous, so that the exporter's price $P_{kij} = \left[\int p_{kij}(\phi)^{1-\rho_k} d\phi\right]^{\frac{1}{1-\rho_k}}$ is an index. P_{kij} now reflects the continuum of firms active in sector k of country i, that

[10] In Appendix A, we show that in a CES world with heterogeneous σ_k, estimating the conventional specification on aggregated data effectively imposes an additional constraint, namely $\sigma_k = \sigma = \gamma$.

export to country j. ϕ indexes firm productivity. The range of active firms can differ endogenously across countries and sectors. It is allowed to vary endogenously in response to changes in international relative prices. This response may itself differ across sectors and exporters because of differences in transport costs, assumptions on firms distribution, or heterogeneity in substitutability parameters. ρ_k denotes the elasticity of substitution between the varieties produced by firms in a given country. It is assumed to be sector-specific but identical across all countries i. We further assume $\rho_k > \sigma_k > \gamma > 1$ for all k. The varieties produced by two firms from the same country are more substitutable than those produced by firms located in two different countries. These, in turn, are more substitutable than goods from different sectors. We assume complementarities away.

We differ from Chaney (2008) in that we allow for the substitutability ρ_k between varieties produced at the firm-level within each country to differ from the elasticity σ_k between countries. We show the expression for the aggregate elasticity of substitution between domestic and foreign goods then continues to depend on σ_k, and indeed continues to suffer from a heterogeneity bias. Under conventional assumptions, the bias becomes even larger. Importantly, we do not mean to suggest no model exists where introducing firm heterogeneity will alter fundamentally the interpretation of the hetereogeneity bias. In the original theory by Chaney (2008) the bias affects only the parameter that governs firm hetereogeneity. The practical implication of a heterogeneity bias in σ pertains first and foremost to the calibration of open economy models with representative firms, where the quantities traded depend only on the substitutability in preferences. We argue the vast majority of *calibrated* models continue to be of that type. The theory detailed in the Technical Appendix illustrates the possibility that the implications of our paper actually generalize to *some* models with firm heterogeneity.

In Section A of the Technical Appendix, we follow the steps described in Chaney (2008). Like Chaney, we assume the mass of exporting firms active in sector k and country i, is proportional to the aggregate wage bill $\omega_i L_i$ so that larger and richer countries have more entrants, and we impose $\gamma = 1$. Using the definition of P_{kj}, we solve equation (2) for the aggregate elasticity of substitution allowing for firm entry decisions in all markets:

$$\sigma_{Firm} - 1 = \sum_k n_{kjj}(\sigma_k - 1)\Delta_k + \sum_k (n_{kj} - n_{kjj})(1 - w_{kj}^M)(\sigma_k - 1)\Delta_k \qquad (4)$$

where $\Delta_k = \frac{\theta_k}{\sigma_k - 1 + \theta_k \frac{\rho_k - \sigma_k}{\rho_k - 1}}$, and θ_k characterizes the Pareto distribution of firm productivity in sector k.

Note first that σ_{Firm} depends on θ_k when $\rho_k = \sigma_k$. This reproduces the result in Chaney (2008), where the substitution between traded and domestic goods depends only on a technology parameter. Second, equation (4) nests the expression for σ_{NoFirm} obtained in equation (3), when $\Delta_k = 1$. Therefore Δ_k represents a corrective term that captures the effect of firm heterogeneity on the elasticity of substitution in the presence of firm heterogeneity. Assuming $\theta_k > \rho_k - 1$, as in Chaney (2008) or Eaton and Kortum (2002), we have $\Delta_k > 1$. In other words, the elasticity of substitution with firm heterogeneity is systematically larger in absolute value than what is implied by the same model with a representative firm, an argument reminiscent of Ruhl (2005).

Our main purpose in this section is to verify that supply responses across all markets do not alter the existence of an heterogeneity bias in estimates of the aggregate elasticity of substitution. Does the discrepancy between the estimate constrained to homogeneity, $\bar{\sigma}_{Firm}$, and its unconstrained counterpart, σ_{Firm}, differ from that between $\bar{\sigma}_{NoFirm}$ and σ_{NoFirm}? We noted in the previous section that the difference between $\bar{\sigma}_{NoFirm}$ and σ_{NoFirm} increases in the cross-sector correlation bewteen n_{kjj} and σ_k. For positive correlations, a model-implied weighted average of σ_k will take larger values than its counterpart imposing homogeneity $\sigma_k = \sigma$ for all k.

Equation (4) suggests the discrepancy will be even larger with firm entry. A given positive value for the correlation between n_{kjj} and σ_k implies larger values for $\sigma_{Firm} - \bar{\sigma}_{Firm}$ than for $\sigma_{NoFirm} - \bar{\sigma}_{NoFirm}$. The bias implied by the theory with firm entry is larger than its counterpart assuming supply responses away. To see this, manipulate equation (4) to obtain

$$\begin{aligned}\sigma_{Firm} - \bar{\sigma}_{Firm} &= \sigma_{NoFirm} - \bar{\sigma}_{NoFirm} + \sum_k n_{kjj} \left[(\sigma_k - 1)(\Delta_k - 1) - (\bar{\sigma} - 1)\left(\bar{\Delta} - 1\right) \right] \\ &\quad + \sum_k (n_{kj} - n_{kjj})(1 - w_{kj}^M) \left[(\sigma_k - 1)(\Delta_k - 1) - (\bar{\sigma} - 1)\left(\bar{\Delta} - 1\right) \right] \quad (5)\end{aligned}$$

where $\bar{\Delta} = \frac{\bar{\theta}}{\bar{\sigma} - 1 + \bar{\theta} \frac{\bar{\rho} - \bar{\sigma}}{\bar{\rho} - 1}}$ and $\bar{\sigma}$, $\bar{\rho}$ and $\bar{\theta}$ are homogeneous parameters, constrained to be identical across sectors. The first summation increases in the correlation between n_{kjj} and σ_k, which also acts to augment the magnitude of the bias.[11] The second summation

[11] To see this, note that $\Delta_k - 1 = \frac{\theta_k - (\rho_k - 1)}{\rho_k - 1 + \theta_k \frac{\rho_k - \sigma_k}{\sigma_k - 1}}$, which increases in σ_k.

term can take either sign, depending on values for n_{kj} and n_{kjj}. But by analogy with the previous section, the term must be smaller than the first summation, and cannot act to mitigate its impact, even if it takes negative values.

We later introduce an empirical approach to estimate σ_k (and its constrained counterpart $\bar{\sigma}$). For lack of cross-country firm-level data, we cannot obtain any estimates of ρ_k. We are therefore only able to calculate a theory-implied heterogeneity bias without firm entry, $\sigma_{NoFirm} - \bar{\sigma}_{NoFirm}$. In our data, n_{kjj} and σ_k are positively correlated so that the bias will take positive values. Under these conditions, and within the confines of our version of Chaney's model, we just showed that $\sigma_{Firm} - \bar{\sigma}_{Firm}$ is greater than $\sigma_{NoFirm} - \bar{\sigma}_{NoFirm}$. This suggests the heterogeneity bias we document - on the basis of our estimates of σ_k - is a lower bound.

2.2.4 The Price Elasticity of Imports

In most of the literature, the elasticity of substitution is inferred from the price elasticity of imports, at any level of aggregation. Since we want to validate our assumptions by comparing our results to conventional estimates (at the micro or macroeconomic level), it is important that we verify the bias we discuss continues to prevail in estimates of the price elasticity of imports. Fortunately, the exercise follows arguments that are analogous to what was just discussed. The conventional approach to identifying the price elasticity of imports consists in estimating the response of imports to changes in international relative prices (a negative number).[12] In the model, this is equivalent to

$$\eta = \frac{\partial \ln \left[\sum_k \sum_{i \neq j} \; P_{kij} C_{kij} \right]}{-\partial \ln w_j}$$

Demand continues to be given by equation (1). In Section B of the Technical Appendix, we derive the following expression for η_{NoFirm}:

$$\eta_{NoFirm} = \gamma - \sigma_{NoFirm} + \sum_k n_{kjj}(\sigma_k - \gamma) w_{kj}^M - \gamma \sum_k w_{kj} \left(1 - w_{kj}^M\right) \tag{6}$$

where $w_{kj} = P_{kj}C_{kj}/P_jC_j$ is the share of sector k in aggregate consumption. The elasticity contains three terms: a direct linear function of the elasticity of substitution σ_{NoFirm}, and two summation terms, likely to be smaller in magnitude than σ_{NoFirm}, since $n_{kjj}, w_{kj}^M < 1$.

We abstract from both summations terms, and label the remaining term a partial elasticity of imports. It is given by $\gamma - \sigma_{NoFirm}$, or by $\gamma - \bar{\sigma}_{NoFirm}$ when homogeneity

[12] We focus on estimates of the price elasticity of imports abstracting from firm dynamics.

is imposed on σ_k. Constraining all elasticities to be the same, the partial price elasticity of imports and the elasticity of substitution are linearly related, a conventional result. With positive cross-sector correlation between n_{kjj} and σ_k, estimates of the (partial) price elasticity of imports that introduce heterogeneity will be larger in absolute value. In the model, $\sigma_{NoFirm} > \bar{\sigma}_{NoFirm}$ implies that $\eta_{NoFirm} < \bar{\eta}_{NoFirm}$. This can explain why macroeconomic estimates of η are smaller in absolute value than those arising from disaggregated data.

For a given γ, equation (6) implies a value for η_{NoFirm} given structural estimates of σ_k, and a value for $\bar{\eta}_{NoFirm}$ given a structural estimate for $\bar{\sigma}$. Whether a heterogeneity bias can account for the discrepancy between estimates of η obtained from micro and macroeconomic data clearly does not depend on a choice for γ. In fact, we later calibrate γ at its lowest value of 1, which corresponds to an upper bound (in absolute value) for the price elasticity of imports. Our hope is the implied value for $\bar{\eta}_{NoFirm}$ continues to be in line with macroeconomic estimates, i.e. close to zero.

3 Identification

We adapt the methodology in Feenstra (1994) to our purposes. Identification is structural, but requires a CES demand system, with constant markups. Our results are predicated on these assumptions. We first discuss the econometrics involved in estimating σ_k for all sectors k in the US economy. We emphasize how we accommodate common effects across all sectors and measurement error. We then turn to the estimation of $\bar{\sigma}$, a measure of elasticity constrained to be identical across sectors. We close with a description of our data.

3.1 Microeconomic Estimates

We identify the substitutability between domestic and foreign varieties using the observed cross-section of traded quantities and prices across exporters to one destination. This is afforded by the crucial assumption of an Armington aggregator between varieties of each good, irrespective of their origin. The assumption is what makes it possible to use Feenstra's (1994) methodology in the present context, even though we do not observe prices or quantities of domestically produced varieties. We now describe our implementation of his methodology, but keep the development concise and focused on the modifications we introduce.

Demand is given in equation (1), which after rearranging writes:

$$C_{kit} = \left(\frac{P_{kit}}{P_{kt}}\right)^{1-\sigma_k} \frac{\beta_{kit}^{\sigma_k-1} P_{kt} C_{kt}}{P_{kit}}$$

where t is a time index.[13] Feenstra (1994) or Broda and Weinstein (2006) impose a simple supply structure, with prices fixed in local currency and inclusive of trade costs

$$P_{kit} = \tau_{kit} \exp(\upsilon_{kit}) C_{kit}^{\omega_k}$$

where υ_{kit} denotes a technological shock that can take different values across sectors and exporters, τ_{kit} is a trading cost and ω_k is the inverse of the price elasticity of supply in sector k.[14] The potential aggregate effects of the nominal exchange rate are soaked up by the shock υ_{kit}. It will be important to implement appropriate econometric tools to control for any potential common effects in our estimated system.[15]

In practice, the approach uses expenditure shares to alleviate measurement error in unit values, following Kemp (1962). We define $s_{kit} = \frac{P_{kit}C_{kit}}{P_{kt}C_{kt}}$ and rewrite demand as $s_{kit} = \left(\frac{P_{kit}}{P_{kt}}\right)^{1-\sigma_k} \beta_{kit}^{\sigma_k-1}$.

We do not observe domestically produced consumption. In addition, prices are measured Free on Board. We introduce tilded variables to denote the observed counterparts to theory-implied prices and quantities. We observe $\tilde{P}_{kit} \equiv P_{kit}/\tau_{kit}$. The empirical market shares are therefore given by $\tilde{s}_{kit} \equiv \frac{\tilde{P}_{kit}C_{kit}}{\sum_{i \neq d} \tilde{P}_{kit}C_{kit}} = \frac{s_{kit}}{\tau_{kit}} \left(1 + \frac{P_{kdt}C_{kdt}}{\sum_{i \neq d} \tilde{P}_{kit}C_{kit}}\right) \equiv \frac{s_{kit}}{\tau_{kit}} \mu_{kt}$.

Taking logarithms, it is straightforward to rewrite demand as

$$\Delta \ln \tilde{s}_{kit} = (1 - \sigma_k) \Delta \ln \tilde{P}_{kit} + \Phi_{kt} + \varepsilon_{kit} \tag{7}$$

with $\Phi_{kt} \equiv (\sigma_k - 1) \Delta \ln P_{kt} + \Delta \ln \mu_{kt}$, a time-varying intercept common across all varieties, and $\varepsilon_{kit} \equiv (\sigma_k - 1) \Delta \ln \beta_{kit} - \sigma_k \Delta \ln \tau_{kit}$ an error term that captures random trade cost and taste shocks, via changes in τ_{kit} and β_{kit}. Feenstra (1994) shows this implies the demand system is robust to quality changes in variety i of good k - or indeed to time-varying number of firms producing good k in country i. The estimation is robust to the presence of an extensive margin within exporting economies.

After rearranging, substituting in log-linearized supply yields

$$\Delta \ln \tilde{P}_{kit} = \Psi_{kt} + \frac{\omega_k}{1 + \omega_k \sigma_k} \varepsilon_{kit} + \delta_{kit} \tag{8}$$

[13] Throughout this section we omit the index j, as our end results focus on a single importing country, the US.

[14] We follow Feenstra (1994) and assume all exporters have the same supply elasticity. Whether prices are inclusive of transport costs or not is innocuous for the end estimates, as τ_{kit} enters the residuals of the estimated equation.

[15] υ_{kit} will also absorb any heterogeneity in the extent of the exchange rate pass-through.

with $\Psi_{kt} \equiv \frac{\omega_k}{1+\omega_k\sigma_k}\left[\Phi_{kt} + \Delta \ln \sum_i (\tilde{P}_{kit} C_{kit})\right]$ a time-varying factor common across varieties, which subsumes sector specific prices and quantities. $\delta_{kit} \equiv \frac{1}{1+\omega_k\sigma_k}\Delta \upsilon_{kit}$ is an error term encapsulating movements in the exchange rate or aggregate technological developments in country i and sector k.

Under standard assumptions on taste shocks β_{kit} and technology shocks υ_{kit}, it is possible to identify the system formed by equations (7) and (8). Identification rests on the cross-section of exporters i to the domestic economy, and is achieved in relative terms with respect to a reference country r.[16] We follow Feenstra (1994) and summarize the information contained in the system with the following estimable regression

$$Y_{kit} = \theta_{1k} X_{1kit} + \theta_{2k} X_{2kit} + u_{kit} \tag{9}$$

where $Y_{kit} = (\Delta \ln \tilde{P}_{kit} - \Delta \ln \tilde{P}_{krt})^2$, $X_{1kit} = (\Delta \ln \tilde{s}_{kit} - \Delta \ln \tilde{s}_{krt})^2$, $X_{2kit} = (\Delta \ln \tilde{s}_{kit} - \Delta \ln \tilde{s}_{krt})(\Delta \ln \tilde{P}_{kit} - \Delta \ln \tilde{P}_{krt})$ and $u_{kit} = (\varepsilon_{kit} - \varepsilon_{krt})(\delta_{kit} - \delta_{krt})\frac{(\sigma_k-1)(1+\omega_k)}{1+\omega_k\sigma_k}$. Estimates of equation (9) map directly with the parameters of interest, since

$$\theta_{1k} = \frac{\omega_k}{(\sigma_k - 1)(1+\omega_k)}, \qquad \theta_{2k} = \frac{\omega_k\sigma_k - 2\omega_k - 1}{(\sigma_k - 1)(1+\omega_k)}$$

Equation (9) still suffers from an endogeneity issue. We follow Feenstra (1994), instrument the regressors with country-sector specific fixed effects, and correct the estimation for heteroskedasticity across exporters i. As in Feenstra, identification is therefore based on the cross-sectional dimension of equation (9), and is only valid under CES preferences of the type assumed here. We include an intercept to account for the measurement error arising from using unit values to approximate prices. Given the origin of potential measurement error, we let it prevail at the most granular level afforded by our data.[17]

The system summarized by equation (9) can accommodate developments that are specific to each sector k. But in macroeconomic applications where the universe of economic activities that form Gross Domestic Product is considered, it is important to allow for more general, aggregate influences. Aggregate technology shocks for instance, or movements in the nominal exchange rate, presumably affect prices and quantities jointly in all sectors. If it were a shock in the exporting economy, that would correspond to a common component of υ_{kit} across all k. We allow for such correlated effects in as general and parsimonious a manner as possible. We implement a correction suggested by Pesaran (2006) to purge all "Common Correlated Effects" (CCE) from sector level data, and estimate

[16] In the empirics, we choose a reference country that is present in the US market during the whole observed period.

[17] See Feenstra (1994).

$$Y_{kit} = \theta_0 + \theta_{1k}\hat{X}_{1ki} + \theta_{2k}\hat{X}_{2ki} + \theta_{3k}X_{1it} + \theta_{4k}X_{2it} + u_{kit} \quad (10)$$

where the intercept allows for HS6-specific measurement error, hatted variables are the instrumented versions of X_{1kit} and X_{2kit}, and X_{1it} and X_{2it} control for the time-varying component of Y_{kit} that is common across all sectors. In particular, following Pesaran (2006), X_{1it} and X_{2it} are the cross-sector arithmetic averages of X_{1kit} and X_{2kit}.

Armed with consistent (and sector-specific) estimates of θ_{1k} and θ_{2k}, it is straightforward to infer elasticities. In particular, the model implies

$$\begin{aligned}\hat{\sigma}_k &= 1 + \frac{\hat{\theta}_{2k} + \Delta_k}{2\hat{\theta}_{1k}} \text{ if } \hat{\theta}_{1k} > 0 \text{ and } \hat{\theta}_{1k} + \hat{\theta}_{2k} < 1 \\ \hat{\sigma}_k &= 1 + \frac{\hat{\theta}_{2k} - \Delta_k}{2\hat{\theta}_{1k}} \text{ if } \hat{\theta}_{1k} < 0 \text{ and } \hat{\theta}_{1k} + \hat{\theta}_{2k} > 1\end{aligned}$$

with $\Delta_k = \sqrt{\hat{\theta}_{2k}^2 + 4\hat{\theta}_{1k}}$. Section D of the Technical Appendix details how these are also used to infer standard deviations around these point estimates.

As is apparent, there are combinations of estimates in equation (10) that do not correspond to any theoretically consistent estimates of $\hat{\sigma}_k$. We follow Broda and Weinstein (2006), and use a search algorithm that minimizes the sum of squared residuals in equation (10) over the intervals of admissible values of the supply and demand elasticities.[18]

3.2 Homogeneous Estimates

We invoke equation (3) to aggregate adequately the estimates of $\hat{\sigma}_k$ just obtained. Our purpose is to compare these results to what is obtained when sectoral elasticities are constrained to be homogeneous, as they would in conventional regression analysis based on macroeconomic data. To do so, we impose $\sigma_k = \sigma$ and modify equation (10) into

$$Y_{kit} = \theta_0 + \theta_1\hat{X}_{1ki} + \theta_2\hat{X}_{2ki} + \theta_3X_{1it} + \theta_4X_{2it} + u_{kit} \quad (11)$$

We maintain the assumption of a HS6-specific intercept, to continue to accommodate the possibility that measurement error varies at the disaggregated level. Not doing so will conflate two potential sources of bias, and the one we are pursuing relates only to the estimates of θ_1 and θ_2. These are now constrained to equality across all sectors k. Identification continues to rest on the cross-section of exporters i, but equation (11) is now estimated on the pooled dataset formed by observations on all sectors. It is noteworthy

[18] We use this approach whenever direct estimates of θ_{1k} and θ_{2k} cannot be used to infer $\hat{\sigma}_k$. Whenever CCE are included, we hold constant the estimates of θ_{3k} and θ_{4k} obtained from the direct instrumental variable regression, and search the combination of values for σ_k and ω_k that minimizes the sum of squared residuals in equation (10). The corresponding standard errors are obtained via bootstrapping of the procedure using 1,000 repetitions.

that identification in equations (10) and (11) rests in practice on the *same* dimension of the *same* dataset. It is therefore difficult to ascribe the discrepancy we find to a difference in the extent of an extensive margin. Whether our data (and procedure) capture or not firm dynamics, they do so equally in both estimations.

If sectoral heterogeneity in equation (10) is a good representation of reality, equation (11) suffers from a conventional heterogeneity bias. The bias arises as the coefficient estimates are constrained to homogeneity, so that sectoral differences in $\theta's$ are forced into the residual u_{kit}. As in Pesaran and Smith (1995), ignored heterogeneity creates a potential bias in the estimates of θ_l, $l = 1, 2, 3, 4$. As a result, empirical estimates of the constrained elasticity of substitution $\hat{\sigma}$ may well be away from a simple average of its sectoral counterparts.

We continue to allow for the possibility that aggregate shocks in any country i should affect all sectors simultaneously, and include adequately modified CCE terms. The instrumentation and correction for heteroskedasticity are also modified accordingly. In particular, country-specific effects are used as instruments across the whole panel of sectors. Armed with estimates of θ_1 and θ_2, it is easy to obtain a value for the constrained elasticity of substitution $\hat{\sigma}$. Our model then implies conventional macroeconomic estimates of the elasticity of substitution are given by

$$\bar{\sigma}_{NoFirm} = \hat{\sigma} + (\hat{\sigma} - \gamma) \sum_k (n_{kj} - n_{kjj}) \left(1 - w_{kj}^M\right)$$

with a standard error given by a first-order approximation, as detailed in the Technical Appendix. $\bar{\sigma}_{NoFirm}$ is the constrained total elasticity, and $\hat{\sigma}$ is a constrained partial elasticity.

3.3 Data

The Armington assumption requires that imports from different countries be imperfectly substitutable varieties. The hypothesis is increasingly palatable as the granularity of the data augments. We choose to use disaggregated, multilateral trade data from the Base Analytique du Commerce International (BACI), released by the Centre d'Etudes Prospectives et d'Informations Internationales (CEPII), and available at the 6-digit level of the harmonized system (HS6). The data cover around 5,000 products over the 1996-2004 period for a large cross-section of countries. The database describes bilateral trade at the sectoral level, building on the United Nations ComTrade database with some added effort put in the harmonization of trade flows on the basis of both import and export declarations. The improvement acts to limit measurement error.

Unlike Feenstra (1994) or Broda and Weinstein (2006), we do not estimate elasticities at the most disaggregated level in the main body of the text, but rather partition our data into 56 ISIC (Revision 3) industries where we implement our methodology. The constrained estimation given by equation (11) is effectively estimated on a panel of HS6 sectors, and our partition of HS6 sectors into 56 ISIC categories is entirely innocuous here. In the constrained estimation, homogeneity is imposed across all HS6 sectors, and whether or not they belong to the same ISIC sector is irrelevant. The partition becomes important for the unconstrained estimation, and is performed for lack of detailed information on n_{kj}, n_{kjj}, w_{kj} or w_{kj}^M at such high level of granularity. It corresponds to the assumption that all HS6 goods are equally substitutable within an ISIC category, but not between. This does presumably assume some heterogeneity away, and possibly creates a bias as a result. We conjecture that heterogeneity between ISIC industries is more sizable, and thus creates more of a bias. We do however perform some robustness in section 4.4, using all HS6 goods in the unconstrained estimation. But to do so we have to maintain some rather stringent assumptions on the values for n_{kj}, n_{kjj}, w_{kj} or w_{kj}^M.

The approach adapted from Feenstra (1994) requires relatively little information on traded flows. To estimate equation (10) we only need measures of $\tilde{P}_{kit}$ and the expenditure shares $\tilde{s}_{kit}$. As is conventional, we use unit values to approximate bilateral prices, and divide values of bilateral trade flows by their volume. In BACI, values are denominated in USD and are Free On Board.[19] Quantities are in tons. The empirical model described in section 3.1 is not sensitive to the currency denomination of trade data, nor to the treatment of trade costs, as both are passed into the residuals. Expenditure shares are measured as $\tilde{s}_{kit} = \frac{\tilde{P}_{kit}C_{kit}}{\sum_{i\neq d}\tilde{P}_{kit}C_{kit}}$.

We subject our data to sampling with a view to limiting the role of extreme outliers. These are notoriously frequent in approaches making use of unit values to approximate prices. For instance, tonnage is not always appropriate to capture the traded volumes of all HS6 goods, which can instill artificial (massive) volatility in the resulting time series on prices. In each sector, we exclude annual variations in prices and market shares that exceed five times the median value. In addition, we impose a minimum of 20 exporters

[19] In general, trade data are collected by national customs offices in the currency of the declaring country. These data are then converted in US dollars by the United Nations, using the current nominal exchange rate.

for each HS6 good over the whole observed time period. The cross-section of exporters is what ultimately achieves identification. Measurement error may prevail in estimates of $\hat{\sigma}_k$ that are based on too few exporters, which would translate into biased values for (unconstrained) aggregate elasticities. We require that at least 20 exporters be present to alleviate this concern. Our data ultimately represent 77 percent of the total value of US imports, across 56 ISIC sectors.

In the model, n_{kj} and n_{kjj} depend directly on the import share w_{kj}^M and the expenditure share w_{kj}. In particular, we have $n_{kj} = \frac{w_{kj}\ w_{kj}^M}{\sum_k w_{kj}\ w_{kj}^M}$ and $n_{kjj} = \frac{w_{kj}\ \left(1-w_{kj}^M\right)}{\sum_k w_{kj}\ \left(1-w_{kj}^M\right)}$. Calibration is therefore only needed for w_{kj} and w_{kj}^M. In the main body of the text, we consider the following data sources. The expenditure shares w_{kj} are obtained from the OECD STAN dataset, as the 1997 ratio of sectoral absorption (value added and imports net of exports) relative to the aggregate across sectors. The import shares w_{kj}^M are computed from the US input/output (IO) tables, available in the ISIC (Revison 3) nomenclature, as the 1997 ratio of imports over domestic gross output. Values for n_{kj} and n_{kjj} are calculated accordingly.[20]

In section 4.4, we verify our results do not depend on this specific choice of data sources. We discuss four alternatives. First, we compute w_{kj}^M directly from the BACI dataset used in our main estimation, rather than the IO tables, normalized by a measure of domestic output taken from the OECD STAN data. But we continue to compute both n_{kj} and n_{kjj} on the basis of their model-implied values. Second, the IO tables provide enough information to compute n_{kj} directly, rather than on the basis of a model-implied formula. In our second variant, we do so, and use IO tables to calibrate both w_{kj}^M and n_{kj}. But n_{kjj} continues to be computed according to the model, since we do not have information on domestic production. Our third variant combines both insights. We infer w_{kj}^M from the BACI and STAN dataset, but now also use BACI to calibrate n_{kj}. Finally, we return to our original data sources in our fourth variant, get w_{kj} from STAN and w_{kj}^M from the IO tables. But now, we compute sectoral absorption on the basis of gross output rather than value added.

[20] w_{kj} and w_{kj}^M do not sum to one because of non-traded sectors. Since n_{kj} and n_{kjj} both sum to unity by definition, we normalize each definition so that it is the case.

4 Results and Relevance

We first review the microeconomic estimates, obtained across 56 ISIC sectors, and relate them with existing evidence. We then aggregate the estimates, preserving heterogeneity. We compare the results with our estimation imposing homogeneity - performed using the same data but with the added constraint that $\sigma_k = \sigma$. We discuss the discrepancy, and in particular whether its magnitude is significant economically. We argue the heterogeneous estimates we uncover change dramatically the quantitative and qualitative predictions of a vast range of international macroeconomic models. We close with some robustness.

4.1 Microeconomic Results

Figure 1 reports sectoral estimates of $\hat{\sigma}_k$ for 56 ISIC sectors. On average, $\hat{\sigma}_k$ is equal to 6.7, with values ranging from 3.1 to 28 and a standard deviation of 4.9. The median value is 5.1, reflecting a skewed distribution of elasticities: only 5 out of 56 estimates are above 10. How do our results compare on average with existing studies of the substitutability between foreign and domestic varieties, at similar aggregation levels? If anything, a median value of 5.1 lies at the low end of the range of estimates obtained in the empirical trade literature. Romalis (2007) finds elasticities of substitution between 4 and 13 at the HS6 level. Head and Ries (2001) find values between 7.9 and 11.4 at the 3-digit SIC level. Hanson (2005) finds estimates between 4.9 and 7.6 using data at the US county level. A common denominator across these studies is their focus on disaggregated, microeconomic information on traded quantities and/or tariffs.

Figure 1: Microeconomic estimates of the elasticity of substitution

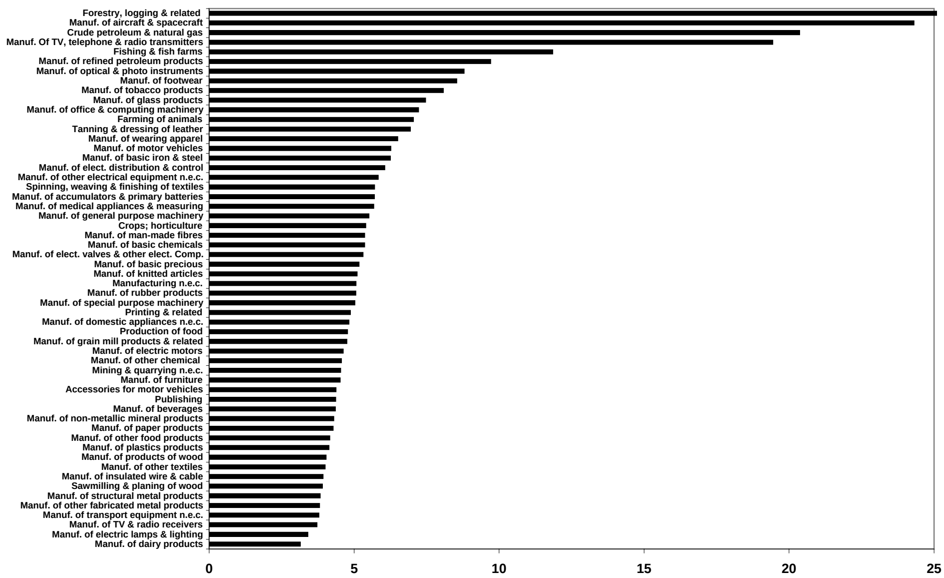

How does the magnitude of our individual sectoral estimates compare with the literature? In theory, the parameters we estimate are comparable with the values obtained in the conventional approach regressing imported quantities on relative prices. The relative price of imports is typically measured with respect to domestically produced varieties, as in Houthakker and Magee (1969) or Kreinin (1967). At a sectoral level, the resulting estimates of the price elasticity of imports map with σ_k according to $\eta_k = 1 - \sigma_k$.

Import elasticities were the object of a vast literature over the 1960s and 1970s. In most instances, the results focused on short-run elasticities, typically for reasons of data availability especially at a disaggregated level. In contrast our estimates correspond to long run elasticities, since the identification is in cross-section. The data available then were also coarser, and focused on just a few sectors. Still, in what follows we strive to ensure our disaggregated estimates are consistent with the existing estimates of import elasticities.

Houthakker and Magee (1969) report in their Table 6 a long run price elasticity in manufactures estimated at -4.05. This is virtually identical to the median value we obtain across our 56 manufacturing sectors, equal to -4.1 $(= 1 - 5.1)$. Kreinin (1967) documents similar estimates, with an elasticity for manufactures equal to -4.71. It is remarkable that such different data sources, coverages and methodologies should yield strikingly similar median estimates. The similarity in estimates across decades suggests there is nothing special to our data, or our approach, relative to existing empirical work. At sector level, we seem to identify the object of interest to many international economists.

The data in Houthakker and Magee (1969) and Kreinin (1967) are much coarser than ours, but they also discuss the relative magnitudes of elasticity estimates across the categories they observe. Manufactures have the higher estimates, followed by semi-manufactures and crude foods and materials. Similar relative rankings come out of the survey in Goldstein and Khan (1985). They summarize their Table 4.4 commenting that "the price elasticity of demand for manufactures is significantly larger than that for non-manufactures. Within non-manufactures, price elasticities for raw materials appear to be larger than those for food and beverages" (pages 1084-1085). A precise mapping is difficult given the differences in granularity, but the ranking is roughly prevalent in our results as well. There are exceptions, but our highest estimates concern finished manufactures, such as aircrafts, TVs, telephones, photo instruments, footwear, motor vehicles or office machinery. At the other end of the spectrum, we find relatively low elasticities for dairy, wood, food, beverages and semi-manufactures like wires or metal products.

Mapping our most disaggregated, individual sector estimates with the literature quickly becomes difficult, once again because of data availability as of 20 or 30 years ago. In fact, not many papers have attempted to estimate sector-specific price elasticities of imports, say at the two or three digit level of aggregation. We were able to identify two exceptions. Stone (1979) presents US estimates at the two digit level. On the whole, his estimates are lower than ours, but that can simply reflect his focus on short run elasticities. A few examples may nevertheless help illustrate the relative similarities in our results. For "Inorganic Chemicals", Stone estimates an import price elasticity of -3.40, as against -3.60 for "Other Chemicals" in Figure 1. He finds -2.32 and -3.71 in "Plastic Materials and Articles" and "Dyeing, Tanning and Coloring Agents", as against -3.1 and -5.9 in "Manufactures of Plastic Products" and "Tanning and Dressing of Leather" in Figure 1. Keeping in mind ours are estimates of long run elasticitites, these values lie in similar ballparks.

Shiells (1991) estimates long run elasticities at the three digit SITC level, but only for 12 US sectors. Once again, an accurate mapping is impossible in most cases. Interestingly however, his estimate in "Newsprint" is -3.6, indistinguishable from our value of -3.4 for "Publishing". He also finds -3.5 in "Steel Plate and Sheet", relatively close to the estimate of -5.2 we find for "Manufacturing of Basic Iron and Steel", even though sector definitions are different. The discrepancies become even less substantial when taking into account Shiells' relatively large standard errors.

This comparison exercise is not meant to suggest we reproduce exactly sector-specific results that were obtained several decades ago in totally different data using drastically different methodologies. Rather, we seek to ascertain the identification strategy we follow is not fundamentally falsified. In particular, the Armington assumption is what exonerates us from having to observe any characteristics of domestically produced goods. With the Armington aggregator, the observed prices and quantities of imports originating from a cross-section of countries are sufficient to identify the elasticity of interest, between domestic and foreign varieties.

From this point of view, it is reassuring that our mean and median estimates should be strikingly close to seminal, fundamental contributions to the literature on imports price elasticities. Virtually all the papers there do make use of domestic prices in their estimations: import prices are evaluated relative to their domestically produced counterpart. That we should find similar results without any information on domestic prices brings support to the Armington assumption. The few punctual comparisons we report at the level of individual sectors do, as well.

There is an obvious comparison absent from our analysis so far. We have implemented a variant of the methodology introduced by Feenstra (1994), just as Broda and Weinstein (2006) have. Our objectives are fundamentally different, as are some of our identifying restrictions and some of the corrections we introduce. Still, Broda and Weinstein estimate the universe of substitution elasticities in disaggregated US data, just as we do. Given the similarities in methodologies, it is not surprising that our estimates should be similar, and they are. Their mean estimate at the three digit level is 4.0, with a standard deviation of 7.9. "Petroleum Oils and Oils from Bituminous Minerals, Crude", "Aircraft and Associated Equipment" or "Fuel Wood" are sectors with relatively high elasticities, whereas "Lighting Fixtures", "Radio-Broadcast Receivers" or 'Motorcycles and Cycles" all rank towards the bottom of their list.

But the comparison is not especially informative in terms of validating our assumptions. What is key here is the Armington assumption that the substitutability between two foreign varieties should equal that between domestic and foreign varieties. If this is true in the data, we can infer directly the price elasticity of imports. Since the latter has been the object of a vast empirical literature, it is with it that we have striven to compare our results.

4.2 Macroeconomic Results

We now turn to macroeconomic estimates of the elasticity of substitution implied by the values in Figure 1. To do so, we apply the aggregation procedure spelled out in the previous sections. Our main point concerns the difference in estimates of η_{NoFirm} and σ_{NoFirm} where $\sigma_k = \sigma$ as against those where σ_k is left unconstrained. Since we have

$$\sigma_{NoFirm} = \sum_k n_{kjj}\ \sigma_k + \sum_k \left(n_{kj} - n_{kjj}\right)\left(\sigma_k - \gamma\right)\left(1 - w_{kj}^M\right)$$

the heterogeneity bias increases in the correlations between σ_k and n_{kjj}. In our data, the correlation is mildly positive, which suggests constraining all elasticities to homogeneity acts to lower estimates of both σ_{NoFirm} and η_{NoFirm}. The second summation has second order effects only, so that the choice of a value for γ has minimal effect on end estimates of σ_{NoFirm} and $\bar{\sigma}_{NoFirm}$. From its definition in equation (6), the calibration of γ does

Table 1: Estimation with common correlated effects

	Import Elasticity η_{NoFirm}	Substitution Elasticity σ_{NoFirm}
Constrained total elasticity	-1.980[a] (.175)	4.124[a] (0.300)
Constrained partial elasticity	-2.738[a] (.262)	3.738[a] (0.263)
Unconstrained total elasticity	-4.508[a] (.745)	7.226[a] (0.962)
Unconstrained partial elasticity	-6.553[a] (1.100)	6.921[a] (0.697)
Number of sectors	56	56
Number of grid searches	11	11

Note: Standard errors in parentheses (obtained by bootstrapping for grid searched sectors), [a] denotes significance at the 1% level.

however have first order effects on the level of η_{NoFirm}, although not on the discrepancy between η_{NoFirm} and $\bar{\eta}_{NoFirm}$. We set γ at its minimal value of 1, so that our results correspond to upper bounds (in absolute values) of η_{NoFirm} and $\bar{\eta}_{NoFirm}$.

Table 1 reports estimates of both aggregate elasticities. We first report estimates of $\bar{\eta}_{NoFirm}$, the total price elasticity of imports, when we impose equal σ_k across sectors. Our point estimate suggests a value for the parameter of -1.98. A confidence interval at standard significance levels implies values ranging roughly from -1.6 to -2.3. This is at the high end of the range of values obtained in conventional estimates of the elasticity based on macroeconomic data. For instance, Goldstein and Kahn (1985) claim that "Harberger's (1957) judgment of 25 years ago that the price elasticity of import demand for a typical country lies in or above the range of -0.5 to -1.0 still seems on the mark". In their Table 4.1, they report estimates for the US between -1.03 and -1.76. The interval is not significantly distinct from the one we estimate, even though the calibration of γ implies an upper bound estimate of $\bar{\eta}_{NoFirm}$. Choosing higher values for γ would only bring our results closer to the literature.

We obtain an aggregate estimate with nothing but import prices, that is not significantly different from one obtained on the basis of relative prices computed with domestic price indices. More importantly, -1.98 is consistent with the choices made in the vast majority of calibration exercises in international macroeconomics. Following the reasoning we develop in the introduction, macroeconomic calibrations infer the elasticity of substitution from $1 - \eta$, where the import elasticity is estimated on aggregate *data*. Here, the implied elasticity of substitution is between two and three, which includes the ranges of values used in, say, Obstfeld and Rogoff (2005) or Backus, Kehoe and Kydland (1994).

Constrained estimates stand in contrast with the value of η_{NoFirm} obtained when σ_k is left unconstrained across sectors. As shown in the Table, the parameter jumps to -4.5, with standard errors that guarantee a significant difference at conventional confidence levels. The right panel of the Table reports the theoretical values for σ_{NoFirm}, which we find is in fact around 7. We argue this is the value that should enter the utility of a representative agent with heterogeneous preferences across sectors. Given the overwhelming evidence that substitutability is heterogeneous across goods or sectors, we contend a value around 7 is preferable from a calibration standpoint. Aggregating the data, instead of aggregating estimates, gives rise to a heterogeneity bias. With heterogeneity, the response of aggregate quantities estimated from aggregated data is not indicative of the average elasticity of substitution.

It is useful to check the heterogeneity bias we document continues to prevail in estimates of the partial import elasticity. We report its constrained and unconstrained values in Table 1. As is patent, a bias continues to prevail, with a constrained estimate at -2.7, jumping to -6.5 when sector specific elasticities are permitted. The adjustment of price indices does not explain our results away. But as expected, it tends to dilute measured elasticity, as partial import elasticities are systematically higher.

We finally verify an aggregated version of our data continues to imply estimates of σ that are commensurate with results in the literature. We implement our estimator on our dataset, aggregated to the country level. Aggregate U.S. imports are given by a simple sum across sectors of the values imported from a given country, and aggregate import prices are computed as the chained Tornqvist index of HS6 specific prices. Identification continues to rest on the cross-section of exporters to the U.S, and requires that aggregate bundles of imported goods be different varieties of the same good, with elasticity of substitution σ. The assumption is palatable at a disaggregated level, but hard to maintain for macroeconomic data, especially if countries are specialized. The bundle of goods exported to the U.S. by developing economies is likely to differ fundamentally from that exported by the developed world, a difference that goes beyond imperfect substitutability between varieties of the same good. We try and alleviate the concern, and focus on a cross-section of 24 high income OECD exporters to the U.S, where the composition of exports is presumably relatively homogeneous.

Based on this cross-section of 24 countries, we obtain an estimate for σ equal to 1.34.

The corresponding price elasticity of imports is close to zero, consistent with conventional aggregate results. The result is reassuring, for it confirms an aggregated version of our data yields perfectly standard estimates. Our data therefore mirror the general conclusions of the literature we have described, with large differences between micro and macroeconomic estimates. Two words of caution are however in order. First, aggregating the data obscures considerably the ceteris paribus nature of our experiment. While estimates of σ_{NoFirm} and $\bar{\sigma}_{NoFirm}$ cannot differ for other reasons than a heterogeneity bias, a number of other parameters are affected by the aggregation of data, rather than of estimates. The dimensionality and sample size of our data are different, as are the variance-covariance properties of import prices and quantities.

Second, aggregating the data effectively forces the substitutabilities between varieties and between sectors to be equal. In our notation, this implies $\sigma = \gamma$, as we show in Appendix A. As all sectors sum up to aggregate imports, it is by construction impossible to distinguish substitution between two varieties of the same sector across countries, or between two different sectors across countries. There are therefore two hypotheses implicit in the use of aggregate data: the homogeneity constraint $\sigma_k = \sigma$ we have discussed, and the additional assumption that $\sigma = \gamma$. On the basis of the vast heterogeneity in existing estimates of σ_k, we have so far speculated the former is quantitatively important. Since our methodology is not equipped to provide structural estimates of γ, we have ignored the latter. In our data, the difference between σ_{NoFirm} and $\bar{\sigma}_{NoFirm}$ - the heterogeneity bias - dwarves that between $\bar{\sigma}_{NoFirm}$ and 1.34. Sample issues notwithstanding, it is the heterogeneity bias that dominates.

4.3 Relevance

Is the correction we document relevant in economic terms? We now discuss the quantitative and qualitative consequences of using heterogeneous estimates to calibrate macroeconomic models with an international dimension. We are careful to focus on theories with CES preferences, since our estimates are predicated on the assumption. We also discuss calibration exercises of models with a representative firm, even though a bias in σ may still be present in *some* theories with heterogeneous firms.

A straightforward implication concerns models directly dealing with the resolution of global imbalances. Most prominently, Obstfeld and Rogoff (2005) use a calibrated model to argue a reversal of the US current account is compatible with a 30% depreciation of the real exchange rate. The calibration sets substitutability at 2. In a slightly simplified two-country version, we obtained depreciation rates of 22 and 21% for values of the parameter of 5 and 7, respectively, down from 31% with an elasticity of 2.[21] The parameter is quantitatively important, and shaves off one third of the "required" depreciation, almost all the way to the 19.3% that obtains for an elasticity of 100. This is true even though Obstfeld and Rogoff's calibration gives prominence to another parameter, the elasticity of substitution between traded and non-traded goods, important in this instance because the US is a largely closed economy. Still, the effects are sizeable and probably important in terms of welfare as well.

Cole and Obstfeld (1991) show the endogenous response of the terms of trade can deliver perfect insurance against country-specific shocks when the elasticity of substitution between domestic and foreign goods is exactly unitary. This is a necessary but not sufficient condition, as perfect insurance also imposes restrictions on other parameters, for instance the intertemporal elasticity of substitution. The result is meant as an illustrative special case of a powerful mechanism. Still, models of international portfolio holdings have drastically different qualitative predictions depending on the magnitude of the parameter. For instance, Heathcote and Perri (2008) show that a model with CES preferences can generate a home equity bias for low enough values of σ. A positive domestic productivity shock will increase the relative return on domestic stocks as long as the terms of trade do not respond too strongly. In a sensitivity analysis, they show portfolio home bias obtains for values of the elasticity of substitution below 4. A contrario, in Coeurdacier (2009), domestic consumers choose to hold foreign assets to insure against shocks to domestic consumption, provided the terms of trade respond strongly enough in response to real shocks, which in his calibration happens for $\sigma = 5$. His sensitivity analysis suggests the conclusion remains robust for values of $\sigma > 1$. Both papers then move on to introduce labor income risk and endogenous production, or incomplete markets. Our purpose here is not to settle the question of the origins of an equity home bias. This depends of course on other calibration choices, not least the intertemporal elasticity of substitution or the existence of trade costs. It is however interesting that most papers in this literature conduct extensive sensitivity analyses to the choice of σ, with far-reaching implications on the models' end predictions.

[21] We are grateful to Cedric Tille for graciously giving us the simulation code.

The policy consequences of international price differences will presumably also depend on the substitutability between domestic and foreign goods. The relevance of the exchange rate in the monetary policy rule developed in Galí and Monacelli (2005) is affected by the parameter. Galí and Monacelli focus on unitary elasticity, so the result is not directly apparent there. But De Paoli (2009) introduce a generalization of their model, and her conclusions point to that direction. As in Cole and Obstfeld (1991) with unitary elasticity, a marginal reduction in the utility value of output is accompanied by an exactly offsetting reduction in the utility value of consumption. This insulates the economy from terms of trade movements. With non unitary elasticity however, policy shocks that affect the terms of trade also affect welfare, in a way that crucially depends on whether the calibrated parameter is above or below one. The fact that we find an aggregate estimate substantially above one must therefore have important policy implications.[22]

The elasticity of substitution of a representative agent living in a one-sector model should be aggregated adequately on the basis of the microeconomic elasticities that we know are heterogeneous. We have shown estimating that average on the basis of macroeconomic data can be misleading. Our recommendation is that σ_{NoFirm} should be preferred to $\bar{\sigma}_{NoFirm}$ in calibrating one-sector theoretical economies if one wants to capture the fact that σ_k *are* heterogeneous in the data. We close with a simple exercise to validate this claim. We construct a two-sector version of a classical model in international economics, due to Backus, Kehoe and Kydland (1994) [BKK], where the sources of heterogeneity are sector-specific elasticities of substitution between domestic and foreign goods, along with w_k and w_k^M. We calibrate this version of the model using our sectoral results, and simulate a J-curve from it. We then compare this prediction with the conventional one-sector model. We ask what value of the (single) elasticity of substitution in the one sector model reproduces the J-curve that is implied by the calibrated two-sector version. We expect an adequately weighted average of the calibrated values of σ_k to come closest to the dynamics implied by the multi-sector version.

Since the workings of the model are well known, we leave a description of the details to Section E of the Technical Appendix. We compare the relative performances of four models. First, a two-sector version of BKK, calibrated on our data and our results. In particular, we choose $(\sigma_1; \sigma_2) = (4.8; 12.9)$, which corresponds to below- and above-mean

[22] Admittedly, other parameters will also determine welfare in this context - not least once again the intertemporal elasticity of substitution.

Figure 2: The J-curve in a two-sector BKK model

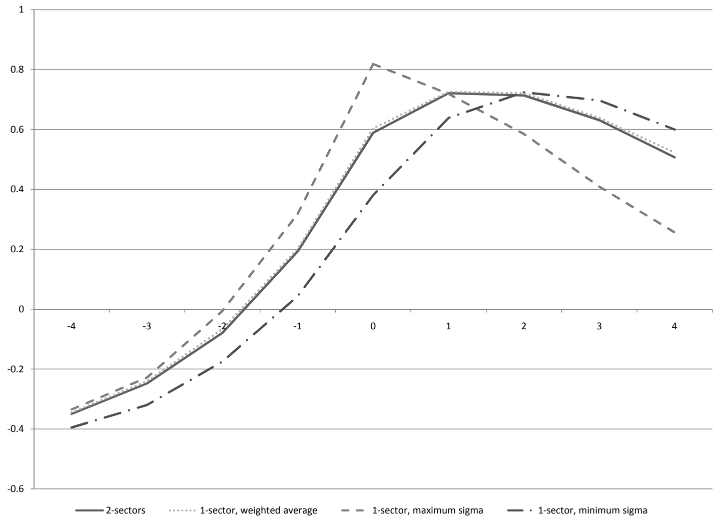

averages of σ_k. The second and third models use σ_1 and σ_2, respectively, in conventional one-sector versions of BKK. Both versions are meant to capture a macroeconomic calibration fully ignorant of heterogeneity issues.[23] Finally, we calibrate a one-sector version of BKK using a weighted average of σ_1 and σ_2, $\sum_{k=1,2} n_{kd}\, \sigma_k$, consistent with the allowance for heterogeneity we have argued matters quantitatively. In all models, we calibrate w_k and w_k^M using our data. Figure 2 reports the J-curves implied by the four models. As is patent, the one-sector version of BKK that best matches the dynamics of the trade balance implied by the two-sector model is one that accounts for heterogeneity in the manner that we have described in this paper.

4.4 Stability

This section verifies the robustness of our results in three dimensions. First, we ascertain our results do not depend on a particular choice of data source in computing w_{kj} and w_{kj}^M. Second, we investigate the importance of "Common Correlated Effects" in obtaining estimates of σ_k. Third, we relax our assumption that elasticities of substitution be identical across the HS6 categories regrouped in each ISIC sector. Instead, like Broda and Weinstein (2006) we estimate a value of σ_k for each HS6 category. We discuss the necessary shortcuts this requires in terms of aggregation.

Table 2 compares the constrained and unconstrained values of the total elasticities σ_{NoFirm} and η_{NoFirm} using different weighting vectors. The data sources and computa-

[23]We also experimented with the arithmetic average of σ_1 and σ_2 with similar conclusions. The arithmetic average is not necessarily an adequate proxy for $\bar{\sigma}$. We have discussed in section 3.2 the possibility that macroeconomic estimates of $\bar{\sigma}$ should suffer from a heterogeneity bias of an econometric nature.

Table 2: Variants on the weights

	Import elasticity		Substitution Elasticity	
	Unconstrained	Constrained	Unconstrained	Constrained
Benchmark	-4.51	-1.98	7.22	4.12
Variant 1	-5.17	-2.21	6.93	4.05
Variant 2	-4.38	-2.08	7.36	4.02
Variant 3	-4.60	-2.15	6.77	4.06
Variant 4	-4.41	-2.10	7.27	4.12

Note: Benchmark: w_{kj}^M using imports and output from IO tables, w_{kj} using STAN sectoral interior demand. Variant 1: w_{kj}^M using imports from BACI and output from STAN, w_{kj} using STAN sectoral interior demand. Variant 2: n_{kj} and w_{kj}^M using imports and output from IO tables, w_{kj} using STAN sectoral interior demand. Variant 3: n_{kj} and w_{kj}^M using imports from BACI and output from STAN, w_{kj} using STAN sectoral interior demand. Variant 4: w_{kj}^M using imports and output from IO tables, w_{kj} using STAN sectoral interior demand (absorption in terms of value added).

tions behind the four alternative variants we present in the Table are discussed in Section 3.3. The first line repeats the results implied by the benchmark weights we have used so far. Across the four variants, constrained estimates of η_{NoFirm} range around -2 and are not significantly distinguishable from conventional estimates, for instance in Goldstein and Kahn (1985). Unconstrained estimates reach -5. The bias continues to be quantitatively important across these four alternatives.

The inclusion of Common Correlated Effects in the estimation of σ_k is justified by our interest in the macroeconomic implications of the microeconomic values we obtain. After all, the quantities traded at sector level, and their prices, do presumably respond to common, aggregate, macroeconomic influences. When estimating the sector-specific substitutability between domestic and foreign varieties, one wants to ascertain one is not capturing aggregate dynamics. This would amount to double-counting at the time of aggregation. Does this matter in our estimations? Table 3 provides a mixed answer, using again our benchmark weights. Without CCE, the constrained estimate of η decreases slightly, to -2.17, whereas the unconstrained estimate increases slightly, to -4.08. These changes are not strongly significant relative to our benchmark results, and they do not alter the conclusion of a significant heterogeneity bias. But they nevertheless suggest the introduction of a CCE term in equation (10) is not innocuous.

Finally, we relax our assumption that the substitutability between two HS6 categories be identical within each ISIC sector. In other words, we allow for heterogeneity in σ_k even within each ISIC sector. Like Broda and Weinstein (2006), we estimate an elasticity of substitution for each HS6 sector, and then use equation (3) to aggregate them at the macroeconomic level. This raises the question of what values for w_{kj} and w_{kj}^M to use:

Table 3: Estimation without common correlated effects

	Import Elasticity η_{NoFirm}	Substitution Elasticity σ_{NoFirm}
Constrained total elasticity	-2.166^a (.150)	4.442^a (.257)
Constrained partial elasticity	-3.016^a (.225.)	4.016^a (.225)
Unconstrained total elasticity	-4.075^a (.112)	6.584^a (.145)
Unconstrained partial elasticity	-5.946^a (.209)	6.321^a (.138)
Number of sectors	56	56
Number of grid searches	12	12

Note: Standard errors in parentheses (obtained by bootstrapping for grid searched sectors), ^{a} denotes significance at the 1% level.

we do not observe any of these weights at such a refined aggregation level. We choose to impose similar values of the weights for all HS6 categories that belong to one ISIC sector. Clearly, this assumes away some possible source of a bias, but there is simply no alternative. But we know choosing other values for w_{kj} and w_{kj}^M does not matter at the ISIC level. At least, this tests whether the heterogeneity in estimates of σ_k within each ISIC category can be such that our conclusions are altered.

Naturally, the constrained estimates of η_{NoFirm} continue to be identical, for instance at -2.17 without a CCE corrective term. After all, this is an estimation that constrains all coefficients to be identical, within and between ISIC categories. The difference arises for unconstrained estimates. We estimated values for σ_k in 4,021 HS6 categories, and aggregated them using our benchmark (ISIC) weights. We obtained a value of -5.23 for the unconstrained elasticity. Computing standard error bands across this point estimate is not tractable, but the aggregation bias appears to be even stronger. We conclude ignoring heterogeneity within ISIC sectors is not important to our conclusions.

5 Conclusion

The elasticity of substitution between domestic and foreign varieties is central in international economics. But no clear consensus has emerged from a vast empirical literature seeking to pin down the parameter, except for one essential finding. On average, microeconomic data tend to imply substantially more heterogeneous values than macroeconomic aggregates. We propose that this heterogeneity is the reason why aggregate results are close to zero. We compute structural estimates of aggregate substitutability allowing or not for heterogeneity at the sectoral level. We find that imposing homogeneity is enough to obtain aggregate estimates in line with the macroeconomic evidence, even using a disaggregated dataset. Allowing for heterogeneity results in an aggregate parameter value of up to 7. This discrepancy validates the conjecture of a heterogeneity bias in elasticity estimates that goes back at least to Houtakker and Magee (1969). Such high parameter values change dramatically the conclusions of calibrated models in most areas of international economics.

A Aggregation and Homogeneity

We want to replicate what is implied by an estimation of the substitutability between the aggregate bundles of domestic and foreign goods, performed on aggregate data. The nominal demand for the bundle of goods imported from country i is expressed in terms its relative price, and overall demand. In our notation:

$$P_{ij}C_{ij} \equiv \sum_{k}^{K} P_{kij}C_{kij} = \left(\frac{P_{ij}}{P_j}\right)^{1-\sigma} P_jC_j, \tag{A.1}$$

where $P_jC_j = \sum_k^K P_{kj}C_{kj}$, P_{ij} is the ideal price index corresponding to the aggregate quantities imported from country i, and P_j is the aggregate price index in country j. By definition, P_{ij} aggregates sector specific prices, using the elasticity of substitution between sectors γ; and P_j aggregates domestic and import prices, using the international elasticity σ. By construction, it is impossible for σ to have a sector dimension. Cross-sector aggregation is performed first, so that the only international substitutability that can prevail is between aggregate bundles of goods. The very notion of an ideal price index for aggregate imports precludes sector heterogeneity in the cross-country elasticity σ. Equation (A.1) is based on an aggregation of the *data*. Instead, our approach in this paper is based on an aggregation of elasticity *estimates*, which is flexible enough that it can preserve heterogeneity in all elasticities, and in particular in σ_k.

Suppose now the world is well characterized by heterogeneous σ_k. We seek to establish what parametric constraints are imposed through the use of aggregate data. We start from our CES model, where goods produced by different exporting countries are substitutable at the industry level. The nominal demand addressed to producers of good k located in country i writes

$$P_{kij}C_{kij} = \left(\frac{P_{kij}}{P_{kj}}\right)^{1-\sigma_k} P_{kj}C_{kj}, \tag{A.2}$$

where $\frac{P_{kij}}{P_{kj}}$ is the relative price of the variety produced in country i, σ_k is the sector-specific elasticity of substitution and $P_{kj}C_{kj}$ is nominal demand for good k. For simplicity, we set the preference shocks β_{kij} to zero. We seek to replicate the setup in equation (A.1) in the context of our disaggregated model. In logarithms, equation (A.2) implies

$$\begin{aligned}\ln P_{ij}C_{ij} &\equiv \ln\sum_k^K P_{kij}C_{kij} = \ln\sum_k^K\left(\frac{P_{kij}}{P_{kj}}\right)^{1-\sigma_k} P_{kj}C_{kj}\\ &\approx \ln\frac{\sum_k^K P_{kij}^{1-\sigma_k}}{\sum_k^K P_{kj}^{1-\sigma_k}} + \ln\sum_k^K P_{kj}C_{kj},\end{aligned}$$

where we use (twice) the property that $\ln\sum_k^K x_k \approx \sum_k^K \ln x_k$ for large K. We seek to recognize the relative price of aggregate imports in the ratio $\frac{\sum_k^K P_{kij}^{1-\sigma_k}}{\sum_k^K P_{kj}^{1-\sigma_k}}$. A well defined ideal price index for aggregate imports exists if we have $\sigma_k = \gamma$. This effectively embeds two constraints. First, the international elasticities of substitution σ_k are homogeneous across sectors. Second, international and cross-sector elasticities are equal. In the body of the paper, we constrain the estimates of σ_k to homogeneity. But our estimation approach is silent about γ, and so we cannot investigate the effect of that additional constraint. We speculate however heterogeneity in σ_k is empirically most relevant, if only because it is documented in decades of applied work. We verify the claim in our data: The estimate for σ we obtain from an aggregated version of our data is in fact close to the results imposing $\sigma_k = \sigma$ reported in the body of the paper.

References

Anderson, J. & Van Wincoop, E., 2004, "Trade Costs", *Journal of Economic Literature,* 84(1): Volume XLII, September.

Backus, D., P. Kehoe & F. Kydland, 1994, "Dynamics of the Trade Balance and the Terms of Trade: The J-Curve?" *American Economic Review*, 84(1): 84-103.

Blonigen, B. & W. Wilson, 1999, "Explaining Armington: What Determines Substitutability Between Home and Foreign Goods?" , *Canadian Journal of Economics*, 32(1): 1-21.

Broda, C. & D. Weinstein, 2006, "Globalization and the Gains from Variety", *The Quarterly Journal of Economics*, 121(2): 541-585.

Chaney T., 2008, "Distorted Gravity: The Intensive and Extensive Margins of International Trade", *American Economic Review*, 98(4): 1707-21.

Chari, V.V., P. Kehoe & E. McGrattan, 2002, "Can Sticky Price Models Generate Volatile and Persistent Real Exchange Rates?", *Review of Economic Studies*, 69(3): 533-563.

Coeurdacier, N., 2009, "Do trade costs in goods market lead to home bias in equities?", *Journal of International Economics*, 77: 86-100.

Coeurdacier N., R. Kollmann & P. Martin, 2007, "International Portfolios with Supply, Demand and Redistributive Shocks", *NBER Working Papers*, 13424.

Cole, H. & M. Obstfeld, 1991, "Commodity trade and international risk sharing : How much do financial markets matter?", *Journal of Monetary Economics*, 28(1): 3-24.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum. 2008, "Global Rebalancing with Gravity: Measuring the Burden of Adjustment,", *IMF Staff Papers*, 55 (3): 511-540.

De Paoli, B., 2009, "Monetary Policy and Welfare in a Small Open Economy", *Journal of International Economics*, Vol.77, No.1, February.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography, and Trade." Econometrica, 70(5): 1741–79.

Feenstra, R., 1994, "New Product Varieties and the Measurement of International Prices", *American Economic Review*, 84(1): 157-177.

Gagnon, J., 2003, "Productive capacity, product varieties, and the elasticities approach to the trade balance", *International Finance Discussion Papers*, 781.

Gali, J. & Monacelli, T., 2005, "Monetary Policy and Exchange Rate Volatility in a Small Open Economy", *Review of Economic Studies*, 72(3): 707-734.

Goldstein M. & M. Kahn, 1985, "Income and price effect in foreign trade", In: J. Ronald and P. Kenen, Editors, Handbook of International Economics, North-Holland, Amsterdam, 1042–1099.

Hanson, G., 2005, "Market Potential, Increasing Returns and Geographic Concentration", *Journal of International Economics*, 67(1): 1-24.

Harberger, A., 1957, "Some Evidence on the International Price Mechanism", *Journal of Political Economy*, 65(6): 506-521.

Head, K. & J. Ries, 2001, "Increasing Returns versus National Product Differentiation as an Explanation for the Pattern of U.S.-Canada Trade", *American Economic Review*, 91(4): 858-876.

Heathcote, J. & F. Perri, 2008, " The International Diversification Puzzle Is Not As Bad As You Think", *NBER Working Papers*, 13483.

Hooper, P., K. Johnson & J. Marquez, 1998, "Trade elasticities for G-7 countries", *International Finance Discussion Papers*, 609.

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

Kemp, M., 1962, "Error of Measurements and Bias in Estimates of Import Demand Parameters", *Economic Record*, 38: 369-372.

Kreinin, M., 1967, "Price Elasticities in International Trade", *Review of Economics and Statistics*, 49: 510-16.

Krugman, P., 1980, "Scale Economies, Product Differentiation, and the Pattern of Trade", *American Economic Review*, 70(5): 950-59.

Marquez, J., 1990, "Bilateral Trade Elasticities", *The Review of Economics and Statistics*, 72: 70-77.

Melitz, M., 2003, "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity", *Econometrica*, 71(6): 1695-1725.

Obstfeld, M. & K. Rogoff, 2000, "The Six Major Puzzles in International Finance: Is There a Common Cause?", *NBER Macroeconomics Annual*, 15.

Obstfeld, M. & K. Rogoff, 2005, "Global Current Account Imbalances and Exchange Rate Adjustments", *Brookings Papers on Economics Activity*, 2005(1): 67-146.

Orcutt, G., 1950, "Measurement of Price Elasticities in International Trade", *The Review of Economics and Statistics*, 32(2): 117-132.

Pesaran, H., 2006, "Estimation and Inference in Large Heterogeneous Panels with a Multifactor Error Structure", *Econometrica*, 74(4): 967-1012.

Pesaran, H. and R. Smith, 1995, "Estimating Long Run Relationships from Dynamic Heterogeneous Panels", *Journal of Econometrics*, LXVIII: 79-113.

Romalis, J., 2007, "NAFTA's and CUSFTA's Impact on International Trade", *The Review of Economics and Statistics*, 89(3): 416-435.

Ruhl, K., 2005, "Solving the Elasticity Puzzle in International Economics", mimeo.

Shiells, J., 1991, "Errors in Import-Demand Estimates Based upon Unit-Value Indexes", *The Review of Economics and Statistics*, 73(2): 378-382.

Stone, J., 1979, "Price Elasticities of Demand for Imports and Exports: Industry Estimates for the U.S., the E.E.C. and Japan", *The Review of Economics and Statistics*, 61(2): 306-312.