Elasticity Optimism*

Jean Imbs [†] Isabelle Méjean[‡]

May 2011

Abstract

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

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

JEL Classification: F41, F32, F21.

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

†Paris School of Economics and CEPR. Corresponding author: Maison des Sciences Economiques, 106-112 bvd de l'Hopital, 75013 Paris France. +33 1 44 07 82 01. jeanimbs@gmail.com. www.jeanimbs.com.

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

1 Introduction

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

Unsurprisingly, the value of the elasticity draws from decades of empirical work. But little consensus has emerged from the effort, except for two broad conclusions. First, elasticity estimates inferred from aggregate data, often in time series, are barely positive. Long time ago, Orcutt (1950) wrote of an “elasticity pessimism”, which he related to such low estimates in the aggregate. Second, there are large differences between goods, with mean estimates typically larger at the sector level than in the aggregate - an “elasticity puzzle” in the terminology introduced by Ruhl (2008).

Here we show an estimate of the elasticity of substitution obtained from macroeconomic data does in fact impose identical elasticities across the sectors that compose the aggregate. Under benign conditions, elasticity estimates based on macroeconomic series, or obtained from disaggregated data constrained to sectoral homogeneity, are actually equal. Inasmuch as sectoral heterogeneity prevails in reality, such a homogeneity constraint creates a classic heterogeneity bias.[1] The bias can drive a difference between macro (i.e. homogeneous) and microeconomic (i.e. heterogeneous) elasticity estimates. We provide a theoretical framework in which the aggregate substitutability is a weighted average of sector-specific elasticities. Within this framework, preserving sectoral heterogeneity, the aggregate substitutability is significantly larger than conventional estimates based on the time dimension of aggregate data. We find a US aggregate elasticity of up to 6 or 7. This is more than double the typical value obtained from aggregate data.

In contrast, when all elasticities are forced to be equal across sectors, the *same* estimator and the *same* data imply an aggregate substitutability for the US between 2.5

[1] A heterogeneity bias arises when homogeneity is forced on coefficient estimates that are in reality heterogeneous. The magnitude of the bias is an empirical matter. See for instance Imbs, Ravn, Rey and Mumtaz (2005).

and 3. Such a range overlaps with conventional estimates arising from aggregate data, as surveyed for instance in Goldstein and Kahn (1985). Estimates from macroeconomic data are typically favored in the calibration of macroeconomic models. For instance, Obstfeld and Rogoff (2005) use a value of 2; Backus, Kehoe and Kydland (1994) use 1.5, as do Chari, Kehoe and McGrattan (2002), and it is set between 0.6 and 2 in Coeurdacier, Kollmann and Martin (2007). Our homogeneous estimates are not significantly different from most of these values. In US data, a significant discrepancy exists therefore between heterogeneous and homogeneous elasticity estimates, which demonstrates the existence of a heterogeneity bias.

The bias matters for calibration. The assumption of a representative sector is a practical modeling shortcut. It remains valid even if sectoral heterogeneity prevails in reality, provided the predictions of a one-sector aggregative model are close to what would be implied by a more sophisticated - but less parsimonious - multi-sector alternative. The heterogeneity bias we document implies such a close link can break down if the aggregative model is calibrated on macroeconomic data. For the US at least, this paper implies the calibration of a one-sector model on macroeconomic data has significantly different predictions from a calibration based on heterogeneous elasticity estimates. The question is which one is closer to what would be implied by a multi-sector model where heterogeneity is modeled explicitly. We investigate this question in a workhorse model in international macroeconomics where a multi-sector extension is tractable, due to Backus, Kehoe and Kydland (1994) [BKK]. We construct a simple multi-sector version of BKK, calibrated with heterogeneous elasticities. We find the one-sector conventional version of BKK must be calibrated to disaggregated, heterogeneous data if it is to mimic its multi-sector, heterogeneous counterpart. Calibration on macroeconomic data ignores heterogeneity, and implies dynamics at odds with the multi-sector model. It is difficult to generalize such reasoning to other models in international economics, for most are not easily amenable to sectoral heterogeneity. The workhorse model in international macroeconomics continues to be aggregative, probably for such tractability reasons. But there is no reason to expect our reasoning does not apply generally. We provide some illustrative discussion spanning several areas of international macroeconomics.[2]

[2] A recent literature has argued elasticities depend in fact only on the distribution of firm heterogeneity, a supply parameter. This happens in Chaney (2008), Eaton and Kortum (2002), Arkolakis, Costinot and Rodriguez-Clare (2009), or Dekle, Eaton and Kortum (2008). Our results suggest a heterogeneity

We argue the elasticity of substitution σ is a preference parameter, whose calibration ought to be motivated by adequate structural estimates. The alternative is to consider σ a free parameter, whose calibration is determined by the model's ability to match particular moments of interest in the data. Most prominently, σ is often chosen in such a way that the model matches the response of aggregate imports to shocks in relative prices. Clearly, the corrected estimate we propose will not perform well along this margin. This should not be surprising. The paper's point is precisely to argue heterogeneity drives a wedge between estimates of the price elasticity of *aggregate* imports and an elasticity of substitution that accounts for sectoral heterogeneity.

Why is an estimation on macroeconomic data equivalent to constraining the sectoral elasticities to homogeneity? The intuition is straightforward. Consider the simplest approach to estimating sectoral elasticities of substitution: a regression of imports M_{kt} in sector k on their relative prices P_{kt}. Suppose the elasticity estimate η is constrained to be identical across sectors. An aggregated version of the sectoral regression can be written as:

$$\sum_k m_k \, \Delta \ln M_{kt} = \eta \, \sum_k m_k \, \Delta \ln P_{kt} + \sum_k m_k \, \epsilon_{kt} \tag{1}$$

where m_k represents the weight of sector k in overall imports. The dependent variable is an import-weighted average of changes in sectoral imports, i.e. the change in aggregate imports. The regressor is an import-weighted average of sectoral price changes, i.e. changes in the measured aggregate price of imports. η estimates the price elasticity of aggregate imports, obtained from aggregate data. It has no reason to equate a weighted average of sectoral estimates, $\sum_k m_k \eta_k$; a heterogeneity bias will exist as soon as it does not. Of course, endogeneity bias can plague equation (1), if the import weights m_k vary over time systematically with sector prices, or if ϵ_{kt} correlates with P_{kt}. But these are issues with the estimation approach, not with aggregation. We show that the argument generalizes to the structural estimator we implement in this paper, borrowed from Feenstra (1994).

How do we aggregate heterogeneous elasticity estimates? We impose specific, nested preferences, where sector consumption is given by a constant elasticity of substitution

bias will also affect the parameter governing the distribution of firms. If distributions are heterogeneous across sectors, an estimate based on the whole panel of firms will suffer from a heterogeneity bias.

(CES) aggregate of varieties produced in different countries, including the domestic economy. The final consumption good, in turn, is given by a Cobb-Douglas aggregation over a continuum of sectors.[3] Crucially, the elasticity of substitution can vary at sector level. We use these preferences to construct a measure of aggregate substitutability consistent with a representative agent choosing between country-level aggregates of domestic and foreign quantities. We characterize the discrepancy between conventional macroeconomic estimates of the elasticity of substitution, imposing equal elasticities across goods, and aggregate estimates allowing for heterogeneity.

The aggregator that we impose at sector level is key. By assumption, the international elasticity of substitution, between domestic and foreign goods, is equal to the substitutability between foreign varieties. This assumption has pervaded international macroeconomics since Armington (1969). It is still used to this day, in BKK for instance, or more recently in Melitz (2003). It is what exonerates us from having to observe domestic production when estimating an international elasticity. We estimate the elasticity between foreign varieties on the basis of data on foreign goods only, using the cross-section of traded quantities and prices across producing countries. By assumption this is identical to the international elasticity.

Even though it is perfectly conventional, such an assumption is restrictive. Alternatives exist that also propose to explain the "elasticity puzzle". For example, Feenstra, Obstfeld and Russ (2010) eschew the key Armington assumption. The substitutability between foreign varieties is left unconstrained, and can be different from the international elasticity. Using data on both traded goods and domestic production, they estimate low values for the international elasticity, and high ones across imported varieties. The discrepancy is akin to the elasticity puzzle. But the generalization comes at a cost: identification is achieved across sectors, and it is difficult to ascertain whether international elasticities are in fact heterogeneous across sectors under their specification. A heterogeneity bias could still be at work, and both explanations complementary.

In fact, there is one set of empirical predictions that help validate the Armington pref-

[3] In earlier versions, we allowed for CES preferences across sectors. Such generalization does not alter our conclusions. The cross-sector elasticity has second order effects on the wedge between heterogeneous and homogeneous estimates of the cross-variety elasticity. The heterogeneity bias remains virtually unaffected.

erences we assume at sectoral level. As we later show, they are best suited to replicate the dispersion in sectoral estimates of the price elasticity of imports. Since Houthakker and Magee (1969), empirical studies have repeatedly established price elasticities of imports display considerable heterogeneity across sectors. Relative to prominent alternatives, and for plausible calibration, our nested preferences are best equipped to mimic such cross-sectional dispersion.

There are alternative explanations for the existence of a discrepancy between the elasticities inferred from microeconomic or macroeconomic data. Ruhl (2008) argues macroeconomic data are observed in time series, and incorporate imperfectly firms entry and exit decisions in export markets. Microeconomic data, on the other hand, are typically in cross section, and focus on the long run. Drozd and Nosal (2010) argue firms price strategically in order to build a match with their customers. Market shares and pricing are persistent, which drives a wedge between short and long run price elasticities of trade. In our paper, the comparison between macro and microeconomic estimates focuses squarely on the importance of heterogeneity. We implement the same estimator, on one data set, with a specific dimension and a specific frequency. The sole difference is whether we impose or not a homogeneity constraint. Alternative explanations based on estimates obtained from different frequencies of the data are effectively held constant.[4]

We next present the model used to guide the aggregation of industry specific elasticities. Section 3 discusses the identification of sector specific parameters, their aggregation and the data involved. Section 4 reports our results, at both micro and macroeconomic levels. We discuss their implication for calibration. We also ascertain their robustness. Section 5 concludes.

2 A Two-Tier CES Demand System

We specify the sector-level demand system that underpins the aggregation of heterogeneous elasticities. The model also maps the international substitutability in goods with

[4]When we perform our estimation on an aggregated version of our *data*, we find estimates for aggregate substitutability that are lower still, near one. This is close to the results obtained with a homogeneity constraint, and still close to conventional macroeconomic estimates. But the comparison is less valuable. Aggregating the data obscures the ceteris paribus nature of our experiment. Aggregation suppresses the sectoral dimension of the data, but unconstrained results are still obtained on the basis of sector-level estimates.

the price elasticity of imports, at both micro and macroeconomic levels. We compare our preferences with prominent alternatives in terms of their ability to reproduce the observed dispersion in estimated price elasticities of imports.

2.1 The Armington Elasticity

In each sector k of country j, consumers demand a continuum of varieties indexed by i, which may be imported or not. Consumption in sector k is given by

$$C_{kj} = \left[\sum_{i \in I} \left(\beta_{kij} \; C_{kij} \right)^{\frac{\sigma_k - 1}{\sigma_k}} \right]^{\frac{\sigma_k}{\sigma_k - 1}}$$

where $i \in I$ indexes varieties, or equivalently producing countries, including j, the variety produced domestically. The parameter β_{kij} lets preferences vary exogenously across varieties. This can reflect for instance differences in quality or home bias in consumption. The sectors that verify $\beta_{kij} = 0$ for all $i \neq j$ are effectively non-traded.

Two assumptions are key. First, the elasticity of substitution σ_k is specific to each industry. The paper revolves around this heterogeneity. Second, substitutability is assumed identical across all varieties, imported or not. By definition, the international substitutability is no different from the elasticity between two imported varieties. The identification scheme revolves around this assumption, which exonerates us from combining information on domestic and imported varieties. We estimate σ_k from the cross-section of imported varieties into the US.

Aggregate consumption in country j combines demand in each sector $k = 1, ..., K$ according to

$$C_j = \prod_{k \in K} \frac{C_{kj}^{\alpha_{kj}}}{\alpha_{kj}^{\alpha_{kj}}}$$

where α_{kj} denotes an exogenous preference parameter. We assume unitary elasticity of substitution between sectors. In earlier versions, we kept the constant elasticity of substitution between sectors unconstrained. Its value has only second-order consequences on the magnitude of the heterogeneity bias.

This structure of demand is classic in international economics. The key assumption for our purposes is equal substitutability between two varieties, no matter their origin.

Introducing the assumption is largely what opened the door to the new trade literature, pioneered by Krugman (1980), and laid the foundation for the more recent models of trade with heterogeneous firms, starting with Melitz (2003).

The representative maximizing agent chooses her consumption allocation on the basis of Cost, Insurance, Freight prices, labeled in local currencies.[5] Utility maximization implies that demand for variety i in each sector k is given by

$$C_{kij} = \beta_{kij}^{\sigma_k - 1} \left(\frac{P_{kij}}{P_{kj}} \right)^{-\sigma_k} \alpha_{kj} \frac{P_j}{P_{kj}} C_j \tag{2}$$

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

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

Second, we focus on uniform shocks to the relative price of imported goods, across all exporters $i \neq j$ 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.[6] A natural candidate is a domestic shock to relative production costs, fully passed-through to relative

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

[6]The 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.

prices. It will change the international price of domestic goods by an identical amount across all sectors k and exporters $i \neq j$. We let φ_j denote such an aggregate shock, to the relative price of all imports of every sector k into country j.

Consider the definition of an aggregate elasticity of substitution σ_j between bundles of domestic and foreign goods in country j. By definition,

$$\sigma_j - 1 = \frac{\partial \ln \sum_k P_{kjj} C_{kjj} - \partial \ln \sum_k \sum_{i \neq j} \ P_{kij} C_{kij}}{\partial \ln \varphi_j}$$

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

The aggregate shock affect the international relative price in each sector k. Using equation (2), simple algebra implies

$$\sigma_j - 1 = \sum_k m_{kj}(\sigma_k - 1) - \sum_k (m_{kj} - d_{kj}) (\sigma_k - 1) \frac{\partial \ln P_{kj}/P_{kjj}}{\partial \ln \varphi_j}$$

i.e.,

$$\sigma_j = \sum_k m_{kj} \sigma_k - \sum_k (m_{kj} - d_{kj}) (\sigma_k - 1) \lambda_j^k \tag{3}$$

with $m_{kj} = \frac{\sum_{i \neq j} P_{kij} \ C_{kij}}{\sum_{k \in K} \sum_{i \neq j} P_{kij} \ C_{kij}}$ the share of sector k in the total consumption of foreign goods, $d_{kj} \equiv \frac{P_{kjj} \ C_{kjj}}{\sum_{k \in K} P_{kjj} \ C_{kjj}}$ the share of sector k in the consumption of domestic goods, and $\lambda_j^k = \frac{\sum_{i \neq j} P_{kij} \ C_{kij}}{P_{kj} \ C_{kj}}$ the share of imports in total expenditures on good k.

The aggregate elasticity σ_j contains two terms: an import-weighted average of industry-specific elasticities, and second-round responses of industry specific price indices P_{kj} to the aggregate shock. Since by assumption the relative price of good k changes identically across all source economies $i \neq j$, the composition of the ideal price index in sector k changes significantly in response to the shock considered. Taking inspiration from Orcutt's (1950) terminology, σ_j is a total elasticity, one that takes the response of price indices into account. A partial elasticity assumes aggregate price responses away.[7] We note the second term in equation (3) is small. By definition of d_{kj}, $\sum_k (m_{kj} - d_{kj})(\sigma_k - 1)\lambda_j^k < \sum_k m_{kj}\sigma_k\lambda_j^k$. The difference between partial and total elasticities is bounded above. Relative to the partial elasticity, the upper bound contains an extra multiplicative term $\lambda_j^k < 1$ for all k. Partial and total elasticities will differ by small amounts. They will in fact be virtually identical if in addition sector allocations of expenditures are similar for domestic and foreign goods, i.e. whenever $m_{kj} \simeq d_{kj}$.

Focus for now on the partial elasticity. Equation (3) illustrates the possibility of a classical heterogeneity bias. Consider an estimate of σ_k constrained to homogeneity across sectors, $\bar{\sigma}$. A discrepancy can exist between the partial elasticity defined in equation (3), and a constrained estimate of the partial elasticity, which happens to equal $\bar{\sigma}$ since $\sum_k m_{kj} = 1$. Such discrepancy corresponds to a heterogeneity bias. Its direction and magnitude increase with the correlation between m_{kj} and σ_k across sectors. If import shares increase with goods substitutability, a heterogeneous estimate of the partial elasticity takes larger values than $\bar{\sigma}$. The reasoning extends readily to total elasticity σ_j, because the second term in equation (3) is at least one order of magnitude smaller than the first.

2.2 The Price Elasticity of Imports

In most of the literature, the elasticity of substitution is inferred from the price elasticity of imports, at various levels of aggregation. It is therefore important that we verify a heterogeneity bias also prevails in estimates of these trade elasticities. By definition, the price elasticity of aggregate imports in country j is given by

$$\eta_j = \frac{\partial \ln \sum_k \sum_{i \neq j} P_{kij} C_{kij}}{\partial \ln \varphi_j}$$

Under two-tier CES preferences, demand is characterized by equation (2), and the price elasticity of imports is given by

$$\eta_j = 1 - \sum_k m_{kj}\sigma_k + \sum_k m_{kj}(\sigma_k - 1)\lambda_j^k \tag{4}$$

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

where the second summation is once again smaller in magnitude. A positive heterogeneity bias in estimates of the elasticity of substitution means $\sum_k m_{kj}\sigma_k > \bar{\sigma}$. Equation (4) suggests the intuition carries through to the price elasticity of imports: it will take lower, negative values whenever a positive heterogeneity bias plagues estimates of σ_j. If imported goods tend to be substitutes, a price elasticity of imports allowing for heterogeneity in σ_k takes larger values (in absolute value) than one based on $\bar{\sigma}$. This can explain the elasticity puzzle, provided $\bar{\sigma}$ is in fact the estimate arising from macroeconomic data.

Equation (4) maps the price elasticity of imports with a preference parameter. When supply decisions are modeled explicitly, a recent literature has argued trade elasticities only depend on a supply parameter. Chaney (2008), Eaton and Kortum (2002), Arkolakis, Costinot and Rodriguez-Clare (2009), or Dekle, Eaton and Kortum (2008) all demonstrate the price elasticity of imports is determined by the distribution of firm productivity, typically the exponent of a Pareto distribution. The heterogeneity bias we document here can in fact be interpreted in the context of such alternative theory. Suppose the distributions of firm productivity are different across sectors, with heterogeneous Pareto exponents. Then estimating a single Pareto parameter, on the basis of the total universe of firms pooled across sectors, gives different results from a weighted-average of sector-specific estimates. Such heterogeneity bias may have important consequences on the calibration of multi-sector models with firm heterogeneity.[8]

2.3 Alternative Preferences

The specific two-tier preferences we assume are key to this paper. The international elasticity of substitution is sector-specific, and can be estimated with parsimonious data requirements. But they are restrictive. For instance, as clearly shown in equation (3), the aggregate elasticity of substitution is not a deep parameter that enters utility directly. It is a weighted average of sector level utility parameters. This may in fact explain the well known cross-country dispersion in trade elasticities, documented for instance in Houthakker and Magee (1969). In our preferences, such dispersion would come from

[8] In an earlier version, we solved a slightly generalized version of Chaney (2008), with firm heterogeneity. We allowed for the elasticity of substitution across countries to differ from the elasticity across firms. In such a case, we showed the price elasticity of imports continued to depend on preference parameters. A heterogeneity bias in σ continued to occur under the conditions described in this paper. In fact, the heterogeneity bias in this paper is a lower bound to its counterpart in such a model with endogenous firm entry.

international differences in the specialization of trade. But alternatives exist, that can accommodate explicitly some sectoral heterogeneity, while preserving the existence of an aggregate elasticity of substitution in preferences.

Consider the following simple example. Consumption in country j is a CES aggregate of domestic and foreign goods, given by

$$C_j = \left[\left(C_j^D \right)^{\frac{\sigma-1}{\sigma}} + \left(C_j^F \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \tag{5}$$

where the foreign bundle of goods C_j^F is itself a CES aggregate of imported varieties, defined as

$$C_j^F = \prod_k \frac{1}{\alpha_{kj}^{\alpha_{kj}}} \left[\left(\sum_{i \neq j} \left(C_{kij} \right)^{\frac{\rho_k - 1}{\rho_k}} \right)^{\frac{\rho_k}{\rho_k - 1}} \right]^{\alpha_k} \tag{6}$$

The elasticity of substitution between imported varieties ρ_k is sector specific. But it is different from the elasticity of substitution between domestic and foreign goods σ, which is explicitly macroeconomic in nature. Such utility can accommodate sectoral heterogeneity in the elasticity of substitution between imported varieties, which under our assumptions pins down the international elasticity. But here the international elasticity is not related with ρ_k, nor indeed with its weighted average. It is macroeconomic in nature.

Clearly, in such alternative preferences, the international elasticity of substitution cannot suffer from any heterogeneity bias. We argue however it misses out on one important empirical dimension. The sector level price elasticity of imports implied by equations (5) and (6) is given by

$$\begin{aligned} \eta_{kj} &= \frac{\partial \ln \sum_{i \neq j} P_{kij} C_{kij}}{\partial \ln \varphi_j} \\ &= (1 - \sigma)(1 - \lambda_j) m_{kj} \end{aligned}$$

where $\lambda_j = \frac{\sum_k \sum_{i \neq j} P_{kij} \; C_{kij}}{\sum_k P_{kj} \; C_{kj}}$ is the share of imports in total consumption. Under such preferences, any dispersion in the price elasticity of imports at the sector level has to originate from import weights. The elasticity of substitution in sector k does not matter, since it is σ everywhere.

There is little doubt that price elasticities of imports vary across sectors. In the US, Houthakker and Magee (1969) find values of η_{kj} between 0 and -4.05, for finished manufactures. Shiells (1991) finds values down to -6 for "tires and tubes" and "toys and games". Kahn (1975) identifies a similar range in Venezuela. The generalized two-tier preferences summarized in equation (5) and (6) imply sectors with price elastic imports are just ones with a large share of total imports m_{kj}. A casual reading of a venerable empirical literature suggests m_{kj} is not a sufficient statistic for the ranking of estimated η_{kj}.

Contrast this result with our preferences, which imply an elasticity given by

$$\eta_{kj} = (1 - \sigma_k)(1 - \lambda_j^k) \tag{7}$$

The price elasticity of imports will now be larger in absolute value for highly substitutable goods.[9] When we use equation (7) to compute the values of η_k implied for the US by our sectoral estimates of σ_k, we find a ranking that maps well with conventional estimates of the sectoral price elasticity of imports.

Feenstra, Obstfeld and Russ (2010) [FOR] generalize equations (5) and (6) further, and let σ vary across sectors. The elasticity across imported varieties continues to differ from the international elasticity, and both are allowed to vary across sectors. With such general preferences, the price elasticity of imports at the sectoral level is identical to ours, and will vary with σ_k, as is plausible and intuitive. But such generality comes at a cost. FOR identify ρ_k across exporting countries, but σ_k in the cross-section of sectors.[10] As a result, it is difficult to obtain sector-level estimates of σ_k. FOR have data on a total of 113 goods, which they partition into 10 sectors. They use the variation across goods within each one of these 10 sectors to identify separate elasticity estimates. They find little heterogeneity in σ_k across 10 sectors. Of course, inasmuch as each sector is itself an aggregate constrained to homogeneity, there is a possibility these apparently homogeneous sectoral estimates reflect a heterogeneity bias. Given the identification scheme, little more can be said, and FOR's explanation to the "elasticity puzzle" is likely to remain complementary to ours. More to the point however, if the preferences in FOR

[9] If the price shifter φ_j represents a shock to import prices relative to an aggregate index (CPI, e.g.), we find the conventional result that $\eta_{kj} = 1 - \sigma_k$. The main text considers shocks to the price of imports relative to domestic price indices, both at the sectoral level.

[10] They also need data on domestic production, which have to be mapped with the level of disaggregation observed in trade data.

imply no sectoral heterogeneity in σ_k, then it is difficult to rationalize the well known cross-sector heterogeneity in estimates of η_k.

3 Identification

We adapt the methodology in Feenstra (1994) to identify the values of σ_k across sectors, which are then used to obtain aggregate elasticity estimates σ_j. Identification is structural, but requires a CES demand system with constant markups. We first discuss the econometrics involved in estimating σ_k for all sectors k in the US economy. We emphasize how we accommodate common effects across all sectors and measurement error. We then turn to the estimation of $\bar{\sigma}$, a measure of elasticity constrained to be identical across sectors. We show under what conditions $\bar{\sigma}$ would also be implied by macroeconomic data. We close with a description of our data.

3.1 Microeconomic Estimates

Thanks to the Armington assumption, σ_k is *both* the elasticity across imported varieties *and* the international elasticity. We can estimate it with no information on domestic prices, just with data on imported quantities and prices across origin countries. Rearranging equation (2), demand can be rewritten

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

where t is a time index.[11] Feenstra (1994) or Broda and Weinstein (2006) impose a simple supply structure, with prices fixed in local currency and inclusive of trade costs

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

where υ_{kit} denotes a technological shock varying across sectors and exporters, τ_{kit} is a trading cost and ω_k is the inverse of the price elasticity of supply in sector k.[12] The

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

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

potential aggregate effects of the nominal exchange rate, for instance, are soaked up by the shock υ_{kit}. In our estimation, it will be important to control for any such common effects across sectors.

Following Kemp (1962), expenditure shares are used to alleviate measurement error in unit values. Define $s_{kit} = \frac{P_{kit}C_{kit}}{P_{kt}C_{kt}}$, the share of expenditures in good k imported from country i. Prices are measured Free on Board, i.e. $\tilde{P}_{kit} \equiv P_{kit}/\tau_{kit}$, where a tilde denotes observed variables. As we do not observe domestically produced consumption, empirical market shares are given by $\tilde{s}_{kit} \equiv \frac{\tilde{P}_{kit}C_{kit}}{\sum_{i \neq US} \tilde{P}_{kit}C_{kit}} = \frac{s_{kit}}{\tau_{kit}} \left(1 + \frac{P_{kUSt}C_{kUSt}}{\sum_{i \neq US} \tilde{P}_{kit}C_{kit}}\right) \equiv \frac{s_{kit}}{\tau_{kit}}\mu_{kt}$. Taking logarithms, it is straightforward to rewrite demand as

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

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 capturing trade cost and taste shocks.

Substituting equation (8) in log-linearized supply yields

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

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. $\delta_{kit} \equiv \frac{1}{1+\omega_k\sigma_k}\Delta\upsilon_{kit}$ is an error term encapsulating movements in the exchange rate or aggregate technological developments in country i and sector k.

Under standard assumptions on fundamental shocks, it is possible to identify the system formed by equations (8) and (9). Identification rests on the cross-section of exporters i to the US, and is achieved in relative terms with respect to a reference country r.[13] Feenstra (1994) summarizes the information contained in the system with the following estimable regression

$$Y_{kit} = \theta_{1k} X_{1kit} + \theta_{2k} X_{2kit} + u_{kit} \tag{10}$$

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

where $Y_{kit} = (\Delta \ln \tilde{P}_{kit} - \Delta \ln \tilde{P}_{krt})^2$, $X_{1kit} = (\Delta \ln \tilde{s}_{kit} - \Delta \ln \tilde{s}_{krt})^2$, $X_{2kit} = (\Delta \ln \tilde{s}_{kit} - \Delta \ln \tilde{s}_{krt})(\Delta \ln \tilde{P}_{kit} - \Delta \ln \tilde{P}_{krt})$ and $u_{kit} = (\varepsilon_{kit} - \varepsilon_{krt})(\delta_{kit} - \delta_{krt}) \frac{(\sigma_k - 1)(1+\omega_k)}{1+\omega_k \sigma_k}$. Estimates of equation (10) map directly with the parameters of interest, since

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

Equation (10) still suffers from endogeneity. Under Armington preferences, Feenstra (1994) shows the time averages of the regressors are valid instruments. A putative correlation between u_{kit} and the regressors in equation (10) washes out in time averages. As in Feenstra, identification is therefore based on the cross-sectional dimension of equation (10). We include an intercept to account for the measurement error arising from using unit values to approximate prices. Given the origin of potential measurement error, we let it prevail at the most granular level afforded by our data.

The system summarized by equation (10) can accommodate developments that are specific to each sector k. But some of our estimates are based on data pooled across sectors, so it is important to allow for more general, aggregate influences in the specification. Aggregate technology shocks for instance, or movements in the nominal exchange rate, presumably affect prices and quantities jointly in all sectors. If it were a shock in the exporting economy, that would correspond to a common component of υ_{kit} across all k. We allow for such correlated effects in as general and parsimonious a manner as possible. We implement a correction suggested by Pesaran (2006) to purge all "Common Correlated Effects" (CCE) from sector level data. We estimate

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

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 time-varying components that are common across all sectors. Following Pesaran (2006), X_{1it} and X_{2it} are the cross-sector arithmetic averages of X_{1kit} and X_{2kit}.

Armed with consistent (and sector-specific) estimates of θ_{1k} and θ_{2k}, it is straightforward to infer elasticities. 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 A 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 $\hat{\sigma}_k$. We follow Broda and Weinstein (2006), and use a search algorithm that minimizes the sum of squared residuals in equation (11) over the intervals of admissible values for the supply and demand elasticities.[14]

3.2 Homogeneous and Macroeconomic Estimates

Estimates of $\bar{\sigma}$, constrained to homogeneity across sectors, are obtained from a modification of equation (11). The sectoral dimension in the estimated coefficients is assumed away, which implies:

$$Y_{kit} = \theta_0 + \theta_1 \hat{X}_{1ki} + \theta_2 \hat{X}_{2ki} + \theta_3 X_{1it} + \theta_4 X_{2it} + u_{kit} \tag{12}$$

Aside from the homogeneity constraint, equation (12) is identical to its heterogeneous counterpart. We maintain the assumption of a HS6-specific intercept, to accommodate measurement error. We continue to allow for the possibility that aggregate shocks in any country i should affect all sectors simultaneously, and include corrective CCE terms. The instrumentation and correction for heteroskedasticity are also identical. Identification continues to rest on the cross-section of exporters i. Equation (12) is now estimated on the pooled dataset formed by observations on all sectors. The only difference between regressions (11) and (12) pertains to the homogeneity constraint. The data, the dimension

[14] We use this approach whenever direct estimates of θ_{1k} and θ_{2k} cannot be used to infer $\hat{\sigma}_k$. Whenever CCE are included, we hold constant the estimates of θ_{3k} and θ_{4k} obtained from the direct instrumental variable regression, and search the combination of values for σ_k and ω_k that minimizes the sum of squared residuals in equation (11). The corresponding standard errors are obtained via bootstrapping of the procedure using 1,000 repetitions.

used in identification, the corrections we implement, are all held constant.[15]

Why would an estimate of $\bar{\sigma}$ reproduce what aggregated data imply? To see under what conditions a mapping exists, consider a version of equation (12) aggregated up across sectors using specific weights:

$$\sum_k m_{kit-1}^2 Y_{kit} = \sum_k m_{kit-1}^2 \theta_0 + \theta_1 \sum_k m_{kit-1}^2 X_{1kit} + \theta_2 \sum_k m_{kit-1}^2 X_{2kit} + \sum_k m_{kit-1}^2 u_{kit} \tag{13}$$

where $m_{kit} = \frac{\tilde{P}_{kit}\ C_{kit}}{\sum_k \tilde{P}_{kit}\ C_{kit}}$ is the share of sector k in total imports from country i. For simplicity we omit the CCE corrective terms and focus on the un-instrumented regressors.

Consider now a version of the Feenstra (1994) estimator implemented on aggregate data. By definition, the dependent variable is given by

$$Y_{it} = (\Delta \ln \tilde{P}_{it} - \Delta \ln \tilde{P}_{rt})^2$$

with $\Delta \ln \tilde{P}_{it} = \sum_k m_{kit-1}\ \Delta \ln \tilde{P}_{kit}$ the measured Laspeyres aggregate price index for imports originating from country i. Consider the dependent variable in equation (13):

$$\sum_k m_{kit-1}^2 Y_{kit} = \sum_k m_{kit-1}^2 (\Delta \ln \tilde{P}_{kit} - \Delta \ln \tilde{P}_{krt})^2$$

With trade weights m_{kit-1} approximately equal in countries i and r, the reference exporter, simple algebra implies

$$\sum_k m_{kit-1}^2 Y_{kit} = Y_{it} - \Lambda_{irt}^Y$$

with

$$\begin{aligned}\Lambda_{irt}^Y \equiv\ & \sum_k \sum_{k' \neq k} m_{kit-1} m_{k'it-1} \Delta \ln \tilde{P}_{kit} \Delta \ln \tilde{P}_{k'it} + \sum_k \sum_{k' \neq k} m_{kit-1} m_{k'it-1} \Delta \ln \tilde{P}_{krt} \Delta \ln \tilde{P}_{k'rt} \\ & -2 \sum_k \sum_{k' \neq k} m_{kit-1} m_{k'it-1} \Delta \ln \tilde{P}_{kit} \Delta \ln \tilde{P}_{k'rt}\end{aligned}$$

[15]Note this implies the existence of an additional heterogeneity bias, that pertains to the estimation of the θs. If, as we believe, θ_1 to θ_4 are in fact heterogeneous across sectors, equation (12) suffers from a classic heterogeneity bias. Constrained estimates of θ_1 to θ_4 may well be different from a weighted average of their sector-specific values. This can explain why our estimate of $\bar{\sigma}$ lies outside of the interval of estimated σ_k.

a term that contains import-weighted averages of cross-sector covariances in prices.

Consider again a version of the Feenstra (1994) estimator implemented on aggregate data. The first regressor is defined as

$$X_{1it} = (\Delta \ln \tilde{s}_{it} - \Delta \ln \tilde{s}_{rt})^2$$

with $\tilde{s}_{it} \equiv \frac{\sum_k \tilde{P}_{kit} C_{kit}}{\sum_k \sum_{i \neq US} \tilde{P}_{kit} C_{kit}} = \sum_k m_{kt}\ \tilde{s}_{kit}$. The share of sector k in US imports is defined as before, $m_{kt} = \frac{\sum_{i \neq US} \tilde{P}_{kit} C_{kit}}{\sum_k \sum_{i \neq US} \tilde{P}_{kit} C_{kit}}$. For small enough growth rates in market shares, approximate

$$\begin{aligned}
\Delta \ln \tilde{s}_{it} &\simeq \frac{\sum_k m_{kt}\ \tilde{s}_{kit}}{\sum_k m_{kt-1}\ \tilde{s}_{kit-1}} - 1 \\
&= \sum_k \frac{m_{kt}}{m_{kt-1}} \left(\frac{\tilde{s}_{kit}}{\tilde{s}_{kit-1}} - 1 \right) \frac{m_{kt-1} \tilde{s}_{kit-1}}{\sum_k m_{kt-1} \tilde{s}_{kit-1}} + \sum_k \left(\frac{m_{kt}}{m_{kt-1}} - 1 \right) \frac{m_{kt-1} \tilde{s}_{kit-1}}{\sum_k m_{kt-1} \tilde{s}_{kit-1}} \\
&= \sum_k m_{kit-1} \frac{m_{kt}}{m_{kt-1}} \Delta \ln \tilde{s}_{kit} + \sum_k m_{kit-1} \Delta \ln m_{kt}
\end{aligned}$$

If the sectoral composition of overall US imports remains approximately stable over time, we have

$$\Delta \ln \tilde{s}_{it} \simeq \sum_k m_{kit-1}\ \Delta \ln \tilde{s}_{kit}$$

Then by analogy with the dependent variable, rearrange the first regressor in equation (13) to obtain

$$\begin{aligned}
\sum_k m_{kit-1}^2 X_{1kit} &= \sum_k m_{kit-1}^2 (\Delta \ln \tilde{s}_{kit} - \Delta \ln \tilde{s}_{krt})^2 \\
&= X_{1it} - \Lambda_{irt}^{X1}
\end{aligned}$$

where we assumed again that $m_{kit-1} \simeq m_{krt-1}$. The term Λ_{irt}^{X1} represents import-weighted averages of cross-sector covariances in market shares, defined as

$$\begin{aligned}
\Lambda_{irt}^{X1} \equiv & \sum_k \sum_{k' \neq k} m_{kit-1} m_{k'it-1} \Delta \ln \tilde{s}_{kit} \Delta \ln \tilde{s}_{k'it} + \sum_k \sum_{k' \neq k} m_{kit-1} m_{k'it-1} \Delta \ln \tilde{s}_{krt} \Delta \ln \tilde{s}_{k'rt} \\
& -2 \sum_k \sum_{k' \neq k} m_{kit-1} m_{k'it-1} \Delta \ln \tilde{s}_{kit} \Delta \ln \tilde{s}_{k'rt}
\end{aligned}$$

Finally, following identical reasoning, we have

$$\begin{aligned}\sum_k m_{kit-1}^2 X_{2kit} &= \sum_k m_{kit-1}^2 (\Delta \ln \tilde{s}_{kit} - \Delta \ln \tilde{s}_{krt})(\Delta \ln \tilde{P}_{kit} - \Delta \ln \tilde{P}_{krt}) \\ &= X_{2it} - \Lambda_{irt}^{X2}\end{aligned}$$

where now

$$\begin{aligned}\Lambda_{irt}^{X2} \equiv & \sum_k \sum_{k' \neq k} m_{kit-1} m_{k'it-1} \Delta \ln \tilde{P}_{kit} \Delta \ln \tilde{s}_{k'it} + \sum_k \sum_{k' \neq k} m_{kit-1} m_{k'it-1} \Delta \ln \tilde{P}_{krt} \Delta \ln \tilde{s}_{k'rt} \\ & - \sum_k \sum_{k' \neq k} m_{kit-1} m_{k'it-1} \Delta \ln \tilde{P}_{kit} \Delta \ln \tilde{s}_{k'rt} - \sum_k \sum_{k' \neq k} m_{kit-1} m_{k'it-1} \Delta \ln \tilde{P}_{krt} \Delta \ln \tilde{s}_{k'it}\end{aligned}$$

The corrective term involves import-weighted averages of cross-sector covariances between market shares and prices.

An estimation of Feenstra (1994) performed on pooled sectoral data, with coefficients (θ_1, θ_2) constrained to homogeneity across sectors implies therefore:

$$Y_{it} = \theta_0' + \theta_1 X_{1it} + \theta_2 X_{2it} + v_{it} \tag{14}$$

where $\theta_0' = \sum_k m_{kit-1}^2 \theta_0$ and $v_{it} = \sum_k m_{kit-1}^2 u_{kit} + \Lambda_{irt}^Y - \theta_1 \Lambda_{irt}^{X1} - \theta_2 \Lambda_{irt}^{X2}$. The implication holds under the approximations that (a) the sectoral compositions of imports in countries i and r are similar, (b) the overall sectoral composition of US imports changes little year by year, and (c) the growth rate in sectoral market shares remains reasonably close to zero. By definition, equation (14) is equivalent to an estimation of θ_1, θ_2 and ultimately σ that is performed on aggregated data. By construction, we know the estimated coefficients are precisely what is implied by a constrained estimation on sectoral data. In other words, $\bar{\sigma}$, an elasticity of substitution estimated on sectoral data and constrained to homogeneity is precisely what macroeconomic data would imply.

The precision of such a mapping depends on whether v_{it} is well behaved in equation (14). In addition to approximations (a), (b), and (c), it is possible that v_{it} is in fact correlated with the regressors in equation (14). The inclusion of CCE terms in the sectoral, constrained estimation begins to limit this possibility. CCE terms purge the sectoral data from any cross-sector covariances in prices and market shares: they set $\Lambda_{irt}^Y =$

$\Lambda_{irt}^{X1} = \Lambda_{irt}^{X2} = 0$. Thus they bring constrained sectoral estimates closer to conventional ones obtained from macroeconomic data.

Still, endogeneity will survive in equation (14) if m_{kit-1} drives a systematic correlation between v_{it} and X_{1it} or X_{2it}. The instrumentation of both variables on the basis of their time averages, suggested by Feenstra (1994), takes care of systematically correlated dynamics between v_{it} and X_{1it} or X_{2it}. But a correlation can survive in the cross-section of exporting countries, indexed by i. This happens if values of m_{kit-1} correlate systematically with realizations of $\hat{X}_{1i} = \frac{1}{T}\sum_t(\sum_k m_{kit-1}\ \Delta \ln \tilde{s}_{kit} - \sum_k m_{krt-1}\ \Delta \ln \tilde{s}_{krt})^2$ or of $\hat{X}_{2i} = \frac{1}{T}\sum_t(\sum_k m_{kit-1}\ \Delta \ln \tilde{s}_{kit} - \sum_k m_{krt-1}\ \Delta \ln \tilde{s}_{krt})(\sum_k m_{kit-1}\ \Delta \ln \tilde{P}_{kit} - \sum_k m_{krt-1}\ \Delta \ln \tilde{P}_{krt})$. Under assumption (a), such outcome is unlikely.[16]

3.3 Data

We 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. Bilateral trade flows are reported at the sectoral level, building on the United Nations ComTrade database. Additional effort went into the harmonization of trade flows on the basis of both import and export declarations. The improvement limits measurement error.

In the main body of the text, we do not estimate elasticities at HS6 level. Instead, we partition our data into 56 ISIC (Revision 3) industries. Such partition is entirely innocuous for the constrained estimation in equation (12), since homogeneity is imposed across all HS6 sectors, whether or not they belong to the same ISIC category. But the partition becomes important for the unconstrained estimation. It is performed for lack of detailed information on m_{kj}, d_{kj} and λ_j^k at HS6 level. It effectively assumes that all HS6 goods are equally substitutable within an ISIC category, but not between. The assumption can create a heterogeneity bias. We conjecture that the heterogeneity

[16] The dimension of the data used to identify the coefficients of interest is of course different between (constrained) sectoral data and macroeconomic series. This is apparent from the intercept in equation (14), which varies across exporting countries i. The intercept in sectoral data is specific to each HS6 category. In other words, aggregating the data has other implications, e.g. in terms of power, than just imposing a homogeneity constraint.

between ISIC industries is largest, and thus induces the bulk of a heterogeneity bias. We do however perform some robustness in section 4.4, estimating elasticities for all HS6 goods. But then the ISIC weights used in aggregation are assumed to apply for all HS6 goods in each category.

To estimate equations (11) and (12), we only need measures of $\tilde{P}_{kit}$ and $\tilde{s}_{kit}$. 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.[17] Quantities are in tons. The empirical model described in section 3.1 is not sensitive to the currency denomination of trade data, nor to the treatment of trade costs, as both are passed into the residuals. Expenditure shares are measured as $\tilde{s}_{kit} = \frac{\tilde{P}_{kit} C_{kit}}{\sum_{i \neq j} \tilde{P}_{kit} C_{kit}}$.

We subject our data to sampling with a view to limiting the role of extreme outliers. They are notoriously frequent in approaches making use of unit values to approximate prices. Tonnage is not always appropriate to capture traded volumes, which can artificially instill massive volatility in the resulting unit values. In each sector, we exclude annual variations in prices and market shares that exceed five times the median value. Since the cross-section of exporters is what ultimately achieves identification, we also impose a minimum of 20 exporters for each HS6 good over the whole observed time period. Our data ultimately represent 77 percent of the total value of US imports, across 56 ISIC sectors.

In the model, the two relevant weights can be rewritten $m_{kj} = \frac{w_{kj} \ \lambda_j^k}{\sum_k w_{kj} \ \lambda_k^j}$ and $d_{kj} = \frac{w_{kj} \ \left(1-\lambda_j^k\right)}{\sum_k w_{kj} \ \left(1-\lambda_j^k\right)}$, where we have defined $w_{kj} \equiv \frac{P_{kj} C_{kj}}{\sum_k P_{kj} C_{kj}}$, the share of sector k in aggregate expenditures. Calibration is therefore only needed for w_{kj} and λ_j^k. In the main body of the text, the expenditure shares w_{kj} are obtained from the OECD STAN dataset. We compute the 1997 ratio of sectoral absorption (value added and imports net of exports) relative to its aggregate across sectors. The import shares λ_j^k are obtained from the US input/output (IO) tables, available in the ISIC (Revision 3) nomenclature. We compute the 1997 ratio of imports over domestic gross output. Values for m_{kj} and d_{kj} are calculated

[17] 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.

accordingly.[18]

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

4 Results and Implications for Calibration

The first sub-section reviews elasticity estimates allowing for heterogeneity. We discuss the estimates of $\hat{\sigma}_k$ obtained across 56 ISIC sectors, and relate them with existing evidence. Then we compute the implied aggregate elasticity. The second section presents estimates constrained to homogeneity, and compares them with conventional results making use of macroeconomic data. Both comparisons, at sectoral and aggregate levels, are sometimes drawn on the basis of the price elasticity of imports, the object of a decades-old literature. The third section discusses the heterogeneity bias and its consequences for calibration. A one-sector model of the US economy, with an elasticity calibrated to macroeconomic data, can imply dynamics at odds with what would arise from a multi-sector version. We show this to be the case in the model proposed by Backus, Kehoe and Kydland (1994). The elasticity should be calibrated according to our heterogeneous estimates if the one-sector model is to replicate the predictions of a multi-sector version. We illustrate how a similar argument can pervade a range of models in international

[18] w_{kj} does not sum to one because of non-traded sectors. Since m_{kj} and d_{kj} both sum to unity by definition, the weights are normalized.

macroeconomics. We close with some robustness.

4.1 Unconstrained Estimates

Figure 1 reports sectoral estimates of $\hat{\sigma}_k$ for 56 ISIC sectors. They range from 3.1 to 28, with average 6.7 and standard deviation 4.9. The median value, 5.1, reflects a skewed distribution: only 5 out of 56 estimates are above 10. How do our results compare with existing studies that identify $\hat{\sigma}_k$ directly at similar aggregation levels? Using various approaches, Romalis (2007) finds elasticities of substitution between 4 and 13 at the HS6 level. The top decile of estimates in Broda and Weinstein (2006) is 14.1 at TSUSA/HTS level. Hummels (2001) find values above 50 for some 5-digit SITC categories. Clearly there is large heterogeneity in $\hat{\sigma}_k$, which we are able to reproduce in Figure 1.

Figure 1: Microeconomic estimates of the elasticity of substitution

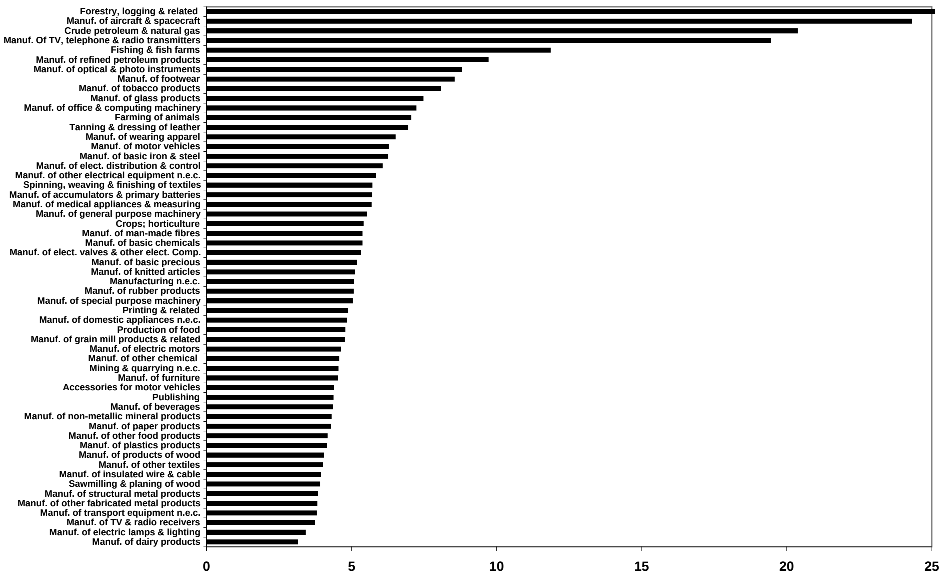

In theory, the parameters we estimate are comparable with the values obtained in the conventional approach regressing imported quantities on relative prices. The relative price of imports is typically measured with respect to an aggregate domestic (wholesale or producers) price index, as in Houthakker and Magee (1969). With this definition, the price elasticity of imports maps with σ_k according to $\eta_k = 1 - \sigma_k$.

Houthakker and Magee (1969) report in their Table 6 a long run price elasticity in manufactures estimated at -4.05. This is virtually identical to the median value we obtain across our 56 manufacturing sectors, equal to -4.1 $(= 1 - 5.1)$. Kreinin (1967) documents similar estimates, with an elasticity for manufactures equal to -4.71. It is remarkable that different data sources, coverages and methodologies should yield such similar median estimates. Importantly, Houthakker and Magee (1969) and Kreinin (1967) use information on *domestic* price indices, which we do not. Still we find similar estimates. Such convergence is reassuring from the standpoint of the key Armington assumption that underpins our estimation. At sector level, we seem to identify the object of interest to many international economists.

Houthakker and Magee (1969) and Kreinin (1967) also discuss the relative magnitudes of elasticity estimates across the broad 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). A precise mapping is difficult given differences in aggregation levels, 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 individual sector estimates one by one with the literature is difficult, because of large differences in data availability and sectoral classification. Not many papers have attempted to systematically estimate sector-specific price elasticities of imports at the two or three digit level of aggregation. When they did, they typically focused on short-run elasticities. Since our identification is in cross-section, this paper's estimates correspond to long run elasticities. They can be expected to be larger in absolute value.

Stone (1979) presents US estimates of import price elasticities at the two digit level. His estimates are lower than ours, as they are short run elasticities. For "Inorganic Chemicals", he estimates -3.40, as against -3.60 for "Other Chemicals" in Figure 1. He finds -2.32 and -3.71 in "Plastic Materials and Articles" and "Dyeing, Tanning and Coloring Agents", as against -3.1 and -5.9 in "Manufactures of Plastic Products" and "Tanning and Dressing of Leather" in Figure 1. Shiells (1991) estimates long run elasticities at the three digit SITC level, but only for 12 US sectors. His estimate in "Newsprint" is -3.6, close to our value of -3.4 for "Publishing". He also finds -3.5 in "Steel Plate and Sheet", as against -5.2 for "Manufacturing of Basic Iron and Steel" in Figure 1. Given Shiell's large standard errors, most of the differences are presumably insignificant.

At sectoral level, there is little doubt estimates of price elasticity of imports display considerable heterogeneity. Houthakker and Magee (1969) find import price elasticities estimated on aggregate data are close to zero for fifteen developed economies. They reject heterogeneity across countries. Then in a separate section, they estimate import elasticities for five US commodity classes broken down by degree of processing. They find values ranging from 0 to -4.05, and significant heterogeneity. Such findings have been reproduced numerous times since. Shiells (1991) for instance finds values below -6.

Thanks to the CES preferences we impose, the heterogeneity in $\hat{\sigma}_k$ that we identify in Figure 1 carries through into the price elasticity of imports. Our estimates of η_{kj} display considerable sectoral dispersion, because they inherit that from $\hat{\sigma}_k$. Such sectoral heterogeneity is compatible with findings in the literature, in spite of large differences in data and methodology. This will not be the case with alternative preferences, with an explicit international, single, macroeconomic elasticity of substitution, as we discuss in Section 2.2. Such alternative preferences imply the heterogeneity we document in Figure 1 has no consequence at all on the sectoral price elasticity of imports.

4.2 Constrained, Macroeconomic Estimates

Using the values from Figure 1, an estimate of the aggregate elasticity of substitution in country j follows directly from equation (3):

$$\sigma_j = \sum_k m_{kj}\hat{\sigma}_k - \sum_k \left(m_{kj} - d_{kj}\right)\left(\hat{\sigma}_k - 1\right)\lambda_j^k$$

An estimate of the constrained, aggregate elasticity of substitution is given by the same expression, with homogeneous sectoral substitutability:

$$\bar{\sigma}_j = \bar{\sigma} - (\bar{\sigma} - 1)\sum_k \left(m_{kj} - d_{kj}\right)\lambda_j^k$$

A heterogeneity bias will exist as soon as $\sigma_j \neq \bar{\sigma}_j$. To a first-order approximation, the magnitude of the bias increases with the cross-sector empirical correlation between

Table 1: Estimation with common correlated effects

	Import Elasticity η_j	Substitution Elasticity σ_j
Constrained total elasticity	-1.820[a] (.175)	3.351[a] (0.226)
Constrained partial elasticity	-2.738[a] (.263)	3.738[a] (0.263)
Unconstrained total elasticity	-4.348[a] (.745)	6.617[a] (0.962)
Unconstrained partial elasticity	-6.553[a] (1.10)	7.553[a] (1.10)
Number of sectors	56	56
Number of grid searches	11	11

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

m_{kj} and $\hat{\sigma}_k$. Since estimates of $\bar{\sigma}_j$ approximate what is implied by aggregate data, the empirical discrepancy between σ_j and $\bar{\sigma}_j$ maps directly with the elasticity puzzle.

Table 1 reports estimates of both aggregate elasticities, along with the implied values for price elasticities of imports. In the macroeconomic literature, no direct estimates of $\bar{\sigma}_j$ exist. The preference parameter is always calibrated on the basis of the aggregate imports price elasticity, η_j, estimated from macroeconomic data. Table 1 reports the constrained value of η_j, computed using equation (4) with $\sigma_k = \bar{\sigma}$. We find a macroeconomic import elasticity of -1.82. A confidence interval at standard significance levels implies values ranging from -1.5 to -2.2. This is at the high end of the range 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 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. Such small estimates have largely remained unchanged since, even though the econometrics have become considerably more sophisticated. Thus, Marquez (1990) implements a frequency domain estimator, Gagnon (2003) instruments import prices using the real exchange rate, and Hooper, Johnson and Marquez (1998) use co-integration techniques. Estimates of η_j from aggregate data rarely fall below -2. They are not significantly different from what is implied by sectoral data constrained to homogeneity.

The conventional inference in international macroeconomics is to calibrate accordingly the elasticity of substitution between aggregate bundles of domestic and foreign goods, i.e. close to zero as well. For instance, Backus, Kehoe and Kydland (1994) note that "there is some uncertainty about what value [of the elasticity of substitution] is indicated by the data. The most reliable studies seem to indicate that for the United States the elasticity is between 1 and 2" (p.91). They use a value of 1.5, with some sensitivity analysis. Obstfeld and Rogoff (2005) use a value of 2. Our estimation, based on sectoral data constrained to homogeneity, implies values in the same ballpark. This is in spite of the restrictive Armington CES preferences we impose. Without ever observing domestic prices or quantities, we are able to replicate estimates of macroeconomic elasticities that typically require data on domestic production and prices. Such convergence brings support to our assumptions.

The price elasticity of imports jumps to -4.3 when σ_k is left unconstrained across sectors. The difference is significant. The corresponding, unconstrained aggregate elasticity of substitution σ_j exceeds 6.5. We argue this is the value that should calibrate a one-sector model whose predictions are meant to replicate what would happen in a heterogeneous, multi-sector world. Calibrating the model on the basis of aggregate data simply assumes heterogeneity away. With heterogeneity, the response of aggregate quantities estimated from aggregated data is not indicative of the average elasticity of substitution.

The unconstrained elasticity lies in the same ballpark as estimates obtained in the trade literature. Hummels (2001) analyses HS6 data, and estimates a distribution of disaggregated elasticities of substitution with a mode around 9. Using disaggregated tariff data, Romalis (2007) finds values up to 11 between the US and Mexico. Perhaps most prominently, Eaton and Kortum (2002) find values for θ, a supply parameter that is equivalent to the elasticity of substitution in their Ricardian model, ranging from 3.60 to 12.86.[19] In the words of Anderson and Van Wincoop (2004), "overall the literature leads us to conclude that the elasticity is likely to be in the range of 5 to 10" in disaggregated data (p.716). Such range is consistent with our unconstrained results.

The heterogeneity bias continues to prevail in estimates of partial import elasticities. Constrained and unconstrained values are reported in Table 1. A bias is apparent: the

[19] The parameter θ is the estimated response of aggregate trade shares to price changes, measured by the *maximal* price difference across the 50 goods they observe. Goods are not imperfect substitutes in Eaton and Kortum. The relative price that is pertinent in driving demand is simply the largest price difference across disaggregated goods. Interestingly, in Section 5.2. Eaton and Kortum obtain $\theta = 3.60$, their smallest estimate, when price changes are proxied by changes in total manufacturing wages, a (semi-) aggregate variable.

constrained estimate, -2.7, jumps to -6.5 when sector specific elasticities are permitted. The heterogeneity bias is not driven by the difference between partial and total elasticities, i.e. by price indices adjustments.

We finally verify an aggregated version of our data continues to imply conventional estimates. We maintain the assumption of Armington CES preferences. The single, macroeconomic, elasticity of substitution σ_{agg} between foreign and domestic goods is still assumed identical to the elasticity between imported (macroeconomic) bundles. They are one and the same by assumption. Identification continues to rest on the cross-section of exporters to the US, and requires that aggregate bundles of imported goods be different varieties of the same good, with elasticity of substitution σ_{agg}. The assumption is hard to maintain for macroeconomic data, especially if countries are specialized. Clearly the specialization of trade towards the US varies with the level of development. To alleviate the concern, we focus on a cross-section of 24 high income OECD exporters to the U.S, where the composition of exports is relatively homogeneous.[20]

We implement our estimator on this cross-section of 24 countries, and obtain an estimate for σ_{agg} equal to 1.34. The corresponding price elasticity of imports is closer still to zero than what is implied by constrained sectoral data. But our estimates suggest $\sigma_j - \bar{\sigma}_j$ is much larger than $\bar{\sigma}_j - \sigma_{agg}$. In US data, the heterogeneity bias is first order.[21]

4.3 Implications for Calibration

In a world of sectoral heterogeneity, a multi-sector model is desirable but not always tractable, especially in an open economy context. A one-sector version is an acceptable shortcut, provided its predictions are close to what a hypothetical multi-sector calibration would imply. Of course, this is difficult to verify in practice, precisely because multi-sector models are often intractable in international macroeconomics. A prominent exception is the workhorse model in international real business cycles, BKK.

We construct a two-sector version of BKK, where σ_k varies by sector. We calibrate this

[20] Aggregate U.S. imports are given by a simple sum across sectors of the values imported from a given country, and aggregate import prices are computed as the chained Tornqvist index of HS6 specific prices.

[21] The difference between estimates of $\bar{\sigma}_j$ and σ_{agg} is small. But such a direct comparison is somewhat diagonale. The two estimates are not obtained from the same data. They are estimated on different cross-sections, with different dimensionality and sample size. They can in fact differ for many other reasons than approximations (a), (b) or (c) from Section 3.2.

version of the model using our sectoral results, and simulate a J-curve from it. We then compare the predicted J-curve to what arises from the conventional one-sector model, calibrated either using σ_j or $\bar{\sigma}_j$. The findings in this paper suggests the multi-sector J-curve is best replicated in a one-sector version of BKK calibrated using σ_j.

Since the workings of the model are well known, we leave a description of the details to Appendix B. We compare the relative performances of three models. First, a two-sector version of BKK, calibrated on our data and our results. We choose $(\sigma_1; \sigma_2) = (4.8; 12.9)$, which corresponds to averages of σ_k computed below and above the overall mean. The second and third models use $\bar{\sigma}_j$ or σ_j in conventional one-sector versions of BKK. The former captures a calibration fully ignorant of heterogeneity. The latter is consistent with the allowance for heterogeneity we have argued matters quantitatively. In all models, we calibrate w_{kj} and λ_j^k using our data. Figure 2 reports the J-curves implied by all three models. As is patent, the one-sector version of BKK that best matches the dynamics of the trade balance implied by the two-sector model is based on σ_j.

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

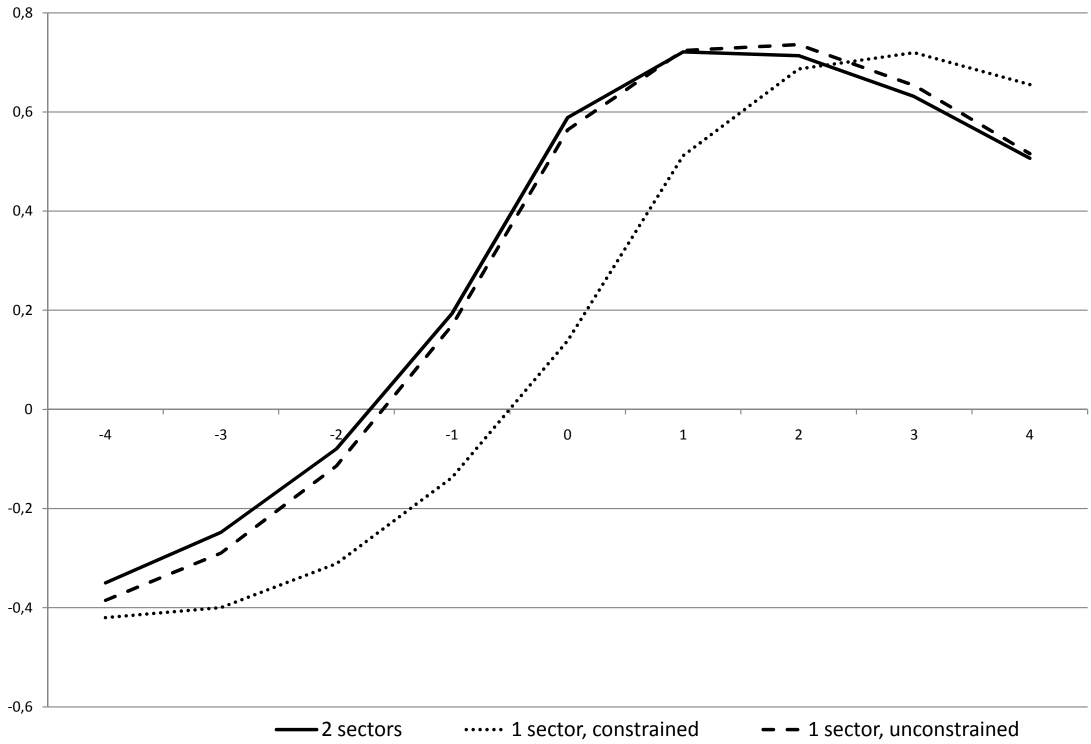

Note: The "2-sector" version is calibrated with $(\sigma_1; \sigma_2) = (4.8; 12.9)$. "One-sector, constrained" is calibrated using $\bar{\sigma} = 3.4$, and "One-sector, unconstrained" uses $\sigma = 6.6$.

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

We conjecture calibrations based on σ_j have novel predictions across a broad range of topics in international macroeconomics. A prominent example concerns models of global imbalances. 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 a depreciation rate of 21% for $\sigma_j = 7$, down from 31% with an elasticity of 2.[22] The parameter is quantitatively important, and shaves off one third of the "required" depreciation, almost all the way to the 19.3% that obtains for an elasticity of 100.

Cole and Obstfeld (1991) show the endogenous response of the terms of trade can deliver perfect insurance against country-specific shocks when the elasticity of substitution between domestic and foreign goods is exactly unitary. Consequently, models of international portfolio holdings have diametrically opposite predictions depending on whether $\sigma_j \gtrless 1$. Heathcote and Perri (2008) show that a model with CES preferences can generate a home equity bias for low enough values of σ_j. A contrario, in Coeurdacier (2009), domestic consumers choose to hold foreign assets to insure against shocks to domestic consumption, provided the terms of trade respond strongly enough in response to real shocks, which in his calibration happens for $\sigma_j = 5$. The result depends on other calibra-

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

tion choices, not least the intertemporal elasticity of substitution or the existence of trade costs. But the calibrated value of σ_j is key to portfolio choice in general equilibrium, as testified by extensive sensitivity analyses to the choice of σ_j in this literature.

The policy consequences of international price differences will presumably also depend on the substitutability between domestic and foreign goods. The relevance of the exchange rate in the monetary policy rule developed in Galí and Monacelli (2005) depends on σ_j, which they assume is equal to 1. De Paoli (2009) generalizes their model for non-unitary elasticities. When $\sigma_j = 1$, a marginal reduction in the utility value of output is accompanied by an exactly offsetting reduction in the utility value of consumption. This insulates the economy from terms of trade movements. With non unitary elasticity however, policy shocks that affect the terms of trade also affect welfare, in a way that crucially depends on whether the calibrated parameter is above or below one. Once again, other parameters also determine welfare - not least the intertemporal elasticity of substitution. But the calibration of σ_j has direct and important welfare consequences.

4.4 Robustness

This section verifies the robustness of our results in four dimensions. First, we ascertain our results do not depend on a particular choice of data source in computing w_{kj} and λ_j^k. Second, we investigate the importance of "Common Correlated Effects" in obtaining estimates of σ_k. Third, we relax our assumption that elasticities of substitution are 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. Finally, we verify a heterogeneity bias exists in the most simplistic version of trade elasticity estimates, introduced in equation (1).

Table 2 compares the constrained and unconstrained values of the total elasticities σ_j and η_j using different weighting vectors. The four variants we present in the Table are discussed in Section 3.3. The first line repeats the results implied by the benchmark weights we have used so far. Across the four variants, constrained estimates of η_j range around -1.9 and are not significantly distinguishable from conventional estimates. Unconstrained estimates are below -4. The bias continues to be quantitatively important across these four alternatives.

Table 2: Variants on the weights

	Import elasticity		Substitution Elasticity	
	Unconstrained	Constrained	Unconstrained	Constrained
Benchmark	-4.35	-1.82	6.62	3.35
Variant 1	-4.48	-1.93	6.41	3.33
Variant 2	-4.21	-1.92	6.29	3.41
Variant 3	-4.43	-1.93	6.35	3.39
Variant 4	-4.17	-1.87	6.27	3.36

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

Table 3: Estimation without common correlated effects

	Import Elasticity η_j	Substitution Elasticity σ_j
Constrained total elasticity	-2.005^a (.150)	3.590^a (.193)
Constrained partial elasticity	-3.016^a (.225.)	4.016^a (.225)
Unconstrained total elasticity	-3.915^a (.112)	6.057^a (.145)
Unconstrained partial elasticity	-5.946^a (.209)	6.946^a (.209)
Number of sectors	56	56
Number of grid searches	12	12

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

The inclusion of Common Correlated Effects is justified by our interest in the macroeconomic implications of the microeconomic values we obtain. As shown in Section 3.2, a constrained estimation without CCE may diverge from estimates arising from aggregated data. This happens because then the terms Λ_{irt}^{Y}, Λ_{irt}^{X1}, and Λ_{irt}^{X2} are present in the residual v_{it} of equation (14). They create a potential bias between constrained sectoral and aggregate estimates. Table 3 reports estimates without CCE. As expected, the constrained estimate of η decreases slightly, to -2.01, away from what is implied by aggregate data. But the changes are second order relative to our benchmark results. They do not alter the conclusion of a significant heterogeneity bias.

We now relax the assumption that the substitutability between two HS6 categories is identical within each ISIC sector. We allow for heterogeneity in σ_k within each ISIC sector, and estimate an elasticity of substitution for each HS6 sector. But we do not observe values for w_{kj} and λ_j^k at such a refined aggregation level, which complicates aggregation. We choose to impose identical values of the weights for all HS6 categories that belong to one ISIC sector. We estimate values for σ_k in 4,021 HS6 categories, and aggregate them using the benchmark (ISIC) weights in equation (3). The unconstrained import elasticity falls down to -5.06, and the constrained value is equal to -2.01. The heterogeneity bias appears to be even stronger.

Finally, we implement the simple regression analysis described in equation (1) in the introduction. We regress sectoral imports on sectoral prices, using dispersion at the HS6 aggregation level (and the time dimension) to identify a trade elasticity η for each ISIC category k. In practice, we implement for each k the Mean Group estimator introduced by Pesaran, Shin and Smith (1999) across trade partners i. This accommodates any potential heterogeneity in elasticities across trade partners, and focuses on the object of the paper, heterogeneity across sectors k. If estimates of η_k are constrained to sectoral homogeneity, we find an overall US import elasticity of -0.52. If instead sector-specific estimates of η_k are aggregated using observed US import shares m_k, we find an import elasticity equal to -2.18. Such regressions are famously problematic, as prices are obviously endogenous to quantities. This continues to be the case here, and instrumentation is both delicate and crucial for end estimates.[23] These infamous econometric difficulties are the reason why the main text builds on estimates from a structural model. But a heterogeneity bias survives even in such simplistic approach.

5 Conclusion

Estimates of the elasticity of substitution between domestic and foreign varieties are large and heterogeneous in microeconomic data, but small in macroeconomic series. We show estimations based on macroeconomic data constrain in fact the parameter to sectoral homogeneity. As a result, the elasticity puzzle can be the manifestation of a classic heterogeneity bias. It appears to be the case in US data, where the bulk of the difference in estimates can be explained by a heterogeneity bias.

[23]We instrument $\Delta \ln P_{kit}$ with HS6 intercepts, the USD nominal exchange rate in country i, and the HS6 unit value charged by country i for exports to other developed countries than the US. In general, the magnitude of the coefficients is sensitive to the set of instruments. If we use instead ISIC effects and the nominal exchange rate as instruments, the constrained estimate is -2.29, and the unconstrained jumps to -14.41.

This interpretation of the elasticity puzzle has consequences on calibration. A one-sector model calibrated on macroeconomic data may not replicate what would be implied by a multi-sector version, simply because aggregate data assumes heterogeneity away. We construct a simple multi-sector version of a workhorse model in international macroeconomics. Its predictions are best replicated in a one-sector version calibrated to disaggregated US estimates. We conjecture this intuition extends readily to other models in the field.

A 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:

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

where:

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

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

$$\sigma_j = \hat{\sigma}_j - \sum_k [m_{kj} - (m_{kj} - d_{kj})\lambda_j^k](\sigma_k - \hat{\sigma}_k)$$

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

$$\begin{aligned}
Var(\hat{\sigma}_j) &\equiv E(\sigma_j - \hat{\sigma}_j)^2 \\
\Rightarrow Var(\hat{\sigma}_j) &= \sum_k [m_{kj} - (m_{kj} - d_{kj})\lambda_j^k]^2 Var(\hat{\sigma}_k) \\
&+ \sum_k \sum_{k' \neq k} [m_{k'j} - (m_{k'j} - d_{k'j})\lambda_j^{k\prime}][m_{kj} - (m_{kj} - d_{kj})\lambda_j^k] Cov(\hat{\sigma}_k, \hat{\sigma}_{k'})
\end{aligned}$$

Since we control for common correlated effects in the estimation of the σ_ks, $Cov(\hat{\sigma}_k, \hat{\sigma}_{k'})$ is effectively zero, and the estimated variance is given by

$$Var(\hat{\sigma}_j) = \sum_k [m_{kj} - (m_{kj} - d_{kj})\lambda_j^k]^2 Var(\hat{\sigma}_k)$$

In the constrained case, the same reasoning gives

$$Var(\hat{\hat{\sigma}}_j) = \left[1 - \sum_k (m_{kj} - d_{kj})\lambda_j^k\right]^2 Var(\hat{\hat{\sigma}})$$

The variance of the constrained and unconstrained import elasticities are given by:

$$\begin{aligned} Var(\hat{\eta}_j) &= \sum_k m_{kj}^2 (1-\lambda_j^k)^2 Var(\hat{\sigma}_k) \\ Var(\hat{\hat{\eta}}_j) &= \left[\sum_k m_{kj}(1-\lambda_j^k)\right]^2 Var(\hat{\hat{\sigma}}) \end{aligned}$$

B A Two-Sector Version of BKK[24]

Each country $i = 1, 2$ is inhabited by a large number of identical agents and labor is internationally immobile. Our main departure from BKK is that each country produces two goods, a and b. Preferences of the representative agent in country i are characterized by utility functions of the form

$$E_0 \sum_{t=0}^{\infty} \beta^t U(c_{it}, 1 - n_{it})$$

where $U = \log c + \xi \frac{(1-n_t)^{1-\eta}}{1-\eta}$ and c_{it} (n_{it}) denote aggregate consumption (hours worked). Aggregate consumption is a Cobb-Douglas function of sector-specific consumption

$$c_{i,t} = \frac{{c_{i,t}^a}^{\alpha_i} \, {c_{i,t}^b}^{1-\alpha_i}}{\alpha_i^{\alpha_i}(1-\alpha_i)^{1-\alpha_i}}$$

where $c_{i,t}^a$ and $c_{i,t}^b$ are the consumption baskets of good a and b, and α_i is the share of sector a in nominal aggregate consumption. The same structure prevails for aggregate investment:

$$i_{i,t} = \frac{{i_{i,t}^a}^{\alpha_i} \, {i_{i,t}^b}^{1-\alpha_i}}{\alpha_i^{\alpha_i}(1-\alpha_i)^{1-\alpha_i}}$$

Sectoral output is produced with capital k and labor n following a Cobb-Douglas

[24]We are grateful to Jean-Olivier Hairault for sharing his codes to solve the one-sector version of BKK.

function:

$$y_{i,t}^k = z_{i,t}\left(k_{i,t}^k\right)^{\theta}\left(n_{i,t}^k\right)^{1-\theta}, \quad i=1,2, \quad k=a,b$$

The quantity $y_{i,t}^k$ denotes country i's production of good k, in units of the local good. In equilibrium, it is equal to domestic sales $c_{i,t}^{ik}+i_{i,t}^{ik}$ plus exports $c_{i,t}^{i'k}+i_{i,t}^{i'k}$. The vector $z_t=(z_{1,t},z_{2,t})$ is a shock to productivity. Importantly, productivity shocks are assumed symmetric across sectors. The cross-sector symmetry assures that, in each country, producer prices are homogenous. In what follows, domestic prices are normalized to unity and the relative price of foreign goods is denoted P.

Sectoral consumption and investment, $c_{i,t}^k$ and $i_{i,t}^k$ are composites of foreign and domestic goods:

$$\begin{aligned} c_{i,t}^k &= \left[\left(\beta_i^k c_{i,t}^{i'k}\right)^{\frac{\sigma_k-1}{\sigma_k}} + \left((1-\beta_i^k)c_{it}^{ik}\right)^{\frac{\sigma_k-1}{\sigma_k}}\right]^{\frac{\sigma_k}{\sigma_k-1}} \\ i_{i,t}^k &= \left[\left(\beta_i^k i_{i,t}^{i'k}\right)^{\frac{\sigma_k-1}{\sigma_k}} + \left((1-\beta_i^k)i_{it}^{ik}\right)^{\frac{\sigma_k-1}{\sigma_k}}\right]^{\frac{\sigma_k}{\sigma_k-1}} \end{aligned}$$

The elasticity of substitution between foreign and domestic varieties σ_k is sector-specific. The weights β_k^i are related to the share of imports in the sectoral consumption of good k. In the calibration, they are assumed symmetric across countries but potentially different across sectors.

The aggregate capital stock evolves in each country according to:

$$k_{i,t+1} = (1-\delta)k_{i,t} + i_{i,t}$$

where δ is the depreciation rate. Adjustment costs for capital are given by:

$$\mathcal{C}_{it} = \frac{\Phi}{2}\frac{(k_{i,t+1}-k_{i,t})^2}{k_{i,t}}$$

Finally, fluctuations arise from persistent shocks to aggregate productivity:

$$z_{t+1} = Az_t + \varepsilon_{t+1}^Z$$

where ε^Z is distributed normally and independently over time with variance V^Z. The correlation between the technology shocks, z_1 and z_2 is determined by the off-diagonal elements of A and V^Z.

We can obtain national income and product accounts for each country. Aggregate GDP in country 1 in period t, in units of domestically produced goods, is $y_{1t} = y_{1t}^a + y_{1t}^b$. The resource constraint equates sectoral GDPs to the sum of (domestic and foreign) consumption and investment:

$$y_{1t}^k = c_{1,t}^{1k} + c_{2,t}^{1k} + i_{1,t}^{1k} + i_{2,t}^{1k}, \quad k = a, b$$

National output is related to expenditure components according to:

$$y_{1t} = c_{1t}^{1a} + c_{1t}^{1b} + i_{1t}^{1a} + i_{1t}^{1b} + P_t(c_{1t}^{2a} + c_{1t}^{2b} + i_{1t}^{2a} + i_{1t}^{2b})$$

Finally, the trade balance, defined as the ratio of net exports to output, both measured in current prices, is:

$$nx_t = \frac{c_{2t}^{1a} + c_{2t}^{1b} + i_{2t}^{1a} + i_{2t}^{1b} - P_t(c_{1t}^{2a} + c_{1t}^{2b} + i_{1t}^{2a} + i_{1t}^{2b})}{y_{1t}}$$

and the terms of trade P_t equal the sectoral marginal rate of transformation between the two varieties in country 1, evaluated at equilibrium quantities.

Table B.1 summarizes our calibration. The elasticities of substitution at the sector level are our only deviation from BKK. The two new parameters in the multi-sector version pertain to the calibration of the Armington and Cobb-Douglas aggregators for consumption. To calibrate these, we use the sectoral data in our estimation of the aggregate substitution elasticity. In terms of the model presented in Section 2, α_i, the share of sector a in nominal consumption, is directly related to w_{kj}. In a symmetric steady state with $P = 1$, the β_i^k parameters are linked with the λ_i^k parameters according to

$$\beta_i^k = \left[\left(\frac{\lambda_i^k}{1 - \lambda_i^k} \right)^{\frac{1}{1-\sigma_k}} + 1 \right]^{-1}$$

Table B.1: Benchmark Parameter Values taken from BKK (1994)

Preferences	
[a]Discount rate	$\beta = 0.99$
[a]Labor supply elasticity	$\eta = 5$
[b]Substitution elasticity in *a*	$\sigma_a = 12.9$
[b]Substitution elasticity in *b*	$\sigma_b = 4.8$
[b]Share of *a* in consumption	$\alpha_i = 0.24$
[b]Import share in *a*	$\lambda_i^a = 0.25$
[b]Import share in *b*	$\lambda_i^b = 0.22$
Technology	
[a]Share of capital in total costs	$\theta = 0.36$
[a]Depreciation rate	$\delta = 0.025$
[a]Adjustment cost	$\Phi = 10^{-6}$
[a]SS hours worked	$n_{i,SS} = 0.34$
[a]SS hours worked in *a*	$n_{i,SS}^a = \alpha_i n_{i,SS}$
Forcing processes	
[a]Correlation matrix	$A = \begin{bmatrix} 0.906 & 0.088 \\ 0.088 & 0.906 \end{bmatrix}$
[a]Variance of productivity shocks	$Var\varepsilon_1^Z = Var\varepsilon_2^Z = 0.00852^2$
[a]Cross-country correlation of productivity shocks	$Corr(\varepsilon_1^Z, \varepsilon_2^Z) = 0.258$

Note: [a] indicates a parameter value taken from BKK. [b] indicates a parameter value from our data.

As before, d_{kj} can be expressed as:

$$d_{ki} = \frac{\left(1 - \lambda_i^k\right) w_{ki}}{\sum_k \left(1 - \lambda_i^k\right) w_{ki}}$$

References

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

Arkolakis, C., A. Costinot & A. Rodriguez-Clare, 2009, "New Trade Models, Same Old Gains?", *NBER Working Papers* 15628.

Armington, P.S., 1969, "A Theory of Demand for Products Distinguished by Place of Production", *IMF Staff Papers,* No.16.

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

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

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

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

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

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

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

Corsetti, G., L. Dedola & S. Leduc, 2008, "International Risk Sharing and the transmission of productivity shock", *Review of Economic Studies*, 75(2):443-473.

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

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

Drozd, L. & J. Nosal, 2010, "Understanding International Prices: Customers as Capital", mimeo.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography, and

Trade." Econometrica, 70(5): 1741–79.

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

Feenstra, R., M. Obtsfeld & K. Russ, 2010, "In Search of the Armington Elasticity", mimeo.

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

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

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

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

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

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

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

Hummels D., 2001, "Towards a Geography of Trade Costs", mimeo.

Imbs J., M. Ravn, H. Rey & H. Mumtaz, 2005, "PPP Strikes Back: Aggregation and the Real Exchange Rate", Quarterly Journal of Economics, vol. 120(1): 1-43, January.

Kahn M., 1975, "The Structure and Behavior of Imports of Venezuela", *The Review of Economics and Statistics*, Vol. 57, No. 2, pp. 221-224.

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

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

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

Marquez, J., 1990, "Bilateral Trade Elasticities", *The Review of Economics and Statis-*

tics, 72: 70-77.

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

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

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

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

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

PESARAN, H., Y. SHIN, AND R. SMITH, 1999, "Pooled Mean Group Estimation of Dynamic Heterogeneous Panels", *Journal of the American Statistical Association*, Vol. 94, No 446: 621-634.

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

RUHL, K., 2008, "Solving the Elasticity Puzzle in International Economics", mimeo.

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

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