

ELASTICITY OPTIMISM

Jean Imbs and Isabelle Méjean

HKIMR Working Paper No.24/2008

December 2008



Hong Kong Institute for Monetary Research

(a company incorporated with limited liability)

All rights reserved.

Reproduction for educational and non-commercial purposes is permitted provided that the source is acknowledged.

Elasticity Optimism*

Jean Imbs[†]

HEC Lausanne

Swiss Finance Institute

CEPR

Hong Kong Institute for Monetary Research

and

Isabelle Méjean[‡]

Ecole Polytechnique

December 2008

Abstract

Estimates of the elasticity of substitution between domestic and foreign varieties are small in macroeconomic data, but substantially larger in disaggregated microeconomic studies. This may be an artifact of heterogeneity. We use disaggregated multilateral trade data to structurally identify elasticities of substitution in US goods. We spell out a partial equilibrium model to aggregate them adequately at the country level. We compare aggregate elasticities that impose equality across sectors, to estimates allowing for heterogeneity. The former are similar in value to conventional macroeconomic estimates; but they are more than twice larger - up to 5 - with heterogeneity. The parameter is central to calibrated models in most of international economics. We discuss the difference our corrected estimate makes in various areas of international economics, including the dynamics of external balances, the international transmission of shocks, international portfolio choice and optimal monetary policy.

Keywords: Trade Elasticities, Aggregation, Calibration, Global Imbalances, International Transmission, International Portfolio, Monetary Policy.

JEL Classification: F41, F32, F21.

* Financial support from the National Center of Competence in Research "Financial Valuation and Risk Management" is gratefully acknowledged. 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 CEPR. We thank Nicolas Coeurdacier, Jordi Galí and Alessandro Rebucci for useful conversations.

† 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

‡ Ecole Polytechnique, isabelle.mejean@polytechnique.edu, <http://mejean.isabelle.googlepages.com>

1. Introduction

The substitutability between domestic and foreign goods is central to most calibrated models in international economics. Depending on the value assigned to the parameter, the predictions of virtually any calibration exercise with an international dimension change quantitatively, sometimes dramatically. The parameter value is evidently key in models that seek to quantify the magnitude of a change in international prices consistent with a rebalancing of external balances as in Obstfeld and Rogoff (2005). It is also a crucial element to the mechanisms that underpin the international diffusion of shocks, as in Kose and Yi (2006) or Corsetti, Dedola and Leduc (2008); or indeed to the extent of international risk sharing in Cole and Obstfeld (1991), and the composition of international portfolio holdings, as in Coeurdacier (2005). Models of international price differences have predictions that rest on the parameter value as well, as demonstrated in Mirdigan (forthcoming) or Atkeson and Burstein (forthcoming). The same is true of the importance of exchange rates in the optimal conduct of monetary policy, as in Galí and Monacelli (2005).

This is quite simply one of the most important parameters in international economics. Unsurprisingly, its calibrated value draws from literally decades of empirical work. Unfortunately, little consensus has emerged from the effort, except for two broad conclusions. First, finely disaggregated good-level quantities are more responsive to (international) relative prices than aggregates. Second, there are enormous differences between goods. Long time ago, Orcutt (1950) referred to an “elasticity pessimism”, which he related to the gap between the low observed volatilities in aggregate quantities and the high volatility of international relative prices. He already conjectured that aggregates could obscure more responsive quantities at the microeconomic level. Here we ask whether this very heterogeneity may not actually be cause for optimism. We propose a correction of the elasticity of substitution between domestic and foreign macroeconomic quantities that accounts for heterogeneity at the microeconomic level.

Why should such a correction be meaningful? The robust finding that microeconomic studies uncover systematically higher estimates than aggregate data points to the possibility of an aggregation bias. But this does not mean aggregate estimates are biased: after all it is the response of aggregate quantities that macroeconomists are interested in, and they are undeniably much smoother than international prices. With heterogeneity however, it does not follow that the aggregate elasticity of substitution be low. For instance, Orcutt (1950) reasoned that “goods with relatively low price elasticities can display the largest variation in prices and therefore exert a dominant effect on the estimated aggregate price elasticity, thereby biasing the estimate downwards.” Aggregation can create a bias because the assumption that all quantities are equally substitutable and adjust homogeneously to changes in relative prices is not supported by the data. As a result there might be systematic differences between the aggregate responsiveness of traded quantities and the preference parameter it is meant to capture. It is of course the latter that should enter calibrated models. The bias we discuss matters for calibration purposes.

Our first goal is to develop a general model telling us how to properly aggregate microeconomic estimates. We construct a measure of aggregate substitutability consistent with a representative agent choosing between aggregates of domestic and foreign quantities, and that can also accommodate the well documented fact that substitutability is vastly different across goods and sectors. We simulate the

standard homogeneity assumption, under which elasticities are identical across sectors. This puts our theory to the test of its ability to reproduce conventional macroeconomic estimates under conventional macroeconomic assumptions. At the same time, the model maps out an index of substitutability between domestic and foreign aggregates, and a weighted average of its disaggregated counterparts. We let theory tell us which weights must be used for aggregation.

Our second goal is to estimate the disaggregated elasticities thus weighted. To do so, we borrow from a methodology introduced by Feenstra (1994a) and recently implemented by Broda and Weinstein (2006). Substitutability between two varieties is identified via the observed cross-country variation in the trade flows towards a given importer. We adapt the approach to our purposes imposing an Armington assumption. Demand at a sectoral level is given by an aggregator of domestic and foreign goods varieties. Crucially, the substitutability between two imported varieties is assumed to be identical to that between foreign and domestic types. Domestic and foreign varieties do differ in terms of transport costs and preferences, but not substitutability. It then becomes possible to use the properties of traded quantities and prices to identify the substitutability of interest. 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.

In practice, the degree of substitutability is usually inferred from the responsiveness of traded quantities to movements in a relevant price. This is possible because, in a conventional framework, a direct, linear mapping exists between a trade elasticity and the substitutability between domestic and foreign goods. For instance, Orcutt (1950), or Houthakker and Magee (1969) deduce the elasticity of substitution from the estimated response of imports to changes in the relative price of imported goods. This raises obvious endogeneity issues, and many approaches have been developed since to address these. But few are structural in nature. This is an attractive feature of what is done here.

Recently, Chaney (2008) showed the mapping between elasticity of substitution and trade elasticity is obscured in the presence of firm heterogeneity and entry dynamics. Traded quantities respond to altered price conditions in a way that does not only reflect consumers' preferences anymore, but also entry decisions on the part of supplying firms. As a result, estimates of the price elasticity of imports are essentially silent about the substitutability implied by preferences. Fortunately, the methodology developed by Feenstra (1994a) directly estimates the preference parameter of interest, without any reference to trade elasticities, and it can accommodate heterogeneous firm dynamics. Our estimates are therefore immune to the criticism, but not their interpretation in terms of a decade-old empirical literature concerned with import elasticities. We use our model to compute theory-implied price elasticities of imports, allowing for firm entry - or not. We are able to spell out the conditions under which firm dynamics act to reinforce an aggregation bias in conventional macroeconomic estimates of the elasticity of substitution between domestic and foreign goods.

In macroeconomic applications, calibration exercises typically favor values of the parameter that are inferred from aggregate estimates of imports price elasticity. This is done for lack of a consensus, often because they are construed 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) use 1.5. The elasticity is (crucially) unitary in Cole and Obstfeld (1991), 0.9 in Heathcote and Perri (2008),

1.5 in Chari, Kehoe and McGrattan (2002), and 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 around 2.5 for the US. This is within the ballpark of the calibrated values used in the macroeconomic literature, and suggests there is nothing special to our data or our approach relative to existing empirical work. We seem to identify the object of interest to many international economists.

With heterogeneity however *aggregate* elasticity estimates more than double, with values in the US up to 5. This reflects the well known fact that microeconomic elasticities tend to be high on average, and also quite dispersed. It is also the result of an appropriate, theory-implied weighting of disaggregated estimates. Does this correction matter economically? We discuss some illustrations in areas as diverse as the resolution of global imbalances, the international transmission of shocks, international risk sharing, portfolio choice, models of the real exchange rate and optimal monetary policy. In short, the parameter is central to most of international macroeconomics. Whether it is in fact 1 or 5 does make a quantitative difference, and calibrated equilibrium responses often change sizably. In some instances, the correction is also relevant qualitatively: the international response to productivity shocks, the justification for a home bias in portfolio investment or the Pareto-ranking of different macroeconomic policies can all alter profoundly with an elasticity substantially above one.

The impact of heterogeneity is significant and robust. Differences between standard and corrected elasticity estimates are statistically significant at conventional confidence levels, and they prevail for a variety of alternative measures or econometric procedures. Our conclusions withstand controls for common components in prices and quantities, and a battery of alternative data sources used to aggregate up microeconomics estimates. Point estimates are minimally affected, and heterogeneity always retains an economically and statistically significant impact.

In what follows, we first present the partial equilibrium 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. 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 larger on average but heterogeneous in disaggregated data. We then lay out a partial equilibrium theory which 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 trade elasticities in ways that may or may not allow for heterogeneity at the good's or sector's level. It is constructed to accommodate heterogeneous firm dynamics.

2.1 Practice

In the absence of sophisticated supply response - for instance firm entry and exit - the substitutability between foreign and domestic goods maps directly with the price elasticity of imports. Estimating the latter, or trade elasticities in general, is an old business in economics. Early contributions used simple regression analysis to evaluate the response of aggregate imported quantities to changes in international relative prices. Houthakker and Magee (1969) for instance, estimate import elasticities between -0.5 and -1.5 depending on the economy considered.¹ The decades that followed saw a flurry of econometric refinements, that purported to alleviate the endogeneity issues that go with having prices as independent variables. 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 based on macroeconomic data were found to be weakly negative, not always significant, and certainly no larger than 2 in absolute value. The conventional inference in macroeconomics is to deduce that the elasticity of substitution between bundles of domestic and foreign goods is close to zero as well.

In contrast, the results obtained from disaggregated information are centered around higher averages, and substantially more heterogeneous. For instance, Blonigen and Wilson (1999) document substitution elasticities between zero 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 short, there is every indication that elasticities are higher on average at a disaggregated level. This comes as no surprise. More than fifty years ago, Orcutt (1950) already remarked that "it is widely recognized that the demand schedule for the product of an individual producer has, in general, far greater price elasticity than the aggregate demand schedule for the entire output of the product" (page 118). The importance of aggregation in the context of trade or substitution elasticities has a straightforward, economic intuition.

Ruhl (2005) proposes an alternative. Cross-sectional elasticity estimates are naturally higher, for they embed 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 are in fact fundamentally different, since in practice they do not capture the extensive margin to the same extent.

¹ Most papers in this literature also take interest in the response of quantities traded to changes in income. We do not discuss this parameter here.

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 (1994a) 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 estimators 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. But that does not mean firm dynamics cannot obscure the theoretical link between trade elasticity and the substitutability parameter in preferences. We now turn to a theoretical treatment of this possibility.

2.2 Theory

Consumption in the domestic economy is an aggregate of imperfectly substitutable goods $k=1, \dots, K$. Utility is given by

$$C = \prod_{k \in K} \frac{C_k^{\alpha_k}}{\alpha_k}$$

where α_k denotes an exogenous preference parameter. Consumption in each sector is derived from a continuum of varieties of good k , that may be imported or not, as in

$$C_k = \left[\sum_{i \in I} (\beta_{ki} C_{ki})^{\frac{\sigma_k - 1}{\sigma_k}} + (\beta_{kd} C_{kd})^{\frac{\sigma_k - 1}{\sigma_k}} \right]^{\frac{\sigma_k}{\sigma_k - 1}}$$

Where i indexes imported types and d is the domestic variety of good k . Crucially, the elasticity of substitution σ_k is specific to each industry, and assumed identical across all varieties, imported or not. β_{ki} lets preferences vary exogenously across varieties, reflecting for instance differences in quality. The sectors that verify $\beta_{ki} = 0$ for all $i \neq d$ will effectively be 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).

We allow for heterogeneity in the continuum of firms that produce each variety i in each sector k . We have:

$$C_{ki} = \left[\int_f c_{kif}^{\frac{\rho_k - 1}{\rho_k}} df \right]^{\frac{\rho_k}{\rho_k - 1}}$$

where f indexes the range of firms that are active in exporting country i and sector k . The range of active firms may differ across countries and sectors, and as in Melitz (2003) or Chaney (2008), it responds endogenously to changes in relative prices. ρ_k denotes the elasticity of substitution between the varieties

produced by firms in sector k , which we assume identical across all countries i . We further assume $\rho_k > \sigma_k$ 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.

The representative maximizing agent chooses her consumption allocation keeping in mind that all imported varieties incur a transport cost τ_{ki} , $i \in I$.² Utility maximization implies that the quantities imported from country i in each sector k are given by

$$C_{ki} = \beta_{ki}^{\sigma_k - 1} \left(\frac{P_{ki}}{P_k} \right)^{-\sigma_k} \alpha_k \frac{P_k}{P} C, \quad i \neq d \quad (1)$$

with

$$\begin{aligned} P_{ki} &= \left[\int_f p_{kif}^{1-\rho_k} df \right]^{\frac{1}{1-\rho_k}} \\ P_k &= \left[\sum_{i \in I} \left(\frac{P_{ki}}{\beta_{ki}} \right)^{1-\sigma_k} + \left(\frac{P_{kd}}{\beta_{kd}} \right)^{1-\sigma_k} \right]^{\frac{1}{1-\sigma_k}} \\ P &= \prod_{k \in K} P_k^{\alpha_k} \end{aligned}$$

The demand addressed to each firm is given by:

$$c_{kif} = \left(\frac{p_{kif}}{P_{ki}} \right)^{-\rho_k} C_{ki} \quad (2)$$

where

$$p_{kif} = \tau_{ki} E_i p_{kif}^*$$

Here, p_{kif}^* denotes the free on board price of variety i of good k produced by firm f , labeled in country i 's currency, and E_i is the bilateral exchange rate with exporter i . We assume producer currency pricing, and impose full pass through. In practice this focuses on estimates of the long run substitution elasticity, in that it excludes short-run dynamics in prices and quantities that stem from nominal rigidities. The paper's empirics are performed at the level of good k , allowing for different varieties across exporters i . The firm-level dimension f is introduced so that the theory-implied estimates we develop can accommodate firm entry dynamics.

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

² We could introduce without loss of generality 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.

disturbances to international relative prices of a specific kind. First, we focus on changes in all relative prices, across all sectors k in K . This assumes away reallocation of demand across *industries*, with relative prices changing by the same amount in all sectors. 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 . 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.³ A natural candidate is a domestic shock to the nominal exchange rate, driven for instance by a monetary disturbance. That will change the international price of domestic goods by an identical amount across all sectors k and exporters i . A domestic aggregate productivity shock would do equally well in moving all relative prices simultaneously. We focus on an “exchange rate” shock to economize on notation. What we really mean is a change in international relative prices across all countries i and all sectors k .

The conventional approach to identifying the substitutability between domestic and foreign varieties using macroeconomic data consists in estimating the response of imported quantities to changes in aggregate relative prices. In the model, this is equivalent to

$$\eta = \frac{\partial \ln [\sum_k \sum_i P_{ki} C_{ki}]}{\partial \ln E}$$

where E denotes the effective exchange rate.⁴ We have taken into account that imported quantities (especially at the aggregate level) are expressed in nominal terms in the conventional datasets on the universe of international trade flows (ComTrade or BACI). We also let these nominal values be expressed in the end consumer’s currency, and inclusive of transport costs or tariffs (i.e. Cost, Insurance, Freight). This is innocuous: we later show that the estimation procedure is robust to the presence or otherwise of τ_{ki} in measured trade flows.

Equation (1) can now be used to obtain a model-implied expression for η . In particular,

$$\begin{aligned} \eta = & \sum_{k \in K} \sum_{i \in I} n_{ki} (1 - \sigma_k) \frac{\partial \ln P_{ki}}{\partial \ln E_i} + \sum_{k \in K} \sum_{i \in I} n_{ki} w_{ki}^M (\sigma_k - 1) \frac{\partial \ln P_{ki}}{\partial \ln E_i} \\ & + \sum_{k \in K} \sum_{i \in I} n_{ki} w_{ki}^M w_k \frac{\partial \ln P_{ki}}{\partial \ln E_i} \end{aligned} \quad (3)$$

³ 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 to perform aggregation are just slightly different, and identification becomes more complicated since it relies on the crosssection of all exporters to a given destination.

⁴ Given our assumption of a uniform shock across all i , this is equivalent to $\sum_i \frac{\partial \ln [\sum_k P_{ki} C_{ki}]}{\partial \ln E_i}$ since $\frac{\partial E_i}{\partial E} \frac{E}{E_i} = 1$.

with

$$\begin{aligned} n_{ki} &\equiv \frac{P_{ki} C_{ki}}{\sum_{k \in K} \sum_{i \in C^k} P_{ki} C_{ki}} \\ w_{ki}^M &\equiv \frac{\partial \ln P_k}{\partial \ln P_{ki}} = \frac{P_{ki} C_{ki}}{P_k C_k} \\ w_{ki}^M \cdot w_k &\equiv \frac{\partial \ln P}{\partial \ln E} = w_k^M \cdot \frac{P_k C_k}{P C} \end{aligned}$$

n_{ki} is the share of imports from country i in sector k in their overall value. w_{ki}^M is the weight of imports from country i in sector k 's nominal consumption, and w_k is the share of sector k in aggregate nominal consumption.

2.2.1 No Firm Heterogeneity

We clarify the separate influences of the issues raised by aggregation across all sectors k in K , and firm heterogeneity within each sector k . In particular, we first focus on the former and shut down firm dynamics. Since we are interested in long run estimates and impose full pass-through, in each sector k , each representative exporter i responds identically to changes in the exchange rate. No differences exist across countries in the entry or exit decisions of supplying exporters in response to a shift in relative prices. In short, $\frac{\partial \ln P_{ki}}{\partial \ln E_i} = 1$. Equation (3) simplifies into

$$\eta_{NoFirm} = 1 - \sum_{k \in K} n_k \sigma_k + \sum_{k \in K} n_k w_k^M (\sigma_k - 1) + \sum_{k \in K} n_k w_k^M w_k \quad (4)$$

with

$$\begin{aligned} n_k &\equiv \frac{\sum_{i \in C^k} P_{ki} C_{ki}}{\sum_{k \in K} \sum_{i \in C^k} P_{ki} C_{ki}} \\ w_k^M &= \frac{\sum_{i \in C^k} P_{ki} C_{ki}}{P_k C_k} \end{aligned}$$

The aggregate import elasticity η_{NoFirm} contains three terms. First, a weighted average of industry-specific elasticities, with weights corresponding to the importance of sector k in overall imports. 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 , the composition of the ideal price index in sector k changes significantly in response to the shock considered. The third term reflects the adjustment of the aggregate price index P , in response to an aggregate shock to relative prices affecting all sectors. In the words of Orcutt (1950), η denotes a *total* elasticity, which ‘‘indicates the percentage change in the quantity imported per percentage change in the price of the import, with allowance made for the full response of [...] competitors’’ (page 118).

The second and third terms in equation (4) exist because of our focus on an *aggregate* multilateral shock in relative prices. If instead the shock considered were microeconomic in nature and focused on a specific good k -a change in tariff- then under standard atomistic assumptions and large K , the third

term in equation (4) would disappear. Similarly, w_k^M is non zero because we assume all relative prices across all exporting economies change identically; for large values of I , a change in a given bilateral exchange rate would have negligible effect on P_k .

Two differences arise between the aggregate elasticity of imports to international prices and the preference parameter it is meant to estimate. One is related to the possibility of sectoral heterogeneity in the elasticity parameter, i.e. $\sigma_k \neq \sigma_s, k \neq s$. The other has to do with the effect of macroeconomic shocks on (sectoral and aggregate) price indices. This possibility is well known in the literature. Orcutt (1950) introduces a *partial* elasticity measure, that assumes aggregate price responses away and captures “what would happen if the price of domestic substitutes were held constant” (page 118).

Focus for now on heterogeneity and suppose $w_k^M = w_k = 0$. The aggregate substitutability between foreign and domestic varieties is given by a weighted average of each industry’s corresponding preference parameter. This captures the direct effect of a shock to E on P_{ki} and the resulting immediate change in C_{ki} visible from equation (1). In macroeconomic data, traded quantities are summed up to the country level, and the econometrician asks how these aggregates respond to changes in international prices. In other words, she assumes $\sigma_k = \sigma$ for all k . Under this constraint, $\bar{\eta}_{NoFirm} = 1 - \sigma$ where σ is the elasticity of substitution between domestic and foreign varieties in *all* sectors, and $\bar{\eta}$ denotes the “macroeconomic” value of imports elasticity that obtains under homogeneity. As expected, the mapping between aggregate import elasticity and substitutability is linear and straightforward. In practice however, $\sigma_k \neq \sigma_s, k \neq s$, and the aggregate (partial) price elasticity of imports is a biased measure of the aggregate elasticity of substitution. The magnitude of the bias depends directly on the cross-sector correlation between import shares n_k and their substitutability σ_k .

In addition, both sectoral and aggregate price indices respond to macroeconomic shocks. Consider a positive shock to the nominal exchange rate, driving up the domestic price of imports, p_{ki} . The second term in equation (4) captures the response of P_k , which by definition increases. This has two effects. First, it mitigates how expensive the domestic variety of good k is relative to its substitutes, since good k as a whole has become dearer. This lowers (in absolute value) the *measured* (total) elasticity, since it acts to dilute the initial change in $\frac{P_{ki}}{P_k}$. Second, it induces substitution away from good k , suddenly relatively more pricey, and increases (in absolute value) the *measured* (total) elasticity since it reinforces the initial shift away from the domestic variety of good k . Which effect dominates depends on the sign of $\sigma_k - 1$; with higher substitutability between varieties than between goods, the term is positive. In any event, both effects increase with the importance of imports in sector k , and the importance of sector k in overall imports.

The third integral in equation (4) reflects the increase in the overall price index P , which limits substitutability away from good k . This acts to mitigate the response of the quantities demanded of the domestic variety of good k , and thus lowers *measured* total elasticity (in absolute value). Under plausible parameterizations therefore, the total price elasticity of imports will be lower in absolute value than its partial counterpart. The responses of price indices tend to dilute the initial response of quantities.

Estimates of the price elasticity of imports obtained from aggregate data provide biased values of the true preference parameter that reflects the aggregate substitutability between domestic and foreign varieties. The bias arises because differences in good-specific substitutability are assumed away. This

is true of partial and total elasticity, and we later report results for both concepts, correcting for heterogeneity or not. We do so because, in the words of Orcutt (1950), partial elasticities have “generally been estimated for particular commodities, while total have been estimated for all, or large groups, of imports lumped together” (page 118). In what follows, we infer from our structural estimates of σ and σ_k the corresponding, theory-implied (partial or total) values for $\bar{\eta}_{NoFirm}$ and η_{NoFirm} . The hope is that the implied values for $\bar{\eta}_{NoFirm}$ are in line with classic results from the macroeconomic empirical literature, whereas the implied values for η_{NoFirm} are not.

2.2.2 Allowing for Firm Heterogeneity

We now reintroduce firm heterogeneity in the definition of the aggregate price elasticity of imports, given in equation (3). Exporter’s prices P_{ki} are now indices, reflective of the continuum of firms active in sector k and country i . This range will vary endogenously in response to changes in international relative prices, and this response itself may differ across sectors and exporters because of differences in transport costs or assumptions on firms distribution. The argument is similar to Chaney (2008), who shows the price elasticity of imports and the elasticity of substitution become theoretically disjoint in the presence of an extensive margin. In Appendix A, we go through what is now a standard model of the supply decisions reached by heterogeneous firms in each sector and each country. As we describe in the Appendix, the price charged in country i and sector k is given by

$$P_{ki} = \left[\frac{M_{ki}}{1 - G_{ki}(\bar{\varphi}_{ki})} \int_{\bar{\varphi}_{ki}}^{\infty} p_{kif}(\varphi)^{1-\rho_k} dG_{ki}(\varphi) \right]^{\frac{1}{1-\rho_k}} \quad (5)$$

where M_{ki} is the mass of exporting firms active in sector k and country i . $G_{ki}(\varphi)$ denotes the cumulative distribution function for firms’ productivity in sector k and country i , and $\bar{\varphi}_{ki}$ is an (endogenous) productivity cut-off level, above which firms are active on export markets. Equation (5) clearly implies that $\frac{\partial \ln P_{ki}}{\partial \ln E_i}$ may vary across sectors or countries, with end effects on the definition of the aggregate price elasticity of imports, given in equation (3). In particular, we have

$$\frac{\partial \ln P_{ki}}{\partial \ln E_i} = 1 - \frac{1}{1 - \rho_k} \left(\frac{p_{kif}(\bar{\varphi}_{ki})}{P_{ki}} \right)^{1-\rho_k} M_{ki} g(\bar{\varphi}_{ki}) \bar{\varphi}_{ki} \frac{\partial \ln \bar{\varphi}_{ki}}{\partial \ln E} + \frac{1}{1 - \rho_k} \frac{\partial \ln M_{ki}}{\partial \ln E} \quad (6)$$

The first term measures the response of individual firm prices to exchange rate shocks, i.e. an intensive margin of adjustment. Under full pass-through, it is unitary. The second term captures the response of the productivity cut-off, $\bar{\varphi}_{ki}$, as firms enter export markets in response to shifting relative prices. Finally the third term corresponds to the adjustment in the potential number of firms serving the foreign market. These last two terms capture the extensive margin of price adjustments, equal to zero under firm homogeneity.

In the short run, the mass of firms in country i and sector k is usually taken as given, and the third term in equation (6) is equal to zero. In this case, the extensive margin works through the medium-run response of the productivity cut-off for a constant mass of potential exporters. This is the avenue followed in Chaney (2008), who shows the response is positive, increasing in the degree of firm heterogeneity and

decreasing in the degree of substitutability between varieties. Given our concern with long run elasticities, we depart from this approach and focus instead on the long run impact of changes in relative prices, when the potential number of exporters is endogenous.

When firms productivity follow a Pareto distribution, we show in Appendix A that the productivity cut-off $\bar{\varphi}_{ki}$ depends only on exogenous parameters in a long run equilibrium where M_{ki} adjusts. In other words, it is now the second term in equation (6) that is equal to zero. The extensive margin works through the number of potential entrants. In the Appendix, we derive an expression for M_{ki} which can be used to compute $\frac{\partial \ln M_{ki}}{\partial \ln E}$. It is then straightforward if tedious to substitute the expression back into equation (6) to obtain

$$\frac{\partial \ln P_{ki}}{\partial \ln E_i} = 1 + \frac{\sigma_k - (\sigma_k - 1)w_{ki}^M - w_{ki}^M w_k}{\rho_k - \sigma_k + (\sigma_k - 1)w_{ki}^M + w_{ki}^M w_k}$$

The presence of an extensive margin drives the response of P_{ki} to changes in international prices away from one, i.e. away from the case without firm entry. To gain intuition, it is useful to recognize negligible terms. In particular, w_{ki}^M denotes the share of a given exporter in overall consumption in sector k , presumably a small number for large I . Approximating out the terms in w_{ki}^M , reflective of the adjustment in aggregate price indices, we get

$$\frac{\partial \ln P_{ki}}{\partial \ln E_i} = 1 + \frac{\sigma_k}{\rho_k - \sigma_k} + o(\epsilon^2)$$

where $o(\epsilon^n)$ denotes an approximation residual of order n . Under our assumptions, the extensive margin of price adjustments is positive as in Chaney (2008). It depends on the substitutability between the varieties of good k produced by different firms in a given country, ρ_k , as against the substitutability between the varieties of good k produced in different countries. Because unlike Chaney (2008) we have focused on a long run equilibrium, the correction does not depend on the parameters that govern the distribution of firms productivity.

How does this price response alter the model-implied value of the aggregate price elasticity of imports? We substitute the equilibrium response of prices into equation (3), to obtain

$$\begin{aligned} \eta_{Firm} &= \sum_{k \in K} (1 - \sigma_k) n_k \left(1 + \frac{\sigma_k}{\rho_k - \sigma_k} \right) \\ &+ \sum_{k \in K} (\sigma_k - 1) \left(1 + \frac{\sigma_k}{\rho_k - \sigma_k} \right) \sum_{i \in I} n_{ki} w_{ki}^M \\ &+ \sum_{k \in K} w_k \left(1 + \frac{\sigma_k}{\rho_k - \sigma_k} \right) \sum_{i \in I} n_{ki} w_{ki}^M + o(\epsilon^3) \end{aligned} \quad (7)$$

The analogy with the theory-implied expression for η_{NoFirm} is illuminating. Consider first the (partial) price elasticity of imports in the presence of firm relocation, when all elasticities are constrained to homogeneity across sectors. The model implies

$$\bar{\eta}_{Firm} = (1 - \sigma) \left(1 + \frac{\sigma}{\rho - \sigma} \right)$$

A structural estimate of σ is now insufficient to compute the corresponding theory-implied price elasticity of aggregate imports. There is a bias, which, under our assumptions, is positive. From an empirical standpoint, estimates of the price elasticity of imports should be larger in absolute value when the data embed firm relocation decisions, an argument reminiscent of Ruhl (2005). In as much as some data do embed firm relocation decisions - as does cross-section information, according to Ruhl (2005) - some direct estimates of the aggregate price elasticity of imports should reflect the true value of $\bar{\eta}_{Firm}$. Few if any have estimated values for $\bar{\eta}$ significantly larger in absolute value than -2 , however. That is too low to warrant any elasticity optimism, and it suggests $\bar{\eta}_{Firm}$ cannot exceed -2 . Actually we later identify estimates of $1 - \sigma$ not significantly different from -2 either. This suggests $\bar{\eta}_{Firm}$ is in fact close to $1 - \sigma$ and the empirical counterpart to $\frac{\sigma}{\rho - \sigma}$ is small in our data.

The important question for our purpose is whether the discrepancy between $\bar{\eta}_{Firm}$ and η_{Firm} differs sizably from that between $\bar{\eta}_{NoFirm}$ and η_{NoFirm} . The previous section documented an aggregation bias increasing in the cross-sector correlation between n_k and σ_k . We now examine how the conclusion alters in the presence of firm dynamics. Allowing for heterogeneity in the elasticity parameters, the model with firm dynamics implies

$$\eta_{Firm} = \sum_{k \in K} (1 - \sigma_k) n_k \left(1 + \frac{\sigma_k}{\rho_k - \sigma_k} \right)$$

As in the absence of firm dynamics, the (partial) price elasticity of imports implied by structural estimates of σ_k increases with the cross-sector correlation between n_k and σ_k . There is still an aggregation bias. But with firm entry, the bias is actually magnified. For a given positive correlation between n_k and σ_k , the theory-implied bias $\eta_{Firm} - \bar{\eta}_{Firm}$ is larger than its counterpart without firm entry, $\eta_{NoFirm} - \bar{\eta}_{NoFirm}$.⁵ In our empirical implementation, we have to ignore the extensive margin when inferring values of the price elasticities of imports, because we cannot identify ρ_k . We just showed that if anything, this will under-estimate the extent of an aggregation bias. The value for η_{NoFirm} that we compute from observed n_k and estimated σ_k is actually lower than the true price elasticity accounting for heterogeneity and firm entry, η_{Firm} . Our bias estimate is a lower bound.

⁵ This will be true as long as $corr(n_k, \rho_k) < \frac{\rho_k}{\sigma_k} corr(n_k, \sigma_k)$. The condition is plausible, since the correlation between ρ_k and the import share of sector k is likely to be smaller than that between σ_k and the import share. After all, ρ_k denotes a substitutability between firms within a sector, whereas σ_k captures a substitutability between countries.

As in the previous section, the responses of aggregate and sectoral price indices will tend to dilute the overall response of quantities. Total elasticities are closer to zero than partial, whether firm dynamics are allowed or not. Under our assumptions, the responses of price indices in the second and third terms of equation (7) continue to have the same sign as in the previous section without firm heterogeneity. They also tend to be of smaller order of magnitude than the first term.

3. Identification

We review how the methodology in Feenstra (1994a) is adapted to our purposes. 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 σ , 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 (1994a) 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.

Contrary to a large literature in Industrial Organization, the identification of a structural demand parameter is achieved in the absence of any supply-shifting instruments. We now explain succinctly how this obtains. 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 and $\sigma_k > 0$ the elasticity of substitution between varieties of good k . Unlike Feenstra (1994a) or Broda and Weinstein (2006), we are not seeking to compute ideal price indices and so can allow for complementarities, and values of σ_k strictly below unity.

Feenstra (1994a) imposes a simple supply structure, and assumes

$$P_{kit} = \tau_{kit} E_{it} \exp(v_{kit}) C_{kit}^{\omega_k}$$

where v_{kit} denotes a technological shock that can take different values across sectors and exporters, and ω_k is the inverse of the price elasticity of supply in sector k .⁶ This assumes production decisions are

⁶ We follow Feenstra (1994a) and assume all exporters have the same supply elasticity.

taken on the basis of the price net of transport costs, and labeled in domestic currency. In Appendix B we show how this expression can arise from the model with heterogeneous firm entry developed in the previous sections.

In practice, the approach uses expenditure shares, not least 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

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

Since the data focus on traded goods only, we do not observe domestically produced consumption. In addition, prices are measured Free on Board but in the importer's currency. 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

$$\tilde{s}_{kit} \equiv \frac{\tilde{P}_{kit}C_{kit}}{\sum \tilde{P}_{kit}C_{kit}} = \frac{s_{kit}}{\tau_{kit}} \left(1 + \frac{P_{kdt}C_{kdt}}{\sum \tilde{P}_{kit}C_{kit}} \right) \equiv \frac{s_{kit}}{\tau_{kit}} \mu_{kt}$$

Taking logarithms, it is straightforward to rewrite demand as

$$\Delta \ln \tilde{s}_{kit} = (1 - \sigma_k) \Delta \ln \tilde{P}_{kit} + \Phi_{kt} + \varepsilon_{kit} \quad (8)$$

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 (1994a) 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

$$\Delta \ln \tilde{P}_{kit} = \Psi_{kt} + \frac{\omega_k}{1 + \omega_k \sigma_k} \varepsilon_{kit} + \delta_{kit} \quad (9)$$

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. $\delta_{kit} \equiv \frac{1}{1 + \omega_k \sigma_k} [\Delta \ln E_{it} + \Delta u_{kit}]$ is an error term encapsulating movements in the exchange rate and technological developments in country i and sector k .

Under standard assumptions on taste shocks β_{kit} and technology shocks u_{kit} , it is possible to identify the system formed by equations (8) and (9). 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 :

$$\begin{cases} \Delta \ln \tilde{s}_{kit} - \Delta \ln \tilde{s}_{krt} = (1 - \sigma_k) (\Delta \ln \tilde{P}_{kit} - \Delta \ln \tilde{P}_{krt}) + \tilde{\varepsilon}_{kit} \\ \Delta \ln \tilde{P}_{kit} - \Delta \ln \tilde{P}_{krt} = \frac{\omega_k}{1 + \omega_k \sigma_k} \tilde{\varepsilon}_{kit} + \tilde{\delta}_{kit} \end{cases}$$

where $\tilde{\varepsilon}_{kit} \equiv \varepsilon_{kit} - \varepsilon_{krt}$ and $\tilde{\delta}_{kit} = \delta_{kit} - \delta_{krt}$.⁷ We follow Feenstra (1994a) and summarize the information contained in the system with the following estimable regression

$$Y_{kit} = \theta_{1k}X_{1kit} + \theta_{2k}X_{2kit} + u_{kit} \quad (10)$$

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 $u_{kit} = \tilde{\varepsilon}_{kit}\tilde{\delta}_{kit}\frac{(\sigma_k-1)(1+\omega_k)}{1+\omega_k\sigma_k}$. Estimates of equation (10) map directly with the parameters of interest, since

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

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

The coefficient estimates are left unaffected by either trade costs or exchange rates movements, even though they can differ across source countries and sectors. Empirically, prices can indifferently be measured Free on Board or Cost, Insurance, Freight, and in either importer's or exporter's currency.

Equation (10) still suffers from an endogeneity issue. We follow Feenstra (1994a) and instrument the regressors with country-specific fixed effects, and correct the estimation for heteroskedasticity across exporters i . We also 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, and allow for intercepts specific to each HS6 sector.⁸

The system summarized by equation (10) 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, 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 u_{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

$$Y_{kit} = \theta_0 + \theta_{1k}\hat{X}_{1ki} + \theta_{2k}\hat{X}_{2ki} + \theta_{3k}X_{1it} + \theta_{4k}X_{2it} + u_{kit} \quad (11)$$

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} .

⁷ In the empirics, we choose a reference country that is present in the US market during the whole observed period.

⁸ For the instrumentation to be consistent, there must be some cross-country differences in the relative variance of the demand and supply curves. For an intercept to capture measurement error, its variance must be equal across exporting countries. Of course, it may still be specific to each sector. See Feenstra (1994a, 1994b).

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

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

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

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

As is apparent, there are combinations of estimates in equation (11) that do not correspond to any theoretically consistent estimates of $\hat{\sigma}_k$. Even though our universe of acceptable values is slightly broader than Broda and Weinstein (2006) - whose purpose was different - this is a problem we encounter in our data, as they did. We follow their approach, and use a search algorithm that minimizes the sum of squared residuals in equation (11) over the intervals of admissible values of the supply and demand elasticities. 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_{0k} , 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 (11). The corresponding standard errors are obtained via bootstrapping of the procedure using 1,000 repetitions.

3.2 Macroeconomic Estimates

We invoke equation (3) to aggregate adequately the estimates of $\hat{\sigma}_k$ just obtained. Appendix C details how the corresponding standard deviation is derived, on the basis of a first-order approximation around the true value of the coefficient. This provides the value and confidence intervals for the aggregate parameter of interest, allowing for heterogeneity at the sectoral level. 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 (11) into

$$Y_{kit} = \theta_0 + \theta_1 \hat{X}_{1i} + \theta_2 \hat{X}_{2i} + \theta_3 X_{1it} + \theta_4 X_{2it} + u_{kit} \quad (12)$$

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 (12) is now estimated on the pooled dataset formed by observations on all sectors. It is noteworthy that identification in equations (11) and (12) 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.

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 import elasticity are given by

$$\bar{\eta} = 1 - \hat{\sigma} + (\hat{\sigma} - 1) \sum_{k \in K} n_k w_k^M + \sum_{k \in K} n_k w_k^M w_k$$

with a standard error given by a first-order approximation, as detailed in Appendix C. Following the definitions in Orcutt (1950), $\bar{\eta}$ is the constrained total elasticity, and $1 - \hat{\sigma}$ 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 (1994a) 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 (12) is effectively estimated on a panel of HS6 sectors. Our partition of HS6 sectors into 56 ISIC categories is entirely innocuous here. Homogeneity is imposed across all sectors, and two HS6 goods have to be equally substitutable whether or not they belong to the same ISIC sector. The partition is only relevant for the unconstrained estimation, and is performed for lack of detailed information on n_k , w_k^M or ω_k 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_k , w_k^M and ω_k .

The approach adapted from Feenstra (1994a) requires relatively little information on traded flows. To estimate equation (11) 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.⁹ 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 $\tilde{s}_{kit} = \frac{\tilde{P}_{kit}C_{kit}}{\sum_{i \in I} \tilde{P}_{kit}C_{kit}}$.

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 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 73 percent of the total value of US imports, across 56 ISIC sectors.

The three weights used in aggregation come from a variety of sources. In the main body of the text, we privilege results that correspond to values of n_k , w_k^M and ω_k taken from as few data sources as possible. In particular, n_k and w_k^M are first computed from the US input/output (IO) tables, available in the ISIC (Revision 3) nomenclature. n_k is defined as the 1997 ratio of sectoral over total import values, and w_k^M is the 1997 ratio of imports over domestic gross output. In the main body of the text, ω_k is taken from the STAN database, as the 1997 ratio of sectoral relative to total interior demand (given by gross output and imports net of exports).

In section 4.4, we verify our results do not depend on this specific choice of data sources. We discuss four alternatives. In all cases, ω_k is left unchanged and computed using STAN 1997 data. We focus the sensitivity analysis on alternative measures of n_k and w_k^M , because equation (3) suggests ω_k has second order impact on estimates of η . In a first alternative, the import data necessary to compute n_k and w_k^M are taken from BACI (in 1997), while the output data continue to come from the IO tables. An advantage is that the same data are then (partly) used to compute the relevant weights and to estimate the elasticities. In a second alternative, we still use imports from BACI but the output data are taken from STAN, and the ratios use 1999 data.

In these two cases, the values of w_k^M are computed on the basis of two different datasets. In a third alternative, we compute n_k and w_k^M using a single source, n_k from BACI, w_k^M from the I/O tables, all measured in 1997. Finally, up to now, w_k^M is measured using output data, while the theoretical parameter refers to total sectoral consumption. Our fourth and final alternative defines w_k^M as the 1999 ratio of sectoral imports over interior demand, as implied by gross output and imports net of exports, obtained from STAN.

⁹ 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.

4 Results and Relevance

We first review the microeconomic estimates, obtained across 56 ISIC sectors, and relate them with existing evidence. We then aggregate them, first preserving heterogeneity. We compare the results with our estimation imposing homogeneity - using the same data sources to perform the aggregation. We discuss the discrepancy, and in particular whether its magnitude is significant *economically*. We argue the corrected 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 28 to 3.1 and a standard deviation of 4.9. The median value is only 5.1, reflecting a skewed distribution of elasticities: only 5 out of 56 estimates are above 10. How do our results compare 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. In particular, they report cross-sector averages of microeconomic estimates, rather than estimates based on macroeconomic data.

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). The price elasticity of imports is then given by $1 - \hat{\sigma}_k$. These elasticities were the object of a vast literature spread over the 1960s and 1970s, with two differences. First, in most instances, the results put forward focused on short-run elasticity estimates, typically for reasons of data availability especially at a disaggregated level. In contrast our estimates correspond to long run elasticities. Second, the data used were more coarse, focusing 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. 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. Our assumptions find comfort in this convergence of results.

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, sector-specific estimates with the literature becomes quickly 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 -2.22, -2.32, and -3.71 in “Rubber Products”, “Plastic Materials and Articles” and “Dyeing, Tanning and Coloring Agents”, as against -4.1, -3.1 and -5.9 in “Rubber Products”, “Manufactures of Plastic Products” and “Tanning and Dressing of Leather” in Figure 1. Keeping in mind ours are estimates of long run elasticities, 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 -3.597, indistinguishable from our value of -3.4 for “Publishing”. He also finds -3.489 in “Steel Plate and Sheet”, not that dissimilar from 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 vindicates the Armington assumption. The few punctual comparisons we report at the level of individual sectors do, as well.

There is of course an obvious comparison that has been absent from our analysis so far. We have implemented a variant of the methodology introduced by Feenstra (1994a), 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 introduced. 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 equation (3). Our main point concerns the difference in estimates of η where $\sigma_k = \sigma$ as against those where σ_k is left unconstrained. As is patent, the bias thus induced will increase in the correlations between σ_k and n_k , and between σ_k and $n_k \cdot w_k^M$. The third term in equation (3) has second order effects only.

Table 1 reports both correlations across the five variants we use to compute n_k and w_k^M . In all cases, the correlations are mildly positive. This suggests constraining all elasticities to be the same will act to lower estimates of both η_{NoFim} and η_{Fim} . Under plausible conditions, the bias will be larger on η_{Fim} . In fact, we have selected the benchmark weights that tend to minimize the extent of a putative aggregation bias, with lowest values for both correlations in the Table. Since the extent of the bias increases in both correlations, Table 1 suggests it will be larger in Variants 2 or 3. The third term in equation (3) can of course slightly affect that ranking, for large enough changes in the correlations between n_k , w_k^M , and ω_k across variants.

Table 2 reports estimates of the aggregate price elasticity of imports. We first report estimates of η , the total elasticity implied by equation (3), when we impose that σ_k be equal across all sectors. Our point estimate suggests a value for the parameter of -1.9, with narrow standard error bands. A confidence interval at standard significance levels implies values ranging roughly from -1.7 to -2.1. 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, an interval that is not significantly distinct from the one we estimate.

Once again, this vindicates the assumption of an Armington aggregator. We obtain an aggregate estimate with nothing but import prices, that is virtually identical to one obtained on the basis of relative prices computed with domestic price indices. More importantly, -1.9 is consistent with the choices made in

the vast majority of calibration exercises in international macroeconomics. 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).

This stands in contrast with the value of η obtained when σ_k is left unconstrained across sectors. As shown in the Table, the parameter jumps to -4.2, with standard errors that guarantee a significant difference at conventional confidence levels. This more than doubles the corresponding value for the estimated elasticity of substitution, at 5.2. 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 5 is preferable from a calibration standpoint. Such a high value characterizes adequately the *average* substitutability of a representative agent who has heterogeneous preferences across goods.

Why the discrepancy? Does this not mean that aggregate quantities should be measurably more responsive to aggregate relative prices? Orcutt (1950) already speculated an explanation, reported by Goldstein and Khan (1985). On page 1070, they explain that, “in aggregate trade equations, goods with relatively low price elasticities can display the largest variation in prices and therefore exert a dominant effect on the estimated aggregate price elasticity, thereby biasing the estimate downwards.” We view the results in Table 2 as a confirmation of this decade-old conjecture.

Orcutt (1950) introduced the notion of a partial elasticity that assumes away aggregate price responses in equation (3). From Section 2, this is simply defined as a weighted average of σ_k , with weights given by n_k . Strictly speaking, this weighted average is the direct equivalent of the aggregate elasticity of substitution, although disaggregated data are necessary to calculate it. We have referred to $1 - \eta$ as the natural proxy for aggregate substitutability, because macroeconomic calibrations would typically infer the elasticity of substitution from the price elasticity of imports, estimated *on the basis of macroeconomic data*. It is therefore the right comparison with the literature.

It is nevertheless useful to check the aggregation bias we document continues to prevail in estimates of the partial elasticity. We report its constrained and unconstrained values in Table 2. As is patent, a bias continues to prevail, with a constrained estimate at -2.7, jumping to -6.0 when sector specific elasticities are permitted. The adjustment of price indices in equation (3) does not explain our results away. But as expected, it tends to dilute measured elasticity, as partial elasticities are systematically higher. The Table also says that a “true” estimate of the aggregate elasticity of substitution, computed purely on the basis of sector specific estimates, and holding all (sectoral and aggregate) price indices constant is effectively equal to $1 - (-6.0) = 7.0$.

4.3 Relevance

Is the correction we document relevant in economic terms? We now discuss the quantitative and qualitative consequences of our corrected estimates. The most straightforward implication concerns models directly dealing with the responsiveness of traded quantities to relative prices, and in particular 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 2.¹⁰ The parameter is quantitatively important, and shaves off two-thirds 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) showed 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. The result is meant as an illustrative special case of a powerful mechanism. Still, models of international portfolio holdings will have drastically different qualitative predictions depending on which side of one the parameter lies. For instance, Heathcote and Perri (2008) show that, with complete markets, imperfect substitutability between domestic and foreign goods can generate a home equity bias. 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: there will be portfolio home bias for values of the elasticity of substitution below one.

A contrario, in Coeurdacier (2005), domestic consumers choose to hold foreign assets to insure against shocks to domestic consumption, provided international relative prices (the terms of trade) respond strongly enough in response to shocks in the relative quantities produced. This happens for values of the elasticity above one. 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. Rather, we stress the fact our estimated aggregate elasticity is substantially above one has far-reaching implications in models of international portfolio choice.

Empirically, trade partners display highly correlated business cycles. Burstein, Kurz and Tesar (forthcoming) show this is quantitatively compatible with a model with production sharing. A positive domestic technology shock diffuses internationally because it increases the demand for foreign inputs, so that both trade and cycle correlations are high simultaneously. A key ingredient is the strong complementarity between the domestic and foreign goods involved in domestic production. Vertical linkages in production are absent from our analysis. Still, our universally high elasticity estimates suggest such strong complementarities are absent in sectoral data (though they may well prevail in highly specific within-firm trade flows). This is part of what gives rise to a co-movement puzzle in Kose and Yi (2006).

Atkeson and Burstein (forthcoming) propose to explain observed deviations from PPP in a model with trade costs, imperfect competition and variable markups. Variable profit margins require that the elasticity of substitution between the (foreign and domestic) varieties that form a sector be larger than the substitutability across sectors. In their model, markups then depend on market shares, and can therefore

¹⁰ We are grateful to Cedric Tille for graciously giving us the simulation code.

fluctuate over time. In their calibration, the elasticity equals 10; deviations from PPP virtually disappear for an alternative value set at 3. Mirdigan (forthcoming) uses disaggregated data from Eastern European countries to conclude menu costs are insufficient to explain the pattern of international relative prices. High trade costs or low elasticities of substitution are also needed. On average at least, our results suggest trade costs are of the essence.

It is perhaps not surprising that models of the real exchange rate should have predictions that depend on the substitutability between domestic and foreign goods. A direct corollary is that the policy consequences of international price differences will also depend crucially on the parameter. Presumably, the relevance of the exchange rate in the monetary policy rule developed in Galí and Monacelli (2005) depends on the substitutability between foreign and domestic varieties. Galí and Monacelli focus on unitary elasticity, so the result is not directly apparent there. But Benigno and De Paoli (2006) introduce a generalization of their model, and their 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 corresponding 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.

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 the crucial weights defined in equation (3). Second, we investigate the importance of “Common Correlated Effects” in obtaining estimates of σ_k , and ultimately of η . 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 3 compares the constrained and unconstrained total import elasticities η using different weighting vectors. The data sources and computations 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. As expected given the empirical values for the correlations between σ_k and n_k , and between σ_k and $n_k \cdot w_k^M$ reported in Table 1, the benchmark weights give rise to the smallest heterogeneity bias. Across the four variants, constrained estimates of η are as low as -1.8, and unconstrained estimates almost reach 5. Variant 3 gives rise to the largest discrepancy, a factor of 2.5. 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 “true”, 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 4 provides a mixed answer, using again our benchmark weights. Without CCE, constrained estimates of η increase slightly, to -2.11, whereas unconstrained estimates decrease slightly, to -3.83. 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 (11) 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 n_k , w_k^M and ω_k to use: we do not observe any direct information on 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 heterogeneity bias, but there is simply no alternative. But we know choosing other values for n_k and w_k^M do 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 continue to be identical, for instance at -2.11 without a CCE corrective term. After all, this is an estimation that constrains all coefficients to be identical, within and between HS6 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 -4.84 for the unconstrained elasticity. Computing standard error bands across this point estimate is not tractable, but the value is virtually identical to Variants 1 and 3 in Table 3, and less than 0.6 away from the others. 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 disagreement. On average, microeconomic data tend to imply substantially higher values than macroeconomic aggregates. The point estimates are also much more heterogeneous. 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 5. This discrepancy validates the conjecture of an aggregation bias in elasticity estimates that goes back at least to Orcutt (1950). Such high parameter values change dramatically the conclusions of calibrated models in areas of international economics as varied as the international transmission of shocks, global imbalances, international risk sharing, portfolio choice and optimal monetary policy.

References

- Atkeson and Burstein, "Pricing to Market, Trade Costs, and International Relative Prices," *American Economic Review*, forthcoming.
- Backus, D., P. Kehoe and F. Kydland (1994), "Dynamics of the Trade Balance and the Terms of Trade: The J-Curve?" *American Economic Review*, 84(1): 84-103.
- Benigno, G. and B. De Paoli (2006), "Optimal Monetary and Fiscal Policy for a Small Open Economy," mimeo.
- Bernard A., J. Eaton, J. B. Jensen and S. Kortum (2003), "Plants and Productivity in International Trade," *American Economic Review*, 93(4): 1268-90.
- Blonigen, B. and W. Wilson (1999), "Explaining Armington: What Determines Substitutability Between Home and Foreign Goods?" *Canadian Journal of Economics*, 32(1): 1-21.
- Broda, C. and D. Weinstein (2006), "Globalization and the Gains from Variety," *The Quarterly Journal of Economics*, 121(2): 541-85.
- Burstein, A., C. Kurz and L. Tesar, "Trade, Production Sharing and the International Transmission of Business Cycles," *Journal of Monetary Economics*, forthcoming.
- Chaney, T. "Distorted Gravity: The Intensive and Extensive Margins of International Trade," *American Economic Review*, forthcoming.
- Chari, V. V., P. Kehoe and E. McGrattan (2002), "Can Sticky Price Models Generate Volatile and Persistent Real Exchange Rates?" *Review of Economic Studies*, 69(3): 533-63.
- Coeurdacier, N. (2005), "Do Trade Costs in Goods Market Lead to Home Bias in Equities?" mimeo.
- Coeurdacier N., R. Kollmann and P. Martin (2007), "International Portfolios with Supply, Demand and Redistributive Shocks," NBER Working Papers No. 13424, Cambridge MA: National Bureau of Economic Research.
- Cole, H. and M. Obstfeld (1991), "Commodity Trade and International Risk Sharing: How Much do Financial Markets Matter?" *Journal of Monetary Economics*, 28(1): 3-24.
- Corsetti, G., L. Dedola and S. Leduc, "International Risk Sharing and the Transmission of Productivity Shock," *Review of Economic Studies*, forthcoming.
- Feenstra, R. (1994a), "New Product Varieties and the Measurement of International Prices," *American Economic Review*, 84(1): 157-77.

- Feenstra, R. (1994b), "New Goods and Index Numbers: U.S. Import Prices," NBER Working Papers No. 3610, Cambridge MA: National Bureau of Economic Research.
- Gagnon, J. (2003), "Productive Capacity, Product Varieties, and the Elasticities Approach to the Trade Balance," International Finance Discussion Papers No. 781.
- Gali, J. and T. Monacelli (2005), "Monetary Policy and Exchange Rate Volatility in a Small Open Economy," *Review of Economic Studies*, 72(3): 707-34.
- Goldstein M. and M. Kahn (1985), "Income and Price Effect in Foreign Trade," in J. Ronald and P. Kennen, eds., *Handbook of International Economics*, North-Holland, Amsterdam: 1042-99.
- 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): 506-21.
- Head, K. and 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): 858-76.
- Heathcote, J. and F. Perri (2008), "The International Diversification Puzzle Is Not As Bad As You Think," mimeo.
- Hooper, P., K. Johnson and J. Marquez (1998), "Trade elasticities for G-7 countries," International Finance Discussion Papers No. 609.
- Houthakker, H. and S. Magee (1969), "Income and Price Elasticities in World Trade," *The Review of Economics and Statistics*, 51(2): 111-25.
- Kahn, M. (1975), "The Structure and Behavior of Imports of Venezuela," *The Review of Economics and Statistics*, 57(2): 221-4.
- Kemp, M. (1962), "Error of Measurements and Bias in Estimates of Import Demand Parameters," *Economic Record*, 38: 369-72.
- Kose, M. and K-M. Yi (2006), "Can the Standard International Business Cycle Model Explain the Relation Between Trade and Comovement?" *Journal of International Economics*, 68(2): 267-95.
- Krugman, P. (1980), "Scale Economies, Product Differentiation, and the Pattern of Trade," *American Economic Review*, 70(5): 950-9.
- Marquez, J. (1990), "Bilateral Trade Elasticities," *The Review of Economics and Statistics*, 72: 70-7.

- Melitz, M. (2003), "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71(6): 1695-725.
- Mirdigan, V., "International Price Dispersion in State-Dependent Pricing Models," *Journal of Monetary Economics*, forthcoming.
- Obstfeld, M. and 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): 117-32.
- Pesaran, H. (2006), "Estimation and Inference in Large Heterogeneous Panels with a Multifactor Error Structure," *Econometrica*, 74(4): 967-1012.
- Romalis, J. (2007), "NAFTA's and CUSFTA's Impact on International Trade," *The Review of Economics and Statistics*, 89(3): 416-35.
- Ruhl, K. (2005), "Solving the Elasticity Puzzle in International Economics," mimeo. segura-cayuela and villarubia
- Shiells, J. (1991), "Errors in Import-Demand Estimates Based upon Unit-Value Indexes," *The Review of Economics and Statistics*, 73(2): 378-82.
- 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): 306-12.

Table 1. Correlations between Elasticities and Weights Computed from Different Data Sources

Weight Matrix	$\text{Corr}(n_k, \hat{\sigma}_k)$	$\text{Corr}(n_k w_k^M, \hat{\sigma}_k)$
Benchmark	0.079	0.048
Variant 1	0.113	0.097
Variant 2	0.112	0.161
Variant 3	0.112	0.160
Variant 4	0.112	0.117

Note: ω_k computed using STAN sectoral interior demand for 1997. Benchmark: n_k and w_k^M using imports and output from IO tables. Variant 1: n_k and w_k^M using imports from BACI and output from STAN. Variant 2: n_k and w_k^M using imports from BACI and output from STAN. Variant 3: n_k using imports from BACI, w_k^M using imports and output from IO tables. Variant 4: n_k using imports from BACI, w_k^M using imports and interior demand from STAN.

Table 2. Estimation with Common Correlated Effects

	Estimate
Constrained total elasticity	-1.916 ^a (0.119)
Constrained partial elasticity	-2.738 ^a (0.048)
Unconstrained total elasticity	-4.210 ^a (0.316)
Unconstrained partial elasticity	-6.006 ^a (0.481)
Number of sectors	56
Number of grid searches	11

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

Table 3. Variants on the Weights

	Total Elasticity	
	Unconstrained	Constrained
Benchmark	-4.210	-1.916
Variant 1	-4.872	-2.083
Variant 2	-4.422	-1.981
Variant 3	-4.987	-1.969
Variant 4	-4.157	-1.814

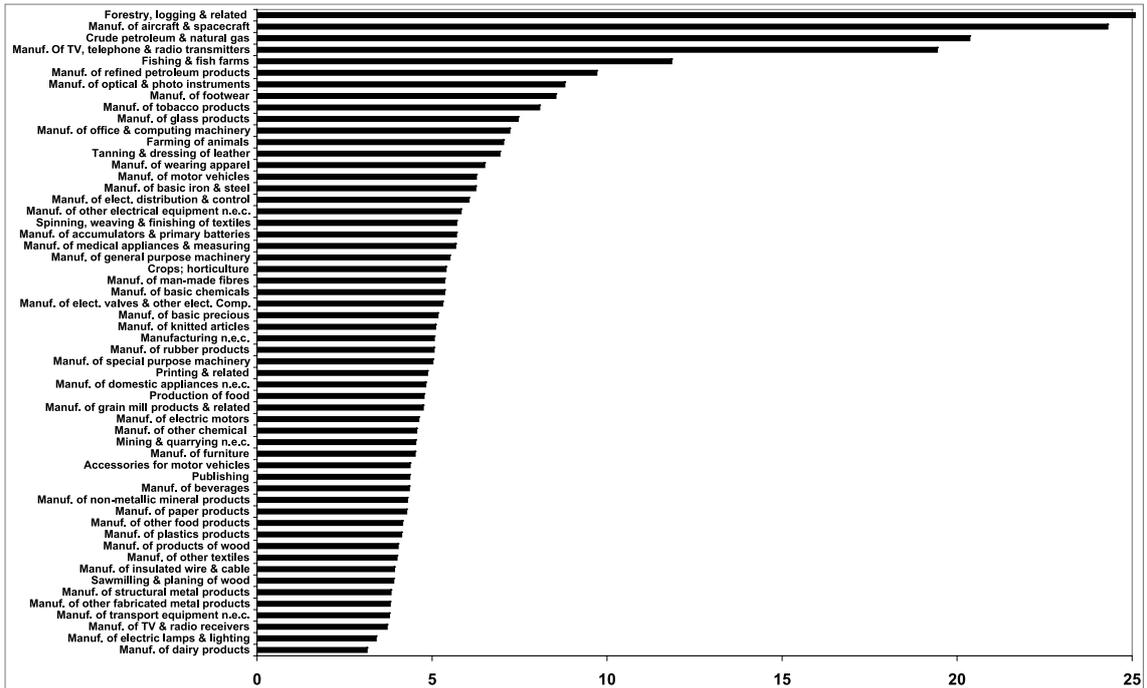
Note: Benchmark: n_k and w_k^M using imports and output from IO tables, ω_k using STAN sectoral interior demand. Variant 1: n_k and w_k^M using imports from BACI and output from STAN, ω_k using STAN sectoral interior demand. Variant 2: n_k and w_k^M using imports from BACI and output from STAN, ω_k using STAN sectoral interior demand. Variant 3: n_k using imports from BACI, w_k^M using imports and output from IO tables, ω_k using STAN sectoral interior demand. Variant 4: n_k using imports from BACI, w_k^M using imports and interior demand from STAN, ω_k using STAN sectoral interior demand.

Table 4. Estimation without Common Correlated Effects

	Estimate
Unconstrained total elasticity	-3.831 ^a (0.122)
Unconstrained partial elasticity	-5.524 ^a (0.206)
Constrained total elasticity	-2.111 ^a (0.027)
Constrained partial elasticity	-3.016 ^a (0.041)
Number of sectors	56
Number of grid searches	12

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

Figure 1. Microeconomic Estimates of the Elasticity of Substitution



Appendix A. Heterogeneous Supply

Here we present the now classic model of production with heterogeneous firms making entry decisions, inspired from Melitz (2003). The theory is written for a given sector k located in country i , and indexes are dropped to facilitate the exposition. All prices expressed in the exporter's currency are denoted with an asterisk *, as in the body of the text. Firms in each sector k and country i are monopolistically competitive, produce differentiated varieties and are heterogeneous in terms of productivity. They decide whether to pay a fixed cost f^e prior to knowing their own productivity, which allows them to export to foreign markets. Once the cost is sunk, productivity φ is drawn from a distribution with probability density function $g(\varphi)$, and an associated cumulated distribution $G(\varphi)$. Productivity is revealed, and the firm decides whether to produce for the foreign market, with a technology featuring constant marginal costs and a fixed overhead export cost f . Both costs are expressed in terms of the unique composite factor of production, whose price is denoted with w^* . Finally, exporting involves an iceberg cost τ , potentially specific to each sector and country. Note this setup departs slightly from the conventional Melitz (2003) model, where the decision to produce for the domestic market predates that to export. For reasons of data availability we abstract from domestic production, and assume production and exporting decisions are effectively disjoint. More precisely, we assume there are fixed costs involved in gathering information on export markets, that are quite separate from those involved in a purely domestic production decision.¹¹

Total costs in exporter's currency are given by

$$TC^*(\varphi) = w^*f + \frac{w^*\tau c(\varphi)}{\varphi}$$

where $c(\varphi)$ is the equilibrium demand addressed to the firm by the foreign market. Individual costs are decreasing in φ , as firms with higher productivity can produce the same amount of output with fewer workers. Firms face a probability δ of exogenous exit in each period. Under these assumptions, the decision to enter export markets and pay the sunk cost f^e is governed by the following maximization of the (probability) discounted value of profits

$$v(\varphi) = \max \left\{ 0, \sum_{t=s}^{\infty} (1-\delta)^{t-s} \pi^*(\varphi) \right\} = \max \left\{ 0, \frac{1}{\delta} \pi^*(\varphi) \right\} \quad (\text{A.1})$$

where profits, expressed in the exporter's currency, are given by

$$\pi^*(\varphi) = \left(p^*(\varphi) - \frac{\tau w^*}{\varphi} \right) c - w^*f$$

¹¹ See for instance Segura-Cayuela and Villarubia (2008).

There is a unique threshold productivity $\bar{\varphi}$ such that $v(\varphi) > 0$ if and only if $\varphi > \bar{\varphi}$. Firms that are productive enough remain in the market after having drawn their productivity and produce in every period. Their export price is set by maximizing profits (A..1) under a demand constraint given in the body of the paper by equation (2), i.e.

$$p^*(\varphi) = \frac{\rho}{\rho - 1} \frac{\tau w^*}{\varphi} \quad (\text{A..2})$$

As Melitz (2003), we introduce a weighted average productivity measure for firms active in the market

$$\tilde{\varphi} = \left[\frac{1}{1 - G(\bar{\varphi})} \int_{\bar{\varphi}}^{\infty} \varphi^{\rho-1} g(\varphi) d\varphi \right]^{\frac{1}{\rho-1}}$$

Then define average profits $\tilde{\pi}^*$:

$$\begin{aligned} \tilde{\pi}^* &\equiv \frac{1}{1 - G(\bar{\varphi})} \int_{\bar{\varphi}}^{\infty} \pi^*(\varphi) g(\varphi) d\varphi \\ &= \frac{1}{1 - G(\bar{\varphi})} \frac{1}{\rho} c \left(\frac{p}{E} \right)^{\rho} \left(\frac{\rho}{\rho - 1} \tau w^* \right)^{1-\rho} \int_{\bar{\varphi}}^{\infty} \varphi^{\rho-1} g(\varphi) d\varphi - w^* f \end{aligned}$$

Since the threshold firm makes no profit, average profits simplify into

$$\tilde{\pi}^* = w^* f \left[\left(\frac{\bar{\varphi}}{\tilde{\varphi}(\bar{\varphi})} \right)^{1-\rho} - 1 \right] \quad (\text{A..3})$$

Equation (A..3) is a “Zero Cutoff Profit” (ZCP) condition. It delineates a relation between average profits $\tilde{\pi}^*$ and the productivity cut-off $\bar{\varphi}$. In particular, when the distribution of productivity follows a Pareto distribution, $G(\varphi) = 1 - \left(\frac{b}{\varphi} \right)^{\kappa}$, the ZCP schedule is flat and given by

$$\tilde{\pi}^* = w^* f \left(\frac{\kappa}{\rho - \kappa + 1} - 1 \right)$$

The ZCP condition characterizes a short run equilibrium, where the threshold productivity $\bar{\varphi}$ is determined in function of the model’s parameters. It is a short run equilibrium in that it takes firms’ location decisions, and thus the potential number of firms active in the market, as given.

A long run equilibrium obtains when a free entry condition holds, that ensures the expected discounted value of profits for a potential entrant equals the fixed cost of entry in the export market. The condition will pin down both threshold productivity and the number of firms active in the export market. Rearranging the condition $\int_0^{\infty} v(\varphi) g(\varphi) d\varphi = f^e$ implies

$$\tilde{\pi}^* = \frac{\delta f^e}{1 - G(\bar{\varphi})}$$

When productivity is distributed according to a Pareto distribution, the Free Entry (FE) condition becomes

$$\tilde{\pi}^* = \delta w^* f^e \left(\frac{\bar{\varphi}}{b} \right)^\kappa$$

The ZCP and FE conditions define a system of two equations in two unknowns, $\tilde{\pi}^*$ and $\bar{\varphi}$. In the particular case of a Pareto distribution, the two conditions give an expression for the productivity cut-off $\bar{\varphi}$ as a function of exogenous parameters:

$$\bar{\varphi} = b \left(\frac{f}{\delta f^e} \right)^{1/\kappa} \left(\frac{\kappa}{\rho - \kappa + 1} - 1 \right)^{1/\kappa} \quad (\text{A..4})$$

Equation (A..4) shows that, once adjustments in the mass of entering firms are accounted for, the productivity cut-off is independent from exchange rate or trade cost shocks. In the long run, firms relocate in response to price changes. When productivity is distributed according to a Pareto-distribution, the threshold productivity $\bar{\varphi}$ remains unchanged. The adjustment along the extensive margin happens exclusively via firm relocation and adjustment in the mass of potential exporters. This differs from Chaney (2008) where the total mass of potential entrants is taken as given.

Finally, the mass M of firms is obtained using the definition of the sectoral price index

$$p = \left[\frac{M}{1 - G(\bar{\varphi})} \int_{\bar{\varphi}}^{\infty} (p^*(\varphi)E)^{1-\rho} g(\varphi) d\varphi \right]^{\frac{1}{1-\rho}}$$

Using optimal pricing, this implies

$$p^{1-\rho} = M \left(\frac{\rho}{\rho - 1} \tau E w^* \right)^{1-\rho} \bar{\varphi}^{\rho-1}$$

Substituting the expression into the definition for average profits implies

$$\tilde{\pi}^* = \frac{1}{\rho} \frac{pc}{EM} - w^* f$$

Using the ZCP condition finally gives

$$M = \frac{\rho - \kappa + 1}{\kappa} \frac{1}{\rho} \frac{1}{w^* f} \frac{pc}{E}$$

Appendix B. Derivation of Supply Curve

Omitting once again indexes for sector k and country i , the definition of the sector price index implies

$$\begin{aligned} p^{1-\rho} &= M \left(\frac{\rho}{\rho-1} \tau w^* E \right)^{1-\rho} \frac{1}{1-G(\bar{\varphi})} \int_{\bar{\varphi}}^{\infty} \varphi^{\rho-1} d\varphi \\ &= M \left(\frac{\rho}{\rho-1} \tau w^* E \right)^{1-\rho} \tilde{\varphi}(\bar{\varphi})^{\rho-1} \end{aligned}$$

Now in the long run we have

$$M = \left(\frac{\bar{\varphi}}{\tilde{\varphi}(\bar{\varphi})} \right)^{\rho-1} \frac{1}{\rho} \frac{1}{w^* f} \frac{pc}{E}$$

Substituting M out gives

$$p = \bar{\varphi}^{\frac{1-\rho}{\rho}} (w^* \rho f E)^{\frac{1}{\rho}} c^{-\frac{1}{\rho}} \left(\frac{\rho}{\rho-1} \tau w^* E \right)^{1-\rho}$$

Now recognize that $\bar{\varphi}$ is constant, and define $\exp(v_{kit}) = (w^* \rho f E)^{\frac{1}{\rho}} \left(\frac{\rho}{\rho-1} w^* \right)^{1-\rho} (\tau E)^{-\rho}$ to obtain the expression in the text. Note that we have assumed constant returns to scale in our theory, but the effective estimation procedure allows for non-constant returns to scale, i.e for non negative values of ω .

Appendix C. Estimated Variances

The variance of $\hat{\sigma}_k$ is computed using the second-order moments of $\hat{\theta}_{1k}$ and $\hat{\theta}_{2k}$ and a first-order approximation of σ_k around its true value:

$$\begin{aligned}\sigma_k &= \hat{\sigma}_k + \left. \frac{\partial \sigma_k}{\partial \theta_{1k}} \right|_{\theta_{1k}=\hat{\theta}_{1k}} (\theta_{1k} - \hat{\theta}_{1k}) + \left. \frac{\partial \sigma_k}{\partial \theta_{2k}} \right|_{\theta_{2k}=\hat{\theta}_{2k}} (\theta_{2k} - \hat{\theta}_{2k}) \\ \Rightarrow \text{Var}(\hat{\sigma}_k) &= \left(\left. \frac{\partial \sigma_k}{\partial \theta_{1k}} \right|_{\theta_{1k}=\hat{\theta}_{1k}} \right)^2 \text{Var}(\hat{\theta}_{1k}) + 2 \left. \frac{\partial \sigma_k}{\partial \theta_{1k}} \right|_{\theta_{1k}=\hat{\theta}_{1k}} \left. \frac{\partial \sigma_k}{\partial \theta_{2k}} \right|_{\theta_{2k}=\hat{\theta}_{2k}} \text{Cov}(\hat{\theta}_{1k}, \hat{\theta}_{2k}) \\ &\quad + \left(\left. \frac{\partial \sigma_k}{\partial \theta_{2k}} \right|_{\theta_{2k}=\hat{\theta}_{2k}} \right)^2 \text{Var}(\hat{\theta}_{2k})\end{aligned}$$

where:

$$\begin{aligned}\left. \frac{\partial \sigma_k}{\partial \theta_{1k}} \right|_{\theta_{1k}=\hat{\theta}_{1k}} &= \frac{1}{\theta_{1k}} \left[1 - \sigma + / - \frac{1}{\sqrt{\theta_{2k}^2 + 4\theta_{1k}}} \right] \\ \left. \frac{\partial \sigma_k}{\partial \theta_{2k}} \right|_{\theta_{2k}=\hat{\theta}_{2k}} &= \frac{1}{2\theta_{1k}} \left[1 + / - \frac{\theta_{2k}}{\sqrt{\theta_{2k}^2 + 4\theta_{1k}}} \right]\end{aligned}$$

Using the same reasoning, the first-order approximation of the aggregate elasticity around its estimated value gives:

$$\eta = \hat{\eta} - \sum_{k \in K} n_k (1 - w_k^M) (\sigma_k - \hat{\sigma}_k) \quad (\text{C..1})$$

The variance of $\hat{\eta}$ is then defined as:

$$\begin{aligned}\text{Var}(\hat{\eta}) &\equiv E(\eta - \hat{\eta})^2 \\ \Rightarrow \text{Var}(\hat{\eta}) &= \sum_{k \in K} n_k^2 (1 - w_k^M)^2 \text{Var}(\hat{\sigma}_k) + \sum_{k \in K} \sum_{k' \neq k} n_k n_{k'} (1 - w_k^M) (1 - w_{k'}^M) \text{Cov}(\hat{\sigma}_k, \hat{\sigma}_{k'})\end{aligned}$$

In the constrained case, the same reasoning gives

$$\text{Var}(\hat{\eta}^c) = \left[\sum_{k \in K} n_k (1 - w_k^M) \right]^2 \text{Var}(\hat{\sigma})$$