The Gains From Input Trade in Firm-Based Models of Importing[*]

[Preliminary - Comments welcome]

Joaquin Blaum[†], Claire Lelarge[‡], Michael Peters[§]

March 2015

Abstract

Trade in intermediate inputs allows firms to lower their costs of production by using better, cheaper or novel inputs from abroad. Quantifying the aggregate impact of input trade, however, is challenging. As importing firms differ markedly in the intensity with which they participate in foreign input markets, results based on aggregate trade models do not apply. In this paper we therefore develop a methodology to quantify the aggregate gains from input trade for a wide class of firm-based models of importing. We provide a powerful sufficiency result: as long domestic and foreign inputs are combined in a CES fashion, the aggregate gains from input trade are fully determined from the observable joint distribution of value added and firms' domestic expenditure shares in material spending. Because our theory does not impose any restrictions on the underlying heterogeneity across sourcing countries, allows for complementarities between firm productivity and input quality and is consistent with any model of the extensive margin, i.e. of how firms find their foreign input suppliers, any firm-based model of importing will have the exact same implication, as long as it is successfully calibrated to the underlying micro data. In an application, we consider a multi-sector general equilibrium trade-model with a rich input-output structure and use data for the population of French importing firms. We find that input trade leads to a 27% reduction in consumer prices in the manufacturing sector and a 9% reduction for the full economy.

[*]We thank Pol Antràs, Costas Arkolakis, Lorenzo Caliendo, Arnaud Costinot, Jonathan Eaton, Pablo Fajgelbaum, Penny Goldberg, Amit Khandelwal, Sam Kortum, Espen Moen, Andrés Rodríguez-Clare, Peter Schott, Jonathan Vogel, David Weinstein, Daniel Xu and seminar participants at Brown, Columbia, LSE, Princeton, Stanford, UCLA, Yale and the BI Norwegian Business School. A previous version of this paper circulated under the title "Estimating the Productivity Gains from Importing".

[†]Brown University. eMail: joaquin_blaum@brown.edu

[‡]Centre de Recherche en Économie et Statistique (CREST). eMail: claire.lelarge@ensae.fr

[§]Yale University. eMail: m.peters@yale.edu

1 Introduction

A large fraction of world trade is accounted for by firms sourcing intermediate inputs from abroad. Trade theory highlights one particular margin how domestic consumers benefit from producers engaging in international sourcing. By providing access to novel and higher quality inputs, input trade reduces firms' unit costs and lowers domestic prices, therefore increasing consumers' purchasing power. Despite its importance, little is known about the magnitude of these effects. In this paper, we develop a methodology to quantify the welfare consequences of global input trade and provide an application to France.

Quantifying the gains from input trade is not straightforward as standard results of recent quantitative trade models in the spirit of Arkolakis et al. (2012) are not applicable. The reason is that importing firms differ substantially in the intensity with which they participate in foreign input markets, rendering the macroeconomic environment non-aggregative. That is, the welfare consequences of input trade cannot be determined from aggregate data and a limited number of parameters, such as an aggregate trade elasticity. A natural approach is therefore to fully specify a structural firm-based model of importing and to use it to estimate firms' unit cost reductions and thus the aggregate gains from trade. An important methodological challenge that arises is that the mathematical structure of firm-based models of importing differs remarkably from that of models of exporting. With complementarities across imported inputs, firms' entry decisions into different sourcing markets are interdependent and optimal import demand is hard to characterize, thus substantially complicating the estimation. An additional concern, common to any structural approach, is how important are the specific assumptions made in shaping the estimated gains. In this paper, we provide a methodology to bypass these concerns and quantify the effect of input trade on consumer prices without taking a stand on important components of the theory, most notably the extensive margin, thus sidestepping the complications of a structural estimation.

Figure 1 below concisely summarizes the difficulties of the study of input trade. It depicts four empirical moments from the population of French importers. The top left panel shows the cross-sectional distribution of firms' "home-shares", i.e. the share of material spending allocated to domestic input suppliers. These differ markedly. While the majority of importers spend less than 10% of their material spending on foreign inputs, some firms are heavy importers with import shares exceeding 50%. It is this heterogeneity in import intensities which renders the macroeconomic environment non-aggregative and forces researchers to resort to firm-based models of importing. The remaining panels of Figure 1 contain further challenges for such models. The top right panel shows the distribution of sales of importers and non-importers in France. While importers are significantly larger, there is ample overlap in the distribution of sales. The bottom left panel depicts the importance of multiple country sourcing. In particular, it shows the number of importing firms, and the share of aggregate imports they account for, as a function of the number of supplying countries within 8-digit products. The data shows that multiple country sourcing is not only prevalent in France, but that aggregate trade flows are shaped by firms who routinely source their narrowly defined inputs from multiple countries. Finally, the last panel analyzes the within-firm distribution of

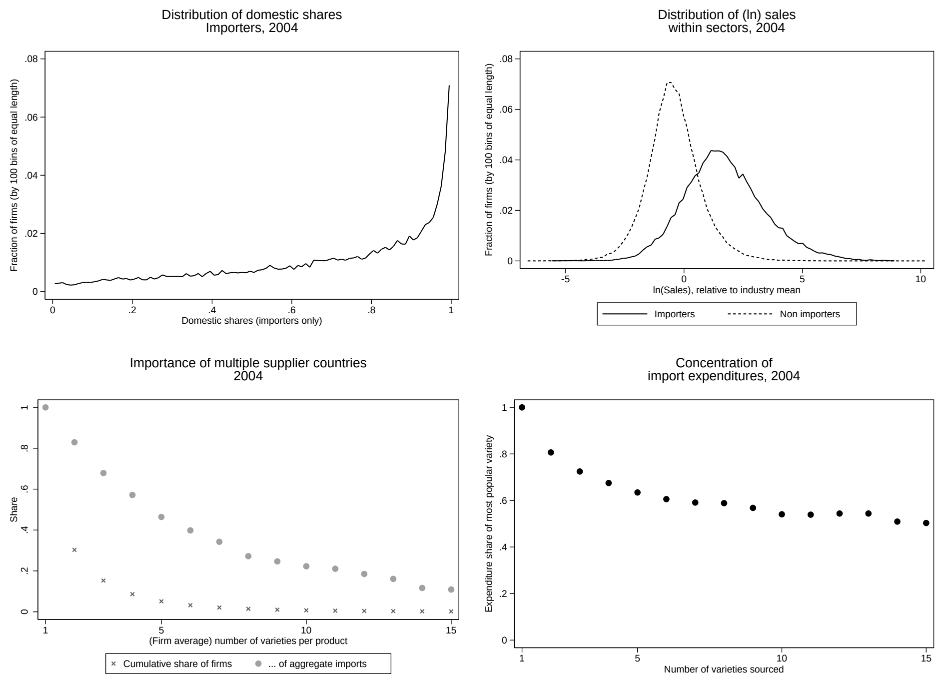

Notes: The top left panel displays the cross-sectional distribution of domestic expenditure shares, i.e. the share of material spending allocated to domestic inputs. The top right panel shows the distribution of log sales by import status. The bottom left panel displays the share of firms, and the share of aggregate imports they account for, which source their 8-digit products from at least n countries. The bottom right panel displays the within-product average expenditure share on firms' top sourcing country as a function of the number of sourcing countries.

Figure 1: Characteristics of Importing Firms in France

expenditure across such source countries. There is a striking degree of concentration: even firms with ten suppliers of a given product spend on average almost 60% of their import budget on their single top supplier country. Taken together, these facts are informative about *positive* aspects of firms' sourcing behavior. A natural question that arises is: which of these empirical patterns is important for the *normative* implications of input trade? In this paper, we argue that for many important questions the data in the top two panels of Figure 1 is all we need.

The main result of this paper is to show that micro-data on domestic expenditure shares and sales at the firm level is sufficient to calculate the aggregate gains of input trade in a wide class of firm-based models of importing and show that any model in this class features the exact same normative implications *regardless* of their particular microstructure, i.e. whether or not they match other moments of the micro data such as the bottom panels of Figure 1. More specifically, all models in this class predict exactly the same changes in consumer prices relative to a situation of "input autarky", where firms can only use domestic goods as inputs into production. This property is attractive as firm-based models of input trade are challenging to characterize, especially once they are sufficiently rich to match the patterns in the data shown above.

Our argument proceeds in two steps. First, we show that the domestic expenditure share is a sufficient statistic for the effect of input trade on the firm's unit cost in a wide class of models. In particular, any model that imposes a CES production function between domestic and foreign inputs features the property that the firm-level unit cost reduction from importing only depends on the domestic share of intermediate spending and two structural parameters: the elasticity of firm output to intermediate inputs and the elasticity of substitution between domestic and international varieties. Intuitively, the static gains from trade at the firm level, which we also refer to as micro gains, are fully summarized by firms facing a lower price index for their input bundle. Conditional on a demand system for imported intermediates, firms' import demand can be simply inverted to determine the change in prices. Importantly, the domestic share, raised to an appropriate trade elasticity, gives the trade-induced change in unit cost holding other firms' equilibrium prices constant - i.e. in partial equilibrium. We can nevertheless use this result to identify the distribution of unit cost changes across firms, and thus study how unequally are the gains from input trade distributed. Second, we show how to aggregate these micro gains to compute the economy-wide consequences of input trade taking into account general equilibrium effects. We consider a canonical multi-sector trade-model with intersectoral linkages, roundabout production and monopolistic competition. We show that the change in a welfare-relevant consumer price index can easily be calculated from microdata on firms' domestic expenditure shares and value added. Thus, as with the micro gains, the macro gains can be essentially read-off from the micro-data. Additionally, we explore whether using an approach based on aggregate data results in a bias in the estimated gains from input trade. We show that indeed aggregate trade models imply gains that are biased relative to the gains computed with our formula.[1] We later quantify such bias.

A remarkable property of this procedure is its generality. Besides the CES production structure between imported and domestically produced inputs, no further assumptions on other structural primitives of firms' import environment are required. Notably, we do not place any restrictions on the underlying distribution of qualities and prices across potential sourcing countries, we can allow for firm productivity and input quality to be complements, giving rise to non-homothetic input demand, and we do not even have to assume that firms share the same production function for imported inputs from abroad. Crucially, we also do not have to take stand on how firms end up with their set of trading partners, i.e. the extensive margin of importing. Hence, our estimates of the aggregate gains from input trade are consistent regardless of whether firms find their trading partners on a spot market, in which case importing might be limited through the presence of fixed costs, a process of network formation or through costly search. Not only will all models within this class deliver the exact same gains from input trade conditional on the micro-data, but calculating the gains is straightforward.

Our methodology relies on micro data and a set of parameters. Of particular importance is the elasticity of substitution between domestically sourced and imported inputs. Instead of identifying this parameter from aggregate trade flows as is common in the literature, we devise a novel strategy to identify this trade elasticity from firm-level variation. More specifically, we exploit the fact that, for a given level of total spending in materials and given quantities of all other inputs, firm revenue

[1]Surprisingly, whether the bias from aggregate models is positive or negative depends only on parameters and not on the microdata.

depends on the domestic expenditure share. The degree to which changes in such share translate into changes in firm revenue depends on the elasticity of interest. Thus, we rely on the panel structure of the micro-data to look at the within-firm relationship between changes in domestic spending and changes in firm revenue. To address the endogeneity concern that unobserved productivity shocks might lead to both higher import spending and higher revenue, we follow Hummels et al. (2011) to use changes in the world supply of particular varieties as an instrument for firms' import spending. Using the micro-data to estimate this elasticity turns out to be important in the French data as we estimate an elasticity of substitution that is close to a value of two - and thus substantially smaller than the corresponding estimates from aggregate trade flows, usually around a value of four. We conclude that approaches that rely on aggregate data to evaluate the normative consequences of input trade will result in a substantial bias through the use of an inappropriate trade elasticity.

We apply our procedure to the population of manufacturing firms in France. We start by estimating the distribution of trade-induced changes in unit costs across firms. We find substantial cross-sectional dispersion in the gains, as well as concentration. While the median firm's unit cost is reduced by 11% through input trade, this number is 22% for the average firm. The gains are larger for exporters, members of a foreign group and larger firms, as expected. As these results do not rely on any assumptions about the microstructure of trade, *any* model that combines domestic and imported inputs in a CES fashion will predict exactly the same distribution of static gains across firms, given the micro-data. We then aggregate these micro-gains to the macro gains from input trade. As discussed above, the aggregate gains are fully determined by the joint distribution of the micro-gains and value added, together with a set of parameters, such as the matrix of input-output linkages, which we estimate. We find that the aggregate gains from input trade for the manufacturing sector are 27% and hence substantially higher than the median gains at the firm-level. This is due to three reasons. First, innate firm productivity and domestic spending are negatively correlated, so that bigger firms benefit more from international trade. Second, firms' domestic expenditure shares show substantial cross-sectional variation, causing the micro-gains from input trade to be quite dispersed. This is beneficial for the aggregate economy given an elastic demand. Finally, there are important interlinkages between firms, whereby non-importers in France source intermediate inputs from importing firms. This structure of round-about production increases the aggregate gains. When we consider the entire economy, including both the manufacturing and non-manufacturing sectors, the gains amount to 9% and hence are more limited. This is due to the fact that the manufacturing sector accounts for only a relatively small share in aggregate consumer spending and that production links between the manufacturing and the non-manufacturing sector, which we assume to be closed to international trade, are limited.

We also quantify the biases that arise from approaches that rely on aggregate data to measure the aggregate gains from input trade. Applying the results in Arkolakis et al. (2012) to our setting, while keeping our estimate of the trade elasticity, leads to over-estimating the gains from trade. The gains for the manufacturing sector and the aggregate economy would be given by 31.5% and 10% respectively. If instead we used a value for the trade elasticity close to four, as suggested by estimation approaches that rely on aggregate data, the gains from input trade implied by aggregative models become up to 50% smaller than our estimates. Thus, while the sign of the bias from using aggregate data depends on the details of the exercise, we find that its magnitude can be substantial.

Finally, we consider the effect of input trade on a broader notion of welfare. Our sufficiency result allowed us to quantify the effect of input trade on consumer prices without fully solving and estimating a structural model. These price-index gains, however, do not take into account the resources spent by firms to attain their equilibrium sourcing strategies. To measure the full welfare consequences of input trade, we therefore need to commit to a particular model and take a stand on aspects that were left unrestricted before - most notably, the extensive margin of trade. In the last section of the paper, we specify a structural model where participation in international markets is limited by fixed costs. We parametrize the distributions of qualities, prices and fixed costs, and impose assumptions that guarantee that firms' extensive margin problem remains tractable. We calibrate the model to moments of the French data, most importantly the joint distribution of domestic expenditure shares and sales. This ensures that the calibrated model matches the price-index gains measured above. To match the far from perfectly negative correlation between domestic expenditure shares and sales observed in the data, the model needs to feature two dimensions of firm heterogeneity - a natural choice is firm productivity and fixed costs. The main result of this exercise is that, according to the calibrated model, the welfare gains from trade are about half of the price-index gains.

We also use the calibrated model to assess the value of the firm-level domestic share data. When such data is not available, our sufficiency result to measure the price-index gains cannot be applied. Using data on firm sales only, we calibrate a version of the structural model with fixed costs and find gains from input trade that are 10-20% larger than the ones obtained before. The reason is that, without information on how firm size and the domestic share correlate, the model generates too strong a negative correlation between physical productivity and the domestic share. This means that firms that are physically more productive experience larger reductions in their unit cost, a feature that tends to make input trade more attractive. This exercise highlights the importance of the micro data on domestic shares for our understanding of the normative implications of input trade.

Related literature. At a conceptual level, this paper is closely related to Arkolakis et al. (2012) and Costinot and Rodriguez-Clare (2014) in that our sufficient statistic for firm productivity is related to their sufficient statistic for aggregate welfare. In particular we also show that conditional on the micro data and a "trade elasticity", a wide class of models will imply the exact same productivity gains, albeit at the firm-level.

Recently, a number of papers have focused on measuring the effect of imported inputs on firm productivity.[2] One strand of the literature provides reduced-form evidence by studying trade liberalization episodes (Amiti and Konings, 2007; Goldberg et al., 2010; Khandelwal and Topalova, 2011; Kasahara and Rodrigue, 2008). Our results are complementary to this literature as we provide a structural interpretation to these reduced form regressions. From the point of view of applied researchers, our sufficiency result provides a convenient way to analyze episodes of trade liberalization or other changes in firms' import environment, without having to fully specify and solve a full struc-

[2]See also De Loecker and Goldberg (2013) for a recent survey about firms in international markets.

tural model of import behavior. The change in firms' share in domestic spending correctly measures the effect of the policy on firm productivity, taking all adjustments into account.[3]

There is also a small literature that takes a more structural approach. Halpern et al. (2011) use Hungarian micro data to estimate a closely related framework of importing.[4] The main difference with our approach is that they do not feature firms' domestic expenditure share as the sufficient statistic for firms' productivity gains. Hence, they have to estimate the entire structural model simultaneously to identify all the structural parameters. Because the firm's extensive margin problem is tractable only under particular assumptions and one has to specify the entire market structure and demand parameters on output markets, their estimating procedure relies on these specifications.[5] Gopinath and Neiman (2014) use a related structural model to measure the aggregate productivity losses during the Argentine crisis. While they also do not feature the sufficiency property of domestic expenditure shares and hence have to simulate the entire structural model, our characterization of the aggregate welfare consequences shares some similarities. Finally, Ramanarayanan (2014) calibrates a fully-specified firm-based model of importing to Chilean data to study the gains from input trade.

On a more technical level, our paper builds on a recent literature (Blaum et al., 2013; Antràs et al., 2014) that stresses that complementarities across inputs of production make the import problem very different from the better known export problem in that firms' extensive margin of trade is - in general - harder to characterize. On the export side[6], recent work studied firms' entry behavior into different markets and developed quantitative models that can come to terms with the evidence (Eaton et al., 2011; Arkolakis, 2010; Arkolakis and Muendler, 2011; Bernard et al., 2011). In contrast, theories that can quantitatively account for the pattern of entry into different import markets are far less developed. A notable exception is the recent contribution by Antràs et al. (2014), who analyze a firm-based model of importing in the spirit of Eaton and Kortum (2002) and embed it into a general equilibrium framework. They adapt the estimation procedure by Jia (2008) and are able to quantitatively account for importers' extensive margin of trade observed in the US firm-level data. While their framework also falls into the class of models covered by our main result, so that conditional on firms' observed domestic spending, the productivity gains from input trade do not depend on firms' extensive margin of trade, their framework is, to the best of our knowledge, the first one to explicitly allow for rich counterfactual policy analysis.

The remaining structure of the paper is as follows. Section 2 lays out the general framework of importing and proves our firm-level sufficiency result. Section 3 derives the main result, namely the sufficient statistic for the aggregate gains from input trade in a rich general-equilibrium trade model, and characterizes the bias viz-a-viz an approach based on aggregate trade data. Our empirical application to France is contained in Sections 4 and 5. In Section 6 we finally calibrate a fully specified

[3] Amiti and Konings (2007) in fact consider one specification where the use the domestic expenditure share as a measure of import behavior. However, they do not offer a structural interpretation.

[4] In fact, our framework nests the one in Halpern et al. (2011) and Gopinath and Neiman (2014).

[5] They also do not measure the aggregate gains from import trade as they do not consider a general equilibrium environment.

[6] The literature on exports is too vast to discuss here. Hence, we refer to the reader to Bernard et al. (2007) and Bernard et al. (2012), which are two recent surveys of the literature.

model of importing to provide a full measure of welfare. Section 7 concludes.

2 Input Trade and Unit Costs

In this section we show that a wide class of firm-based models of importing share a powerful property: the effect of international input sourcing on the firm's unit cost is fully summarized by its domestic expenditure share, i.e. the share of material spending that is accounted for by domestic varieties. Hence, to the extent that data on domestic spending is readily available, the firm-level unit cost reductions associated with input trade - which we refer to as the *micro-gains* - are observable.[7] Crucially, this result allows us to bypass the computation of the firm's optimal import demand as well as the estimation of many structural parameters which would be otherwise required. As will be clear below, this will prove especially useful in the context of firm-based models of importing since their mathematical structure can entail substantial computational complexity.

In Section 2.1 we consider first a simple economy where input trade is limited by the presence of fixed costs and we illustrate how domestic expenditure shares are a sufficient statistic for the unit cost reduction of input trade. Then, in Section 2.2 below, we show the full generality of the result, namely, that it requires no assumptions on the structure of product markets or the mechanics of the extensive margin, and that it holds under virtually unrestricted firm and input heterogeneity. Thus, conditional on the micro-data, every model in the class described in Section 2.2 predicts exactly the same unit cost reductions of international trade.

2.1 A Simple Example

We start by deriving our main unit cost result in the context of a simple environment with fixed costs. Consider an economy populated by a set of firms that have access to the following production structure:

$$\begin{aligned} y &= \varphi f\left(l, x\right) = \varphi l^{1-\gamma} x^{\gamma} && (1) \\ x &= \left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + x_I^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}} && (2) \\ x_D &= q_D z_D && (3) \\ x_I &= \left(\sum_{c\in\Sigma} \left(q_c z_c\right)^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}}. && (4) \end{aligned}$$

Intermediate inputs (x) are combined with primary factors (l), which for simplicity we call labor, in a Cobb-Douglas way to produce output (y). Firms are heterogeneous in their efficiency (φ). Intermediate inputs are a CES composite of a domestic variety (x_D) and a foreign import bundle (x_I). In turn, the foreign bundle is a CES aggregate of many heterogeneous foreign varieties.[8] The

[7] Throughout the paper we use the term "gains" to refer to the change in unit cost between the current trade equilibrium and autarky.

[8] We do not distinguish for now between products and varieties, i.e. imports of a given product stemming from different countries. As will be clear below, we do not need to take stand on this distinction for our main result. In the

effective efficiency of the variety from country c is given by the product of physical units sourced (z_c) and a quality flow (q_c) whose distribution is exogenously given by $G(q)$.

An important endogenous object in the production structure is the firm's sourcing strategy Σ, which is the set of foreign countries the firm sources from. We assume that international sourcing is limited by the presence of fixed costs, which are denominated in units of labor.[9] In particular, we assume that sourcing an input from country c entails paying a fixed cost f_c, which can vary by firm. Firms also need to pay a fixed cost to start importing, f^I. It is precisely the existence of fixed costs at the country level which prevents firms from sourcing from all countries in the world: the optimal Σ is determined by balancing love of variety effects vs fixed costs. Crucially, variation in physical efficiency and fixed costs across firms will map into variation in sourcing strategies: some firms will endogenously import from more, potentially better markets than others.

An important assumption that will be maintained throughout the analysis is that firms are price-takers in input markets. That is, conditional on accessing a country c, firms can source any quantity at price p_c. Note that prices include trade costs and are allowed to be correlated with country quality.

Finally, we do not place any restrictions on the structure of domestic output markets - and hence do not specify whether firms produce a homogeneous or differentiated final good and how they compete.

This environment is a description of a standard firm-based model of trade. In particular, it nests most of the existing contributions directly. Gopinath and Neiman (2014) for example assume that sourcing countries are identical ($q_c = q$, $p_c = p$, $f_c = f$) and that output markets are monopolistically competitive with isoelastic demand. Halpern et al. (2011) consider only a single foreign country and allow for firm-specific fixed costs (f_i) that are uncorrelated with firm efficiency.[10] They also assume monopolistic competition in output markets.

Import Demand Firms choose their size and set of imported varieties, as well as the quantities of all inputs to maximize profits. It is convenient to split the firm's problem into a cost-minimization subproblem where the sourcing strategy is taken as given, and the choice of the optimal firm size y and sourcing strategy Σ given the cost function for foreign materials. Formally,

$$\pi \equiv \max_{\Sigma,y,l} \left\{ py - \Gamma(\Sigma, y, l; \varphi) - wl - w\left(\sum_{c\in\Sigma} f_c + f^I I(\Sigma) \right) \right\} \tag{5}$$

where

$$\Gamma(\Sigma, y, l; \varphi) \equiv \min_{z} \left\{ \sum_{c\in\Sigma} p_c z_c \text{ s.t. } \varphi f(l, x) \geq y \right\} \tag{6}$$

empirical application of Section 4 below we focus mostly on the country dimension by including product fixed effects. We hence use the terms "foreign input", "variety" or "country" exchangeably in what follows.

[9]Section 2.2 below shows that the unit cost result of this Section holds *regardless* of the mechanism by which the extensive margin of importing is determined.

[10]More precisely, Halpern et al. (2011) allow for multiple *products* $\left(x = \prod_k x_k^{B_k}\right)$, each of which is produced according to (2). We explicitly show in the Appendix how our result can be generalized to a multi-product environment.

is the firm's cost function for foreign materials. Here p denotes the demand function and w the wage which are taken as given by the firm. Furthermore, $I(\Sigma)$ is an indicator of the firm's import status, i.e. $I(\Sigma) = 1$ whenever the firm sources *any* foreign variety. We do not take a stand on the nature of competition or market structure on the output side, and hence the demand function p is unrestricted.

It is the solution to the firm's import demand problem which determines the benefits of importing intermediate inputs. Albeit conceptually easy, quantifying these gains presents us with two practical challenges. First, the choice of the optimal sourcing strategy in (5) can be computationally difficult - see Blaum et al. (2013) and Antràs et al. (2014) for a discussion. The reason is the interdependence between entry decisions in different import markets. A particular variety is imported whenever the reduction in the average production cost outweighs the incurred fixed costs. When imported varieties are imperfect substitutes[11], the cost reduction associated with entering a particular foreign market depends on the quantities sourced from all other markets. If foreign inputs differ in both quality and fixed costs, the profit maximization problem in (5) is non-convex and the choice of the optimal sourcing set requires evaluating all possible sourcing strategies, entailing substantial computational burden.[12] This interdependency of entry decisions makes the extensive margin of imports different from that of exports, where the sourcing strategy can typically be solved market by market - see for example Eaton et al. (2011). A related implication is that, unlike in the case of exports, more productive firms need not source their inputs from more countries, unless more restrictions are imposed.[13][14]

A second challenge to quantifying the gains is that solving the firm's problem in (5)-(6) requires knowledge of many structural primitives of the model, including the distribution of prices, qualities and fixed costs across countries as well as the local output demand function. While these parameters can be in principle estimated, such estimation would typically entail solving for the optimal sourcing set in (5).[15] Thus, in general, the extensive margin problem cannot by sidestepped even in cases where the researcher is interested in computing unit cost changes between two states where the sourcing sets are known - e.g. the current trade equilibrium and autarky - as evaluating the cost function in (6) at a particular Σ requires knowledge of all parameters.

The main insight of this paper is that we can bypass both of the challenges described above and measure the effect of international trade on firms' unit cost by focusing on the intensive margin

[11] That is, $\rho < \infty$.

[12] The number of possible combinations of sourcing countries is given by 2^{N_C}, where N_C is the number of available countries.

[13] It might be that unproductive firms source multiple varieties with low fixed costs and low quality flows while high productivity firms concentrate on few fixed-cost-expensive varieties that yield high quality flows.

[14] Note, however, that this does not imply that general results concerning the extensive margin cannot be derived. If for example the demand elasticity exceeds unity and fixed costs are not firm-specific, more productive firms import and more productive firms adopt a sourcing strategy that leads to lower unit costs, i.e. $\gamma(\Sigma(\varphi')) \leq \gamma(\Sigma(\varphi))$ if $\varphi' > \varphi$. Hence, similar to the exporting intuition, more productive firms sell more and thus have a higher incentive to reduce their marginal costs by incurring the fixed costs of importing additional products/varieties. However, this does *not* imply that more productive firms source *more* varieties or products. Antràs et al. (2014) show in a related model that sourcing can be shown to be hierarchical as long as the profit function function has increasing differences.

[15] One prominent example of this approach is Halpern et al. (2011), who impose sufficient assumptions to make the extensive margin problem tractable and structurally estimate a model of importing to then compute changes in unit costs.

problem in (6). We exploit an insight that allows us to dissociate the choice of the optimal sourcing set from the evaluation of the unit cost function.

Intensive Margin of Imports We now derive the firm's unit cost associated with the production structure given by (1)-(4). We start by noting that, given the firm's sourcing strategy Σ, fixed costs are irrelevant for the computation of the optimal quantities sourced. Standard cost minimization calculations imply that optimal import services x_I satisfy $x_I A(\Sigma) = m_I$, where m_I denotes import spending and $A(\Sigma)$ is a *firm-specific* import price index given by:

$$A(\Sigma) \;=\; \left(\sum_{c\in\Sigma}(p_c/q_c)^{1-\rho}\right)^{\frac{1}{1-\rho}}. \tag{7}$$

This price index will play an important role because it is exogenous *given* the sourcing strategy Σ. In a similar vain, total material services x are related to total material spending m via $Q(\Sigma)\,x = m$, where the firm-specific material price index $Q(\Sigma)$ is given by:

$$Q(\Sigma) = \left((p_D/q_D)^{\varepsilon-1} + A(\Sigma)^{\varepsilon-1}\right)^{\frac{1}{\varepsilon-1}}. \tag{8}$$

where p_D is the price of the domestic variety.

The unit cost of firm i is therefore given by:[16]

$$UC_i = \frac{1}{\varphi_i} w^{1-\gamma} Q(\Sigma_i)^{\gamma}. \tag{9}$$

The expression in (9) illustrates the challenges of firm-based models of importing: the firm's unit cost depends on the endogenous sourcing strategy Σ as well as on the distribution of prices and qualities across countries required in (7) and (8). Because these parameters are unknown, the standard approach in the literature consists of structurally estimating the full model, thus coming up with a numerical approach to solve for Σ as (5) does not allow for analytical characterization - see Halpern et al. (2011); Gopinath and Neiman (2014); Ramanarayanan (2014); Antràs et al. (2014).[17] This requires simultaneously estimating all structural primitives of the model beyond the ones directly present in (9) - e.g. the distribution of fixed costs and the structure of output markets.

We take a different approach. Instead of solving for $Q(\Sigma)$ in terms of primitives, we explicitly use the fact that the *unobserved* price index $Q(\Sigma)$ is related to the *observed* expenditure share on domestic inputs via:

$$s_D(\Sigma) \;=\; \frac{(q_D/p_D)^{\varepsilon-1}}{(q_D/p_D)^{\varepsilon-1} + (A(\Sigma))^{\varepsilon-1}} = \left(\frac{q_D/p_D}{Q(\Sigma)}\right)^{\varepsilon-1}. \tag{10}$$

[16] With a slight abuse of notation we suppress the constant $\left(\frac{1}{1-\gamma}\right)^{1-\gamma}\left(\frac{1}{\gamma}\right)^{\gamma}$ in the definition of (9).

[17] Note that this is necessary even to compute the change in the unit cost between the current trade equilibrium and autarky - i.e. two states where Σ is known. The reason is that the standard approach to structurally estimating the model entails solving for the optimal import demand.

Using (10) and (9) gives the following result.

Proposition 1. *Consider the model above. The unit cost of firm i is given by:*

$$UC_i = \frac{1}{\varphi_i} \times (s_{D,i})^{\frac{\gamma}{\varepsilon-1}} \times \left(\frac{p_D}{q_D}\right)^{\gamma} w^{1-\gamma}. \tag{11}$$

Despite its simplicity, Proposition 1 is powerful. It says that we can measure the endogenous reduction in unit cost arising from trade simply from the observed domestic shares s_D and the structural parameters γ and ε, which can be estimated. Hence, firms' domestic expenditure shares are sufficient statistics for the cost reductions associated with international sourcing: conditional on s_D and the two elasticities, neither the extensive margin of trade Σ nor any other underlying structural primitive such as the distribution of import quality q_c, prices p_c, or fixed costs f_c are required to determine the gains at the firm-level. Hence, a large class of models will have the exact same answer for the implied micro gains from trade given firm-level data on domestic spending shares and parameters ε and γ.[18]

The expression in (11) is akin to a firm-level analogue of Arkolakis et al. (2012). Broadly speaking, they show that the domestic expenditure share at the *country level* is a sufficient statistic for welfare in models where the demand system is CES, among other conditions.[19] By analogy, (11) states that *at the firm level* the share of materials spent on domestic intermediates raised to an appropriate trade elasticity is a sufficient statistic for the endogenous component of firm productivity associated to importing. In the same vain as consumers gain purchasing power by sourcing cheaper or complementary products abroad, firms can lower the effective price of intermediate purchases by tapping into foreign input markets.

In the next section we show that Proposition 1 holds in a much more general environment, which arguably covers the vast majority of frameworks that can be employed in quantitative models of imports.

2.2 The General Model

We now enrich the simple model from above in three dimensions. We allow for more flexibility in the sources of *firm heterogeneity,* in *technology* and in the way firms decide about their *extensive margin.*

[18] Additionally, there are other models, which satisfy the Proposition 1. Antràs et al. (2014) for example consider a model of importing in the spirit of Eaton and Kortum (2002). In their model, firms' total productivity is proportional to $\vartheta = \varphi s_D^{-1/\theta}$, where φ also denotes firms' exogenous productivity and θ is the parameter of the Frechet distribution, where suppliers' efficiency levels are dawn from. Hence, as in Arkolakis et al. (2012), the CES parameter $\varepsilon - 1$ changes to the heterogeneity parameter θ, but the basic result of Proposition 1 remains intact: conditional on the *observed* domestic shares and a value of θ, the productivity gains of importing are fully determined. The material share γ is equal to unity in Antràs et al. (2014) as their importing firms do only use intermediary products as an input to production. In our empirical exercise we will estimate ε directly from the micro data using exogenous variation in firms' domestic spending. Hence, conditional on the exclusion restrictions of our instrument, our estimates of ε will also be consistent for θ.

[19] As emphasized below, the firm-based models of importing considered in this paper do not meet the conditions required in Arkolakis et al. (2012) and therefore the aggregate domestic expenditure share will not be sufficient for welfare.

Consider first the case of extra dimensions of *heterogeneity* across firms. Because Proposition 1 was derived at the firm level, we can allow for firm heterogeneity in essentially any of the structural primitives of the firm's problem. For example, we can allow for the distribution of quality across countries to be firm-specific, i.e. firm i faces a schedule of import qualities given by $G_i(q)$. A prominent instance of such firm-specific quality flows is the case where there is a complementarity between the firm's efficiency φ and the input's quality q_c - the effective quality flow of imports from country c is given by $\eta(q_c, \varphi)$.[20] This form of non-homothetic import demand is consistent with the findings reported in Kugler and Verhoogen (2011) and Blaum et al. (2013).[21] On the supply side we can also cover the case of import prices varying at the firm-country level, i.e. firm i facing a price of p_{ci} to import an input with quality q_c. An example of such a pricing function is

$$p_{ci} = r_c(q_c) \times \tau_{ci},$$

where $r(.)$ are the costs of producing quality q_c in country c and τ_{ci} are firm-specific iceberg trade costs.

Now consider the case of *technology.* We can substantially generalize the technology for imported inputs. Instead of the CES production function in (4) we can allow for any firm-specific constant returns to scale (CRS) production function:

$$x_I = h_i\left([q_c, z_c]_{c \in \Sigma}\right).$$

With h_i being CRS, firms still face an import price index $A_i(\Sigma)$ akin to (7), which is exogenous given the extensive margin. And, as the derivation of Proposition 1 shows, this is all that is required for the inversion of the demand system to express the (endogenous) input price $Q_i(\Sigma)$ in terms of the observable domestic share. Also note that we only used the constant output elasticity of intermediate inputs. Hence, we can also cover more general production functions of the form

$$y_i = \varphi_i f_i(l, k, x) = \varphi_i e_i(l, k)^{1-\gamma} x^{\gamma},$$

where $e_i(.)$ is some firm-specific CRS production function between other inputs like labor and capital. In that case, w_i would simply be the constant price for the bundle $e_i(l, k)$, which depends on the prices for l and k.

Additionally, we can allow for a firm-specific import bias in the production function for interme-

[20]Formally, the production function for the foreign input bundle in (4) is now given by $x_I = \left(\sum_{c \in \Sigma} (\eta(q_c, \varphi) z_c)^{\frac{\rho-1}{\rho}}\right)^{\frac{\rho}{\rho-1}}$.

[21]In particular, there is quality-productivity complementarity when h is log-supermodular in (q, φ). If h is log-submodular, firm productivity and product quality are substitutes. Kugler and Verhoogen (2011) provide evidence for the case of Colombian producers of a positive correlation between plant size and input prices. This evidence is consistent with the presence of a complementarity between input quality and firm productivity. Blaum et al. (2013) provide evidence for French manufacturing firms that is also consistent with such complementarity.

diate inputs (2), by which some firms have a higher demand for imported varieties given prices:

$$x_i = \left(\beta_i x_D^{\frac{\varepsilon-1}{\varepsilon}} + (1-\beta_i)\, x_I^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}},$$

where $\beta_i > 0$ is firm-specific.

Finally, we turn to the *extensive margin.* The key behind Proposition 1 was to focus entirely on the intensive margin problem, thus dissociating the characterization of the unit cost from the firm's interaction on output markets and the determination of the sourcing strategy Σ. It follows that this proposition was derived under no assumptions on how the extensive margin of trade comes about. In an economy with fixed costs, we can leave the entire distribution of fixed costs unrestricted, i.e. f_i^I and f_{ci} can vary by country and firm in any way. Beyond fixed costs, we could allow for a dynamic process for the extensive margin whereby Σ evolves according to some mapping $\Sigma' = \phi(\Sigma, e, u)$, where e denotes the firm's expenditures on finding new trading partners (e.g. resources allocated towards the search for new suppliers) and u encapsulates random shocks. These could either be events that cause break-ups in existing supplier relationships or "matches" with new suppliers through existing partners as in Chaney (2013) or Oberfield (2013).

Despite all the added generality, the main result of this Section is that Proposition 1 remains essentially unchanged.

Proposition 2. *Consider the general model above. The unit cost of firm i is given by:*

$$UC_i = \frac{1}{\tilde{\varphi}_i} \times (s_{D,i})^{\frac{\gamma}{\varepsilon-1}} \times \left(\frac{p_D}{q_D}\right)^{\gamma} w^{1-\gamma}, \tag{12}$$

where $\tilde{\varphi}_i \equiv \varphi_i \beta_i^{\frac{\varepsilon\gamma}{\varepsilon-1}}$.

Proof. See Section 8.1 in the Appendix. □

As in Proposition 1, we find that the effect of participating in international input markets on the firm's unit cost is observable given values of the elasticities γ and ε.[22] In particular, the change in unit cost between the current trade equilibrium and autarky is given by the firm's domestic expenditure share raised to an appropriate trade elasticity. An important caveat is that these micro gains are partial equilibrium, as prices p_D and w can change.[23] The term $s_{D,i}^{\gamma/(\varepsilon-1)}$ therefore measures the loss in productivity that firm i would experience if it (and only it) was excluded from international markets. We can nevertheless use Proposition 2 to identify the *distribution* across firms of the gains from input trade and, in particular, to study how unequal this distribution is.[24] Precisely this is

[22]In Section 9.1 in an Online Appendix, we also consider two additional generalizations. We explicitly show that we can derive a local version of (12) if domestic and foreign inputs are not combined in a CES fashion and we derive Proposition 2 for the case where the output elasticity of material inputs is not constant. There we also discuss explicitly what additional information one would require to perform counterfactual analysis in that case.

[23]The distribution of *relative* unit cost changes can, however, be fully identified from the data. That is, the distributional implications of trade in inputs can be fully read from the micro data.

[24]To see this, note that prices $p_D$and w drop out when we normalize the unit cost changes by that experienced by one particular firm, call it firm 1, i.e. $(s_{D,i}/s_{D,1})^{\gamma/(\varepsilon-1)}$.

what we do in Section 5.1 below. How these firm-level unit cost reductions translate into aggregate welfare gains depends on the nature of product markets, i.e. the rate of pass-through to consumers, and on the nature of interlinkages between firms. Section 3 below deals with the aggregation of the micro gains in a general equilibrium setting.

Before doing so, however, we want to stress that Proposition 2 can be used beyond measurement of the gains from trade relative to autarky. We can apply Proposition 2 to analyze the effects of any trade policy or shock, e.g. an episode of past trade liberalization, as long as data on domestic shares is available before and after. The distribution of partial equilibrium changes in firm productivity through access to better or complementary foreign inputs as a result of the policy is given by:

$$\left.\frac{UC_i^{post}}{UC_i^{pre}}\right|_{\varphi} = \left(\frac{s_{D,i}^{post}}{s_{D,i}^{pre}}\right)^{\gamma/(1-\varepsilon)} . \tag{13}$$

Thus, knowledge of the change in the domestic shares is sufficient to analyze the direct, static consequences of a trade policy or shock. Consider for concreteness the productivity effects of an episode of trade liberalization (e.g. Chile in 1980s (Pavcnik, 2002), Indonesia in the late 1980s and early 1990s (Amiti and Konings, 2007) or India in the 1990s (De Loecker et al., 2012)). One can then use (13) to estimate the direct effect of improved access to international inputs on firm productivity.[25] In particular, (13) contains both the exogenous change in foreign prices due to lower trade barriers and tariffs as well as the endogenous change resulting from adjustments in the sourcing pattern.[26]

Finally, we point out that Proposition 2 is not directly amenable to answer counterfactual questions. While the formula is still valid in any scenario, more structure is required to forecast the counterfactual domestic shares.[27]

3 The Aggregate Gains from Trade

In this section we embed the model of firm behavior from Section 2 in a macroeconomic environment and study the aggregate effects of input trade. The key innovation in our methodology is to exploit the availability of firm-level domestic spending data to bypass the specification of the extensive margin. As discussed above, the micro data on domestic expenditure shares identifies the distribution of unit costs up to general equilibrium constants. By specifying how firms compete in output markets and the nature of input/output linkages across firms, we can effectively map those unit costs - see (12) - into aggregate allocations and welfare. Crucially, because we are interested in comparing the current trade equilibrium vs autarky, we do not need to specify the mechanics of how firms choose their

[25] This methodology is subject to the caveat that the domestic shares may have changed for reasons unrelated to the policy under study. This concern, however, is equally relevant for the methodologies in the papers cited above.

[26] Of course, opening up to trade might induce firms' to engage in other productivity enhancing activities like R&D, in which case innate productivity φ would also increase. Such increases in complementary investments are not encapsulated in (13), which only measures the direct gains from trade holding productivity fixed. If one wanted to disentangle these indirect gains from trade from the direct ones, more structure and data is required. See for example Eslava et al. (2014).

[27] In particular, we need to specify a general equilibrium macroeconomic structure to model the firms' interactions, as we do see Section 3, and also fully specify the extensive margin mechanism, as we do in Section 6.1.

sourcing strategies.[28] In this way, we bypass the complications of the extensive margin of imports discussed above and obtain estimates of the aggregate gains from trade that are common to a wide class of models.

Naturally, the choice of the macro-structure depends on the question at hand. For studies on particular sectors, detailed information on the structure of product markets and demand may be available and treating factor prices as parametric is reasonable. For a study on the economy-wide gains from trade, product market competition can arguably be treated in less detail as information on demand is typically not available for all sectors; general equilibrium considerations are, on the other hand, more important. While we follow the second route, we want to stress that the methodology employed in this section is more general.

Set-up. We consider the following multi-sector CES monopolistic competition environment in the spirit of Caliendo and Parro (ming).[29] There are S sectors of production, each comprised of N_s firms. We treat N_s as fixed.[30] There is a unit measure of consumers who supply L units of labor inelastically and whose preferences are given by:

$$U = \prod_{s=1}^{S} C_s^{\alpha_s} \tag{14}$$

$$C_s = \left(\int_0^{N_s} c_{is}^{\frac{\sigma_s - 1}{\sigma_s}} \, di \right)^{\frac{\sigma_s}{\sigma_s - 1}} . \tag{15}$$

Firm i in sector s produces according to the production technology:

$$\begin{aligned} y_{is} &= \varphi_i l^{1-\gamma_s} x^{\gamma_s} \\ x &= \left(\beta_s x_{Ds}^{\frac{\varepsilon_s - 1}{\varepsilon_s}} + (1 - \beta_s) \, x_{Is}^{\frac{\varepsilon_s - 1}{\varepsilon_s}} \right)^{\frac{\varepsilon_s}{\varepsilon_s - 1}} , \end{aligned}$$

[28]Note that without data on domestic expenditure shares we would need to fully specify the macro model (including the extensive margin) in order to estimate many parameters that are necessary to compute unit costs, even when comparing the current trade equilibrium with autarky. Such estimation would require solving both the profit maximization and the cost-minimization problems *simultaneously* - this is the route taken by Halpern et al. (2011); Gopinath and Neiman (2014); Ramanarayanan (2014); Antràs et al. (2014). The micro data on domestic shares, however, allows us to analyze these two problems separately and avoid structurally estimating the full model.

[29]While this (by construction) shuts down any issues of imperfect pass-through, we would need data on firm-specific prices to put discipline on the extent of pass-through. In principle one could extend our analysis along these lines. See also Dhingra and Morrow (2012) and Fabinger and Weyl (2012) for recent contributions on incomplete pass-through in international trade.

[30]We come back to the case of free entry below.

where $x_{Is} = h_{is}\left(\left[q_c z_{cis}\right]_{c\in\Sigma}\right)$ is the firm's imported input bundle and the domestic bundle is given by:

$$x_{Ds} = \prod_{j=1}^{S} Y_j^{\zeta_j^s} \tag{16}$$

$$Y_j = \left(\int_0^{N_j} y_{vjis}^{\frac{\sigma_j - 1}{\sigma_j}}\, dv\right)^{\frac{\sigma_j}{\sigma_j - 1}} \tag{17}$$

where y_{vjis} is the demand for output of firm v in sector j by firm i in sector s, and $\zeta_j^s \in [0,1]$ with $\sum_{j=1}^{S} \zeta_j^s = 1$ for all s. Thus, firms' production structure is as in Section 2 above but allows for input-output linkages across producers as e.g. in Jones (2011). The structure of roundabout production generated by (16) implies that a trade-induced reduction in a firm's unit cost will spread out to the rest of the economy, thus potentially benefiting domestic producers. We allow for differences in ζ_j^s and α_s to capture the possibility that import-intensive sectors might be of different importance for final good and input production.

Importantly, as in Section 2.2 above, we do not assume any particular mechanism by which the extensive margin of trade is determined. To attain a particular sourcing strategy Σ, the firm needs to spend some resources given by the mapping l_Σ, on which we place no restrictions other than it being in units of labor.

Finally, we assume that the market for differentiated varieties is monopolistically competitive. We also require that trade is balanced and that foreign preferences for the country's locally produced goods are given by (14)-(15).[31]

The macro structure laid out above entails several functional form assumptions and specifies a particular structure of output markets. We point out, however, that most of these assumptions can be replaced by alternative ones and, while the macro gains may change, the methodology to compute them remains intact. We also note that the macro model of this section allows for more flexibility than the existing work on firm-based models of importing. Apart from the generality in the extensive margin and heterogeneity discussed in Section 2.2, it allows for multiple sectors and sector-specific technologies $\left[\varepsilon_s, \gamma_s, \zeta_j^s, \sigma_s, \alpha_s\right]$.[32]

We now formally define an equilibrium in this economy.

Definition 1. *An equilibrium is a set of prices* $(w, P, [p_{Ds}]_s, [p_{is}]_{is})$, *labor allocations* $([l_{is}, l_{\Sigma is}]_{is})$, *differentiated product quantities* $([y_{is}]_{is})$, *domestic and international input demands* $\left(\left[[y_{vjis}]_{vj}, [z_{cis}]_c\right]_{is}\right)$, *sectoral output aggregates* (C_s), *foreign demand* (Y_{ROW}) *and sourcing strategies* $([\Sigma_{is}]_{is})$ *such that:*

1. *Prices* $[p_{is}]_{is}$ *and input demands* $\left[l_{is}, [y_{vjis}]_{vj}, [z_{cis}]_c\right]_{is}$ *solve firms' maximization problem,*
2. *Consumers maximize utility given by* (14) *and* (15),

[31] This is equivalent to assuming that the economy exports a final consumption good produced according to (14). We use the two representations exchangeably.

[32] The reason is that, conditional on firms' unit costs, we can characterize the aggregate economy analytically and do not have to estimate the parameters governing output markets simultaneously with the ones governing import behavior.

3. *The aggregate and sectoral price indices P, $[p_{Ds}]_s$ are consistent with firms' prices $[p_{is}]_{is}$,*

4. *Firms sourcing strategies $[\Sigma_{is}]_{is}$ are consistent with the particular mechanism how the extensive margin of trade is determined,*

5. *The resource loss of firms' extensive margin (if any) $l_{\Sigma is}$ is consistent with firms' sourcing strategies Σ_{is},*

6. *Trade is balanced: $PY_{ROW} = \sum_{s=1}^{S} \int_0^{N_s} p_{cs} z_{cis} di$,*

7. *Labor and good markets clear.*

An equilibrium requires firms' sourcing strategies Σ_{is} to be optimal given the specific structure imposed for the extensive margin. In equilibrium, a firm's sourcing strategy depends on other endogenous variables, such as the prices of the domestic inputs (p_{Ds}), which in turn depend on the sourcing strategies of all other firms as these affect their marginal costs. Hence, finding the equilibrium Σ_{is}'s entails solving a fixed point which, depending on the specifics of the extensive margin, may be a computationally difficult task. However, given the equilibrium sourcing strategies $[\Sigma_{is}]_{is}$ there are unique accompanying equilibrium domestic shares $s_{D,i}$ which fully summarize the unit costs.[33] Using the observed distribution of firms' domestic expenditure shares is hence particularly useful to characterize the aggregate implications of input trade - in particular, the *aggregate gains from trade.*

Proposition 3. *Let P be the domestic price index in the trade equilibrium and P^{Aut} the price index in autarky. We define the aggregate gains from input trade as the reduction in the domestic price index relative to autarky, i.e. $G \equiv ln\left(P^{Aut}/P\right)$. Then,*

$$G = \sum_{s=1}^{S} \gamma_s \alpha_s \Psi_s + \sum_{s=1}^{S} \alpha_s \Lambda_s, \tag{18}$$

where

$$\Lambda_s = \frac{1}{1-\sigma_s} ln\left(\int_{i=0}^{N_s} \omega_i s_{D,i}^{\frac{\gamma_s}{1-\varepsilon_s}(1-\sigma_s)} di\right) \tag{19}$$

is observable from the micro-data as $\omega_i = \frac{va_i}{\int_i^{N_s} va_i di}$ denotes firm i's share in value added and $[\Psi_s]_s$ are determined from the system of equations

$$\Psi_s = \sum_{j=1}^{S} \zeta_j^s \Lambda_j + \sum_{j=1}^{S} \gamma_j \zeta_j^s \Psi_j. \tag{20}$$

Alternatively,

$$G = \alpha' \Gamma \left(\mathcal{I} - \Xi \times \Gamma\right)^{-1} \Xi \Lambda + \alpha' \Lambda \tag{21}$$

[33]Note that the converse is *not* true, i.e. there might be different sourcing strategies leading to the same domestic share.

where $\Lambda = [\Lambda_1, \Lambda_2, ..., \Lambda_S]'$, $\Xi = \left[\zeta_j^s\right]$ *is the* $S \times S$ *matrix of production interlinkages,* α *is the* $S \times 1$ *vector of demand coefficients,* $\mathcal{I}$ *is an identity matrix and* $\Gamma = diag\,(\gamma)$*, where* γ *is the* $S \times 1$ *vector of input intensities.*

Proof. See Appendix. □

Proposition 3 is the main result of the paper. It shows very concisely how the information contained in the micro-data on domestic spending and value added is so powerful. Conditional on a set of parameters, the aggregate gains from input trade can be directly read off from the data and hold true under the quite general assumptions laid out in Section 2.2. In particular, regardless of how the extensive margin is determined and how the distribution of input quality varies by firm, the aggregate gains from input trade are given by the expression in (18) above. In this way, all firm-based models in the class considered in Sections 2.2 and 3 feature exactly the same gains conditional on the micro-data.[34] Examples of recent contributions that fit into Proposition 3 include Gopinath and Neiman (2014); Antràs et al. (2014); Halpern et al. (2011) and Ramanarayanan (2014).

To better understand Proposition 3, it is instructive to consider two special cases. Consider first the case of a single sector economy. Expressions (18) and (20) then imply that

$$G = \gamma\Psi + \Lambda = \frac{\Lambda}{1-\gamma}. \tag{22}$$

That is, the gains from input trade are simply given by the aggregate unit cost reduction Λ, inflated by $1/\left(1-\gamma\right)$ to capture the presence of roundabout production. Consider next the case with multiple sectors and a symmetric input-output matrix in the sense that $\zeta_j^s = \zeta_j^k = \zeta_j$. It can then easily be verified that

$$G = \sum_{s=1}^{S}\left(\alpha_s + \zeta_s \frac{\sum_{s=1}^{S}\gamma_s\alpha_s}{1-\sum_{j=1}^{S}\gamma_j\zeta_j}\right)\Lambda_s. \tag{23}$$

That is, the aggregate gains from trade are an appropriately weighted average of the sectoral "productivity" gains Λ_s. This is simply a consequence of the CES structure for consumers and producers. The weights combine the importance of materials (γ_s), which determine the strength of roundabout production, and the importance of the different sectors for consumers (α_s) and producers (ζ_s). Note that when all sectors have the same importance for consumers and intermediate producers, i.e. $\alpha_s = \zeta_s$, the sectoral weights are simply $\alpha_s/\left(1-\overline{\gamma}\right)$, where $\overline{\gamma} = \sum_{j=1}^{S}\gamma_j\alpha_j$ is the average material intensity in production.

Proposition 3, and especially (22) and (23), illustrate how our methodology builds from micro to macro. The first ingredient are the firm-level gains, $s_{Di}^{\gamma_s/(1-\varepsilon_s)}$, which hold under the general

[34] More precisely, Proposition 3 applies to firm-based models of importing with the micro-level restrictions of Section 2, namely (i) CRS technology, (ii) CES import demand, (iii) price taking in input markets, and the macro-level restrictions of Section 3, namely (i) monopolistic competition, (ii) Cobb-Douglas aggregators for the different sectors and CES aggregators for the firms' output, (iii) input-output linkages given by (16). As discussed above, the macro-level assumptions can be easily replaced by alternative ones and were chosen to generate parsimonious expressions for the aggregate gains. In contrast, the micro-level restrictions of Section 2 are crucial for Proposition 3, but at the same time allow for considerable generality as they do not place any restrictions on the extensive margin nor the distributions of qualities and prices which can vary by input and firm in any way.

conditions of Proposition 2 above. Under CES monopolistic competition, these micro gains are then aggregated into a sector-specific statistic, Λ_s, that weights each firm by its value added share. Next, the Cobb-Douglas structure for demand across sectors and input/output linkages determines how to aggregate these sectoral indices Λ_s into a single statistic that measures the economy-wide gains from input trade. We tied into a canonical macroeconomic structure which delivers parsimonious expressions for the aggregate gains. We want to stress, however, that we can derive the aggregate gains under alternative macro structures depending on the particular application.

Proposition 3 also makes clear that, with heterogeneity in firms' import intensity, aggregate data does not fully summarize the gains from trade, and so the results of Arkolakis et al. (2012) do not apply to the models of importing considered here. To see this more precisely, define the aggregate domestic share in sector s as

$$\lambda_s \equiv \int_i^{N_s} \frac{va_i}{\int_i^{N_s} va_i di} s_{D,i} di = \int_i^{N_s} \omega_i s_{D,i} di. \tag{24}$$

We can then write the sectoral gains Λ_s in (19) as

$$\Lambda_s \quad = \quad \underbrace{\frac{\gamma_s}{1-\varepsilon_s} ln\left(\lambda_s\right)}_{\text{Aggregate Data}} \;+ \underbrace{\frac{1}{1-\sigma_s} ln\left(\int_i \omega_i \left(\frac{s_{D,i}}{\lambda_s}\right)^{\frac{\gamma_s}{1-\varepsilon_s}(1-\sigma_s)} di\right)}_{\text{Heterogeneity}}. \tag{25}$$

When domestic spending is equalized across producers, i.e. $s_{D,i} = \lambda_s$, so that the environment is aggregative, we obtain:

$$ln\left(G_s\right) = \frac{\gamma_s}{1-\varepsilon_s} ln\left(\lambda_s\right),$$

which is reminiscent of the results of Arkolakis et al. (2012) and especially Costinot and Rodriguez-Clare (2014): the sector-specific unit cost reductions depend only the share of domestic spending in material purchases and a "trade elasticity" $\gamma_s / \left(1-\varepsilon_s\right)$, which is increasing in the importance of intermediary purchases and decreasing in the substitutability of domestic and imported varieties. However, when there is heterogeneity in firms' import intensity, the aggregate import share is no longer a sufficient statistic for the cost reductions. In particular both the dispersion in domestic expenditure shares and the correlation between domestic spending and firm size matter for the aggregate welfare gains of importing. Importantly, this bias term can be readily computed from the microdata (conditional on parameters). A natural question is therefore whether an "aggregative" trade model in the spirit of Costinot and Rodriguez-Clare (2014) over- or under-estimate the aggregate unit cost reductions. Surprisingly, the answer to this question does *not* depend on the underlying distribution of domestic shares and firm size (conditional on the aggregate data) but only on the parameters γ_s, ε_s and σ_s.

Proposition 4. *Consider the economy above. Then*

$$\Lambda_s > \frac{\gamma_s}{1-\varepsilon_s} ln\left(\lambda_s\right) \;\textit{ if and only if } \sigma_s < 1 + \frac{\varepsilon_s - 1}{\gamma}.$$

Proof. As $\sigma_s > 1$, and $\int_i \omega_i \left(\frac{s_{D,i}}{\lambda_s}\right) di = 1$ by construction of λ_s, $\Lambda_s > \frac{\gamma_s}{1-\varepsilon_s} ln\left(\lambda_s\right)$ if and only if

$$\int_i \omega_i \left(\frac{s_{D,i}}{\lambda_s}\right)^{\frac{\gamma_s}{1-\varepsilon_s}(1-\sigma_s)} di < 1 = \left(\int_i \omega_i \left(\frac{s_{D,i}}{\lambda_s}\right) di\right)^{\frac{\gamma_s}{1-\varepsilon_s}(1-\sigma_s)}. \tag{26}$$

Jensen's inequality therefore implies that (26) holds true if and only if

$$0 < \frac{\gamma_s}{1-\varepsilon_s}\left(1-\sigma_s\right) < 1.$$

□

Propositions 3 and 4 deliver a general characterization of the macroeconomic implications of input trade. While firm-based models of importing are non-aggregative, the availability of micro-data on domestic spending allows us to nevertheless apply the insights of the literature on sufficient statistics for welfare. In particular, we show that domestic spending is a sufficient statistic for the unit cost reduction associated with input trade at the firm level. That this is true under quite general conditions is the content of Proposition 2. We then show how to aggregate the heterogenous effects of input trade using micro-data on domestic spending and value added. The aggregate gains from input trade can be calculated according to Proposition 3. Finally, Proposition 4 provides a condition on the parameters of the model under which the gains implied by aggregate data are biased relative to the true gains. As our microstructure is vastly more general than the existing models of input trade, and the macrostructure is essentially the canonical quantitative model of international trade, our results characterize the gains from input trade for a wide class of models.

4 Empirical Implementation

We now take the framework laid out above to data on French manufacturing firms to quantify the gains from input trade both at the firm and aggregate level. Implementing Propositions 2 and 3 empirically requires a set of parameters. In this section, we deal with the estimation of these parameters. Our approach relies on both micro data and aggregate relationships. In particular, we use the microdata to estimate the parameters governing firms' technologies, namely the material elasticities $\left[\gamma_s\right]_s$, the elasticities of substitution $\left[\varepsilon_s\right]_s$ and the sector-specific demand elasticities $\left[\sigma_s\right]_s$. We identify the input-output structure on the production side $\left[\zeta_j^s\right]_{sj}$ and the aggregate demand parameters $\left[\alpha_s\right]_s$ from aggregate input-output tables. This allows us also to account for non-manufacturing firms and doing so is quantitatively important. While $\left[\zeta_j^s\right]_{sj}$ and $\left[\alpha_s\right]_s$ can essentially be read off from the input-output matrix and we back out $\left[\sigma_s\right]_s$ from revenue-cost margins, we are particularly careful in estimating the technical coefficient related to intermediate inputs in production functions $\left[\gamma_s\right]_s$, as well as the elasticity of substitution $\left[\varepsilon_s\right]_s$. The latter is not only a crucial parameter for the gains from trade, but it is also novel in the literature and has - to the best of our knowledge - not been estimated from firm-level data yet. We then use these estimated parameters to quantify the gains from input trade in Section 5.

4.1 Data

This subsection provides a general overview of the data used in this paper. The main source of information is a firm-level dataset from France.[35] A detailed description of how the data is constructed is contained in Section 8.3 of the Appendix. Because we are interested in the demand for inputs, we restrict the analysis to manufacturing firms. We observe import flows for every manufacturing firm in France from the official custom files. Manufacturing firms account for 30% of the population of French importing firms and 53% of total import value in 2004.[36] Import flows are classified at the country-product level, where products are measured at the 8-digit (NC8) level of aggregation. This means that the product space consists of roughly 9,500 products. Using unique firm identifiers we can match this dataset to fiscal files, which contain detailed information on firm characteristics. The final sample consists of an unbalanced panel of roughly 170,000 firms which are active between 2002 and 2006, 38,000 of which are importers. Table 9 in the Appendix contains some basic descriptive statistics of our data. We augment this data with two additional data sources. First we employ data on input-output linkages in France, which we take from the STAN database of the OECD. This allows us to identify the parameters of the aggregate trade model and has also been used in Caliendo and Parro (ming). Secondly, we use global trade flows from the UN Comtrade Database to measure aggregate export supplies at the origin-product-level. We use this data to construct an instrument for the estimation of the elasticity of substitution ε.

4.2 Identification of $\left[\alpha_s, \zeta_j^s, \sigma_s\right]$

We compute the demand parameters $[\alpha_s]_s$ and the matrix of input-output linkages $\Xi \equiv \left[\zeta_j^s\right]_{sj}$ using data from the French input-output tables on the distribution of firms' intermediate spending and consumers' expenditure by sector.[37] In particular, letting Z_j^s denote total spending on intermediate goods from sector j by firms in sector s and E_s total consumption spending in sector s (net of intermediate purchases by other producers in the economy), our theory implies:

$$\zeta_j^s = \frac{Z_j^s}{\sum_{j=1}^S Z_j^s} \text{ and } \alpha_s = \frac{E_s}{\sum_{j=1}^S E_j}. \tag{27}$$

While this paper is mainly concerned with the importance of input trade in the manufacturing sector, and this is also where we have access to the micro-data, we have to take the non-manufacturing industry into account to make statements about the aggregate economy. We aggregate all non-manufacturing sectors into one residual sector, which we assume is not allowed to trade, and construct its consumption share α_{NM} and input-output matrix ζ_j^{NM} directly from the Input-Output Tables.

Table 1 below contains the results. Column three reports the consumption share α_s for each of

[35]This dataset is not new and has been used in the literature before (e.g. Mayer et al. (2010); Eaton et al. (2004, 2011)). However, these contributions focused almost exclusively on the export side. We have also used this data in our earlier work (Blaum et al., 2013).

[36]In the same year, retail and gross trade account for 53% of importing firms and 39% of total import value.

[37]See Section 9.3 in the Online Appendix for more a detailed description of how we construct the input-output matrix.

Industry	ISIC	α	$\bar{\zeta}$	σ	γ	VA share	λ
Mining	10-14	0.02%	1.50%	2.58	0.33	1.28%	0.90
Food, tobacco, beverages	15-16	9.90%	1.66%	3.85	0.73	15.24%	0.80
Textiles and leather	17-19	3.20%	3.25%	3.35	0.63	3.96%	0.54
Wood and wood products	20	0.13%	2.90%	4.65	0.60	1.67%	0.81
Paper, printing, publishing	21-22	1.37%	3.91%	2.77	0.50	7.96%	0.75
Chemicals	24	2.04%	6.69%	3.29	0.67	12.91%	0.60
Rubber and plastics products	25	0.44%	3.48%	4.05	0.59	5.88%	0.63
Non-metallic mineral products	26	0.24%	2.41%	3.48	0.53	4.54%	0.72
Basic metals	27	0.01%	6.52%	5.95	0.67	2.07%	0.60
Metal products (ex machinery and equipment)	28	0.26%	6.16%	3.27	0.48	9.27%	0.81
Machinery and equipment	29	0.66%	4.15%	3.52	0.62	7.00%	0.69
Office and computing machinery	30	0.43%	2.29%	7.39	0.81	0.35%	0.59
Electrical machinery	31	0.47%	2.46%	4.49	0.60	3.99%	0.64
Radio and communication	32	0.63%	4.27%	3.46	0.62	1.92%	0.64
Medical and optical instruments	33	0.35%	2.58%	2.95	0.49	3.83%	0.66
Motor vehicles, trailers	34	4.31%	2.13%	6.86	0.76	9.99%	0.82
Transport equipment	35	0.37%	2.62%	1.87	0.35	4.72%	0.64
Manufacturing, recycling	36-37	1.79%	0.92%	3.94	0.63	3.42%	0.75
Non-manufacturing		73.39%	40.11%	na	0.41		1

Notes: σ denotes the demand elasticity, which we determine from the observed profit shares in the French microdata. α is the sectoral share in aggregate demand, which is directly available in the Input-Output Tables. γ is the sectoral share of material spending in total costs, which we estimate from the French micro data. $\bar{\zeta}_s = \sum_{j=1}^{S} \zeta_s^j$ is the average intensity, with which sector s is used as an input to production. "VA share" is the sectoral share of value added in manufacturing as observed in the microdata. λ is the aggregate domestic shares, i.e. $\lambda = \sum_{i=1}^{n} s_{D,i} \times va_i$, which in the model is equal to $\frac{M-IM}{M}$, where M is aggregate intermediary spending and IM is total import spending. See Appendix for the details.

Table 1: Structural Parameters by Industry

the 20 sectors in France. The non-manufacturing sectors are important as they account for a large share of the budget of consumers. In column four we report the *average* intensity with which sectors are used as a factor of production in the rest of the economy. Formally, this intensity is given by

$$\bar{\zeta}_j \equiv \frac{1}{S}\sum_{s=1}^{S}\zeta_j^s.$$

The higher $\bar{\zeta}_j$, the higher the spending of other sectors on sector j is. For completeness we also report the full input-output matrix in Section 9.3 of the Appendix.

In the absence of information on firm-specific prices, we use industry-specific average mark-ups to get the demand elasticities $[\sigma_s]_s$. In the model, mark-ups in sector s are equal to $\sigma_s/(\sigma_s-1)$. As in Oberfield and Raval (2014), we identify mark-ups from firms' ratios of revenues to total costs. We include firms' costs of capital and calculate firms' total costs as the sum of material spending, payments to labor and the costs of capital, which we measure as Rk_i, where k_i denotes firm i's capital stock. We take $R = 0.20$ as our benchmark and take sector averages of these mark-ups to identify σ_s. The results are reported in column five of Table 1. Consistent with the literature we find demand elasticities of around 3.

4.3 Estimating $[\varepsilon_s, \gamma_s]$

Of particular importance are the elasticities of substitution ε_s and the intermediate input shares γ_s, as they directly affect firms' unit cost reductions from trade. Our approach builds on production function estimation. To understand our identification strategy, note that the environment of Section 2 implies that firm output can be written as[38]

$$y_{is} = \varphi_i s_{Di}^{-\frac{\gamma_s}{\varepsilon_s-1}} k_i^{\alpha_s} l_i^{\beta_s} m_i^{\gamma_s} \tag{28}$$

where m_i is total material spending by firm i and $\beta_s \equiv 1-\alpha_s-\gamma_s$.[39] After taking into account the optimal allocation of spending between domestic and all foreign varieties, output can be expressed as a function of the quantities of labor and capital, and spending in materials. As quantities of materials are rarely observed, this is a necessary step in most production function estimation exercises. Importantly, the domestic expenditure share appears in expression (28) as a term akin to a productivity shifter. Based on (28), our approach consists of treating the domestic share as an additional input in production and estimating its output elasticity, given by $-\gamma_s/(\varepsilon_s-1)$. In this way, we address the implicit price bias that arises when using expenditure on material inputs as a production function input (De Loecker and Goldberg, 2013).[40]

As is common in the literature, our dataset does not have information on firm-specific prices but only revenues. We bypass this issue by relying on the demand structure assumed in Section 3, and

[38] In this section, we augment the production function considered in Sections 2 and 3 to include capital, i.e. $y_{is} = \varphi_i k^{\alpha_s} x^{\gamma_s} l^{1-\alpha_s-\gamma_s}$. Note that we are abusing notation since α_s was used above to denote the consumer's expenditure in sector s. We will fix this soon.

[39] We omit the constant $(q_D/p_D)^{\gamma}$ from equation (28) for expositional simplicity.

[40] In other words, the domestic share turns out to be the correct deflator for material spending.

express (28) in terms of firm revenue, S

$$ln\left(S_{is}\right)=\delta+\tilde{\alpha_{s}}ln\left(k_{i}\right)+\tilde{\beta_{s}}ln\left(l_{i}\right)+\tilde{\gamma_{s}}ln\left(m_{i}\right)+ln\left(\omega_{i}\right), \tag{29}$$

where the productivity residual ω_i is given by

$$ln\left(\omega_{i}\right) = \frac{1}{1-\varepsilon_{s}}\tilde{\gamma_{s}}ln\left(s_{Di}\right)+\frac{\sigma_{s}-1}{\sigma_{s}}ln\left(\varphi_{i}\right) \tag{30}$$

and $\tilde{\gamma_s}=\frac{\sigma_s-1}{\sigma_s}\gamma_s$ and $\tilde{\alpha_s}$ and $\tilde{\beta_s}$ are defined accordingly.

We use equations (29) and (30) to estimate ε_s and γ_s following three complementary approaches. The first two methods estimate (29) and (30) separately. They only differ in the way in which the the output elasticities $[\alpha_s,\beta_s,\gamma_s]$ are obtained from (29). We consider both a factor shares approach and a proxy method. We then use such elasticities to construct productivity residuals $ln\left(\omega_i\right)$ and use (30) together with data on domestic shares to estimate ε_s. To increase the power of the estimation, we pool firms from all sectors together and estimate a single ε. The third approach treats the domestic share as an additional input and hence estimates all parameters in (29)-(30) simultaneously. In this approach we also allow for sector-specific ε_s.

As a simple and easy-to-implement benchmark, we first consider an approach based on observed factor shares. In particular, the Cobb-Douglas production structure implies that

$$\frac{\sigma_{s}-1}{\sigma_{s}}\gamma_{s}=\frac{m_{i}}{p_{i}y_{i}}. \tag{31}$$

Hence, given a set of estimates for $[\sigma_s]$, we can use (31) to measure $[\gamma_s]$. We can similarly measure $[\alpha_s,\beta_s]$. Next, we construct productivity residuals $ln\left(\omega_i\right)$ from (29) up to an inconsequential constant. Using the estimated coefficients $\tilde{\gamma_s}$ from the previous step, we then use the productivity residuals and data on domestic shares to estimate equation (30).[41]

Clearly, we cannot estimate (30) via OLS as the required orthogonality restriction fails: s_D is not orthogonal to innate productivity φ under most reasonable models of import behavior. In particular, more productive firms are likely to sort into more and different sourcing countries and this variation in the extensive margin of trade will induce variation in firm-specific price indices and hence domestic shares. Hence, we estimate ε from (30) using an instrumental variable strategy. In particular, we follow Hummels et al. (2011) and instrument s_D with shocks to world export supplies, which we

[41]This approach is related to Kasahara and Rodrigue (2008), who use plant level data from Chile and estimate a production function equation similar to (29)-(30). Using the proxy method of Olley and Pakes (1996) and Levinsohn and Petrin (2003) augmented with import status as an additional state variable, they obtain a positive coefficient for the domestic share and conclude that foreign intermediate inputs increase firm productivity. Their point estimates imply that decreasing the domestic share by 100% increases firm productivity by 5.8-27%, depending on the estimator used. Our approach is different in two respects. First, we identify the coefficient for the domestic share from exogenous variation in export supplies. Second, we use the estimated coefficient to back out the trade elasticity, ε, and then use this elasticity together with the *observed* domestic expenditure shares to measure the productivity gains from trade relative to autarky for every firm.

construct from the Comtrade data. More precisely, we construct the instrument

$$z_{it} = \Delta ln \left(\sum_{ck} WES_{ckt} \times s_{cki}^{pre} \right), \tag{32}$$

where WES_{ckt} denotes the total exports for product k of county c in year t to the entire world excluding France, s_{cki}^{pre} is firm i's import share on product k of county c *prior* to our sample, and Δ denotes the change between year $t-1$ and year t. Hence, z_{it} can be viewed as a firm-specific index of shocks to the supply of the firm's input bundle. Movements in this index should induce variation in firms' domestic shares that are plausibly orthogonal to firm productivity. Intuitively, if we see China's exports in product k increasing in year t, French importers that source product k from China will be relatively more affected by this positive supply shock and should increase their import activities. Using this source of variation in import prices at the firm-level, we can identify the elasticity of substitution ε. We estimate (30) in first differences using (32) to instrument the domestic share according to the following specification

$$\Delta ln \left(\hat{\omega}_{ist} \right) = \delta_s + \delta_t + \frac{1}{1-\varepsilon} \times \Delta \tilde{\gamma}_s ln \left(s_{ist}^{D} \right) + x_{ist}' \zeta + u_{ist}, \tag{33}$$

where δ are again the respective fixed effects, x_{ist} are firm-level controls and $\Delta ln\left(\hat{\omega}_{ist}\right)$ and $\Delta\tilde{\gamma}_s ln\left(s_{Dist}\right)$ are the changes in firm residual productivity and domestic shares respectively, which are instrumented by (32). For our baseline results, we define products at the 6-digit level and consider all importing firms in our sample, taking their respective first year as an importer to calculate the pre-sample expenditure shares s_{cki}^{pre}. As stated above, to increase statistical power we estimate a unique ε from (33) by pooling firms from all sectors together.[42]

As an alternative to the factor shares approach, we also employ a proxy method from the production function estimation literature, akin to Levinsohn and Petrin (2012), to obtain the output coefficients in equation (29). We assume labor to be a dynamic input, which seems plausible for the French labor market, and estimate the obtained equation using GMM as in Wooldridge (2009) to arrive at estimates of the vector of coefficients $[\alpha_s, \beta_s, \gamma_s]$. We experiment with the standard Cobb-Douglas specification, as well as a more flexible translog specification where we continue to assume a constant output elasticity for intermediate inputs but allow for second-order terms in capital and labor:

$$ln\left(S_{ist}\right) = \delta_t + \alpha_s ln\left(k_{ist}\right) + \alpha_s^{kk} ln\left(k_{ist}\right)^2 + \beta_s ln\left(l_{ijt}\right) + \beta_s^{ll} ln\left(l_{ijt}\right)^2 + \psi ln\left(k_{ist}\right) ln\left(l_{ist}\right) + \gamma_s ln\left(m_{ist}\right) + ln\left(\omega_{ist}\right) + u_{ist}, \tag{34}$$

The second step is as in the previous approach: we construct productivity residuals $ln\left(\omega_i\right)$ for each firm and estimate ε from (33) using the instrumental approach described above. Hence, if the production function estimation were to give us the same $[\alpha_s, \beta_s, \gamma_s]$ as the factor shares approach, the implied estimate for ε would be numerically identical.

Finally, we estimate firms' production function in a single step with an integrated GMM approach.

[42] To pool firms from different sectors, we use the respective estimated $\tilde{\gamma}_s$ to build a "scaled" domestic shares $\tilde{\gamma}_s ln\left(s_D\right)$.

	Factor shares			2-step GMM					
				Cobb-Douglas			Translog		
	First Stage	ε	N	First Stage	ε	N	First Stage	ε	N
Full sample	-0.019***	2.378***	526,687	-0.017***	1.776***	331,412	-0.016***	1.727***	331,412
	(0.003)	(0.523)		(0.003)	(0.288)		(0.003)	(0.235)	
Importers	-0.010***	2.322**	65,799	-0.008**	1.896**	53,349	-0.008**	1.802**	53,349
	(0.004)	(1.014)		(0.003)	(0.850)		(0.003)	(0.735)	

Notes: Robust standard errors in parentheses with ***, **, and * respectively denoting significance at the 1%, 5% and 10% levels. The first stage column refers to the estimation of (33) with the instrument given in (32). We estimate γ_s based on factor shares, as per (31), or on the proxy method used in Levinsohn and Petrin (2012) and Wooldridge (2009). For the latter we report results based on Cobb-Douglas technology (29) and or Translog (34). For the factor share specification, we use data for the years 2002-2006. For the proxy method we use data for the years 2004-2006, as two lagged values are required to build the appropriate instruments for the estimation of the production function. For the 2-step GMM procedure we construct standard errors via bootstrap to take the sampling variation in the generated regressor $\gamma_s \Delta ln\left(s_D\right)$ into account.

Table 2: Estimating the Elasticity of Substitution ε

Instead of treating (29) and (30) as separate estimation equations, we consider:

$$ln\left(S_{is}\right) = \delta + \tilde{\alpha_s} ln\left(k_i\right) + \tilde{\beta_s} ln\left(l_i\right) + \tilde{\gamma_s} ln\left(m_i\right) + \frac{1}{1-\varepsilon_s}\tilde{\gamma_s} ln\left(s_{Di}\right) + \frac{\sigma_s - 1}{\sigma_s} ln\left(\varphi_i\right). \tag{35}$$

Treating this as a regular production function estimation problem with four inputs we again follow Wooldridge (2009) and estimate the four parameters using GMM. We follow the literature in using lagged values of capital, labor and materials to proxy for φ, and two-years lagged values of intermediate inputs as an instrument for current intermediate inputs (the only static input). We use the trade instrument discussed above to account for the endogeneity of firms' domestic shares.

The results of the three estimation approaches for ε are reported in Table 2 and Figure 2 below. For brevity, we report the results concerning the other production function parameters in Section 9.4 in the Online Appendix. Table 2 contains the estimation results for ε using the factor shares approach and the proxy method based on Levinsohn and Petrin (2012) and Wooldridge (2009). For the latter procedure we report both the results based on the Cobb Douglas assumption and the more general translog specification. In the respective first column, we show the first stage relationship between changes in aggregate export supply z_{it} and firms' changes in domestic spending. Reassuringly there is a negative relationship that is statistically significant, i.e. firms whose trading partners see an increase in their aggregate exports also reduce their domestic spending.[43] Turning to the results for ε, we see that the different procedures yield relatively similar results as the estimates of ε lie between 1.7 and 2.4.[44] In particular, the point estimate is essentially unchanged when we estimate the second stage equation only for importing firms. Note however that the standard error increases substantially

[43] The reason why the first stage results for the 2-step GMM procedure are not numerically equivalent is that the estimated material elasticity is different. Recall that the independent variable is $\Delta\tilde{\gamma}_s ln\left(s^D_{ist}\right)$.

[44] This result is consistent with Antràs et al. (2014). In their Eaton and Kortum (2002)-style model of importing, they estimate the Frechet parameter of the efficiency distribution θ. Above, we showed that - as far as the firm-level productivity consequences are concerned - their model has the same implications as the "love-for-variety" models, where $\theta = \varepsilon - 1$. They estimate $\theta \in [1, 1.8]$.

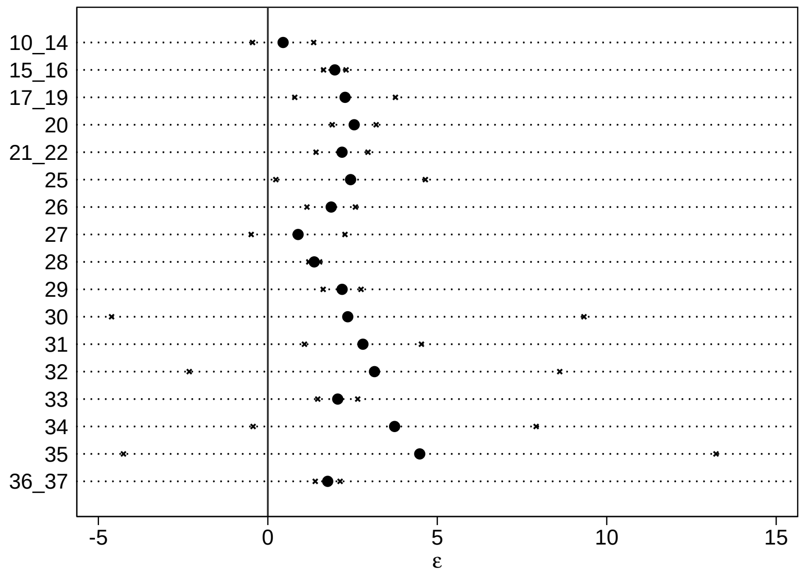

Notes: The figures displays the sector-specific point estimates for ε from the integrated GMM approach based on (35). We also display two standard deviations confidence intervals. The full results are contained in Table 18 in Section 9.4 in the Online Appendix.

Figure 2: Sector-specific ε

as we lose a large amount of data by conditioning on import status.

Finally, Figure 2 turns to the integrated GMM approach based on expression (35). Because we estimate firms' production function for each industry, this procedure gives us sector-specific estimates for ε. We depict both the point estimate and confidence intervals based on two standard deviations. While we lack precision in some industries, the point estimates are mostly in the same ballpark as the pooled results from above.[45] Note that the source of variation is slightly different. While the results based in Table 2 are based on (33), which is estimated in first differences, the 1-step GMM approach treats firms' domestic shares as an explicit input and estimates the production function in levels. It is comforting to see that all these approaches yield consistent results. Conceptually, we prefer the identification strategy in first differences as we find the underlying exogeneity assumptions more plausible. Hence, for the quantitative analysis that follows we take the estimate stemming from the factor shares approach, i.e. $\varepsilon = 2.378$, as the benchmark.[46] While we lock in to this benchmark

[45] Halpern et al. (2011) use a related framework and estimate it on Hungarian micro data. They derive a production function equation analog to (29)-(30), as well as an import demand equation analog to Proposition 7. The main difference with our approach is that they obtain the parameters of their structural model, namely the trade elasticity (analog to ε) and the quality of foreign varieties, by *simultaneously* estimating the production function and import demand equations. Because both of these estimating equations are derived after solving for the extensive margin of trade (i.e. n), they only hold under rather restrictive assumptions, which are akin to our Assumption 1 below. In contrast, we identify ε not within the structural model but by using exogenous variation in input supplies. This allows to estimate ε without having to impose additional assumptions and without having to take a stance how the extensive margin of trade is determined. Halpern et al. (2011) find a much bigger elasticity of substitution between the domestic and foreign variety of 7.3.

[46] Another reason why the factor shares or the two-step GMM approaches may be preferable to the one-step GMM method is that in the latter there are sectors for which we cannot reject $\varepsilon_s < 1$, a feature that leads to the prediction

value for the remainder of the paper, we report confidence intervals for all quantitative results which take into account the sampling variation in this benchmark estimate. Note additionally that our choice of benchmark ε is conservative as far as the magnitude of the gains from trade is concerned, since the unit cost reductions are decreasing in ε. In Section 8.4 in the Appendix we also provide further robustness checks to our estimates of ε, which lead to similar conclusions.[47]

5 Quantifying the Gains from Input Trade in France

With the structural parameters at hand, we now quantify the gains from input trade in France. We proceed as in the theory. In Section 5.1, we follow Proposition 2 and use data on domestic expenditure shares to measure the *distribution* across firms of the unit cost reduction induced by input trade. While domestic share data is sufficient to identify how unequally distributed are the gains from input trade, it is not enough to measure absolute, nor aggregate effects. For that, we need to know how domestic shares correlate with firm size. In Section 5.2 we therefore use Proposition 3 to measure the aggregate price-index gains from input trade in France. There we also exploit our decomposition of the gains from trade in expression (25) to quantify the importance of using the micro data by comparing our results to an analysis based on aggregate data.

5.1 Input Trade and Unit Costs at the Firm Level

Given our estimates of ε and γ_s and the micro data on firms' domestic shares we now implement Proposition 2 and hence consistently estimate the distribution of gains from input trade for the population of French manufacturing firms. Recall that these are simply given by $s_D^{\gamma_s/(1-\varepsilon)}$. We depict these gains in Figure 3 and summarize them in Table 3. We see that there is substantial dispersion in the gains from trade. While the median firm would see its unit cost increase by 10.6% if the economy moved to autarky, firms above the 90th percentile of the distribution would experience losses of 62% or more. As we did not have to make any assumptions on the distribution of fixed costs, the underlying heterogeneity at the country level, the extensive margin of trade, or firms' import demand structure, *any* model with a CES demand system between domestic and imported varieties will arrive at exactly the same conclusions about the distribution of the gains from trade at the micro-level, as long as the model matches Figure 3 and utilizes the same values for γ_s and ε.

We can also learn about firm characteristics that are correlated with such gains. In particular, consider the diagnostic regressions of the form:

$$\frac{\gamma_s}{1-\varepsilon} ln\left(s_{D,ist}\right) = \delta_s + \delta_t + \beta x_{ist} + u_{ist}, \qquad (36)$$

where δ_s and δ_t are industry and time fixed effects and x_{it} are different firm characteristics. To interpret β, consider the homothetic model, which implied that firm-characteristics *only* matter

that all firms in such sector are importers, thus making our theory inconsistent with the data.

[47] While we always estimate a negative effect of domestic shares on productivity in equation (30), in some specifications we lack power to find a significant coefficient. See Table 10 in Section 8.4 of the Appendix.

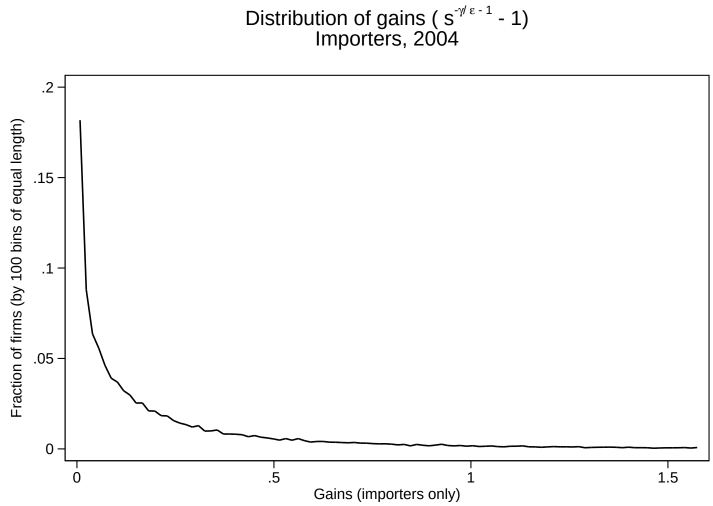

Notes: The Figure displays the distribution of $s_{D,i}^{\gamma/(1-\varepsilon)} - 1$, which is a consistent estimator of the distribution of the gains from input trade holding aggregate prices fixed (see (13)). The data is taken to be the cross-section of importing firms in 2004.

Figure 3: Distribution of the Firm-Level Gains from Input Trade in France

Mean	Quantiles				
	10	25	50	70	90
22.21%	0.64%	2.75%	10.6%	29.07%	61.91%

Notes: The Table reports moments of the empirical distribution of $s_{D,i}^{\frac{\gamma}{1-\varepsilon}} - 1$, which is a consistent estimator of the distribution of the gains from input trade holding aggegate prices fixed (see (13)). The data is taken to be the cross-section of importing firms in 2004.

Table 3: Inequality in the Gains from Input Trade in France

	Dependent variable: Gains from input trade $\frac{\gamma}{1-\varepsilon} ln\left(s_{D,i}\right)$							
ln(Value Added)	0.028***		0.013***	0.005***		-0.008***		-0.029***
	(0.000)		(0.000)	(0.001)		(0.001)		(0.001)
ln(Employment)		0.028***			-0.000			
		(0.000)			(0.001)			
Exporter			0.085***			0.040***		0.024***
			(0.001)			(0.002)		(0.002)
Intl. Group			0.148***			0.138***		0.113***
			(0.003)			(0.003)		(0.003)
ln (Num. Varieties)							0.128***	0.144***
							(0.002)	(0.002)
Sample	Full sample			Importers Only				
Observations	633,240	640,610	633,240	118,799	120,344	118,799	120,344	118,799

Notes: Robust standard errors in parentheses with ***, ** and * respectively denoting significance at the 1%, 5% and 10% levels. All regressions include year fixed effects and 3-digit industry fixed effects. The number of varieties is the number of countries firms source from. A firm is member of a foreign group if at least one affiliate or the headquarter is located outside of France.

Table 4: Cross-Sectional Variation in Gains from Input Trade

via the extensive margin of trade - see (7). As the import price index A is "decreasing" in that a larger sourcing set Σ implies a lower price index for the importing bundle, we can interpret the partial correlations in (36) as reflecting the equilibrium relationship between different firm characteristics, the accompanying sourcing strategy and the resulting unit cost consequences. The results are contained in Table 4 and are intuitive. Bigger firms as measured by either value added or employment see higher gains because they are more likely to endogenously select into more import markets. In column (3) we add two firm characteristics which we expect to be positively related to firms' ability to benefit from international sourcing. We see that both exporters and members of a foreign group have 8.5% and 14.8% higher productivity *through* their import activities. The remaining columns restrict the analysis to the sample of importers. First, note that the correlation between import intensity and size drops substantially. Hence, importing firms are bigger than non-importing firms, as shown in Figure 1 in the introduction, but conditional on importing bigger firms do not source substantially more from abroad.[48] Again, we find larger gains for exporters and members of international groups even conditional on importing.[49] As stressed in Section 2, an important determinant of the gains is the firm's sourcing strategy, which can be measured by the number of international varieties sourced. Columns (4) and (5) show that there is a strong positive correlation between firms' extensive margin of importing and the resulting productivity gains, and that other firm characteristics shrink in importance once the number of varieties is controlled for. In particular, firm-size is now substantially negatively correlated with import spending holding the number of imported varieties fixed. This is

[48]This weak correlation between firm size and import intensity conditional on importing is an important moment for our quantitative model of Section 6 below, as it directly rejects models based on a single source of heterogeneity such as innate productivity.

[49]It is interesting to note that the coefficient on firm size is now negative. This is intuitive. A large firm that decides to neither be an exporter nor a member of a foreign group, is likely to either be disadvantaged on international markets, e.g. through a large fixed cost draw, or see for other reasons small benefits for international activities.

	Manufacturing Sector		Aggregate Economy	
Gains From Input Trade	27.5	*[21.2,35.9]*	9	*[7.1,11.6]*
Aggregate Data	31.4	*[21.3,47.5]*	9.9	*[7.1,14]*
Bias	-3.9	*[-12.3,0.1]*	-0.9	*[-2.6,0]*

Notes: The table reports the reduction in aggregate prices for the manufacturing sector (left panel) and the aggregate economy (right panel). The first row contains the measured gains, taking firm-heterogeneity into account. It is based on (21), with Λ (see (19)) reported in Table 6 and the underlying structural parameters Ξ, γ and α given in Table 1. The second row contains the counterfactual results based on an aggregative model with the identical input-output structure and structural parameters. Hence, we again use (21) but measure the sectoral gains by $ln\left(\lambda^{\gamma/(1-\varepsilon)}\right)$ (see(24)) instead of Λ. The third row report the bias, which is simply the difference. We report 90-10 confidence intervals for all these number in brackets. These are calculated via a bootstrap procedure, which is contained in the Appendix. We estimate the empirical distribution of all statistics, using 200 bootstrap iterations.

Table 5: Aggregate Gains From Trade

intuitive: if a small firm decided to source from a large number of sourcing countries, it is likely that this firm is a proficient importer, which manifests itself in a low share of domestic spending.

5.2 Input Trade and the Aggregate Gains from Trade

We now have all the ingredients required by Proposition 3 to quantify the aggregate price-index gains from input trade. Using firm-level data on domestic expenditure shares and value added, together with the elasticities $[\gamma_s, \sigma_s]$ and ε estimated in Sections 4.2 and 4.3 above, we can construct the sector-specific aggregate cost reductions Λ_s given in (19). With Λ_s and the demand and input-output parameters $\left[\alpha_s, \zeta_j^s\right]$ estimated in Section 4.2 above, we can calculate the aggregate gains from input trade by solving the linear system in (21). The results are contained in Tables 5 and 6. Consider first Table 5 which reports the gains from input trade for the aggregate economy and the manufacturing sector as well. The first row reports the aggregate gains implied by the French micro-data and Proposition 3: these amount to 9% for entire economy. That is, French consumer prices would be 9% higher if French producers were forced to source their inputs domestically. We also see that the price gains are substantially larger in the manufacturing sector, where they amount to 27.5%. Hence, the reason why the economy-wide gains are relatively moderate lies in the non-manufacturing sector, which experiences only a 3% price reduction and accounts for 70% of consumers' budget as seen in seen Table 7.[50] The second and and third rows in Table 5 compare our estimates to those that arise from an approach based on aggregate data in the spirit of Arkolakis et al. (2012) - see (25) in Section 3. An analysis of the French data through the lens of a model with a representative

[50]The economy-wide gains are a weighted average of the gains in the manufacturing and non-manufacturing sectors

$$G = \alpha_{NM} ln\left(\frac{P_{NM}^{AUT}}{P_{NM}}\right) + (1-\alpha_{NM})\, ln\left(\frac{P_M^{AUT}}{P_M}\right),$$

where

$$ln\left(\frac{P_M^{AUT}}{P_M}\right) = \sum_{j=1}^{S-1} \frac{\alpha_j}{1-\alpha_{NM}} ln\left(\frac{P_s^{AUT}}{P_s}\right)$$

is the change in the price-index in the manufacturing industry. The weights are given by consumers' expenditure shares.

firm implies gains of 9.9% and 31.4% in the entire economy and manufacturing sector, respectively. Hence, such aggregative approach results in an over-estimation of the gains from input trade by 1 percentage point for the entire economy and 4 percentage points for the manufacturing sector.[51]

With the micro-data at hand, we can also quantify our confidence in these results. There are two sources of uncertainty. First, we base our analysis on a large but finite sample. Second, the structural parameters ε, γ_s and σ_s are estimated with error. To quantify the magnitude of this uncertainty, Table 5 reports the 90-10 confidence intervals of the bootstrap distribution of the respective statistics in italics.[52] It is clearly seen that the uncertainty about the parameters and the sample itself introduces quite a bit of variation in the magnitudes of our objects of interest. The gains from trade in the manufacturing sector lie between 21 and 36% with 80% probability, and the gains for the entire economy lie between 7.1% and 11.6%.[53] It is also seen that an aggregative approach will almost surely lead to an over-estimation of the gains from input trade given the underlying data. A graphical depiction of this sampling uncertainty is contained in Figure 4. This figure also shows that the distribution of the bias has the majority of its mass on negative numbers.

Table 6 contains the results at the sectoral level. For each sector we report four statistics and their accompanying confidence intervals. The third column reports the sectoral price gains P_s^{AUT}/P_s which measure the total change in the price of the output bundle of that sector. Columns one and two decompose these price gains into the direct unit cost reductions from firms in sector s sourcing internationally, Λ_s, and the indirect gains whereby domestically sourced inputs become cheaper as firms in other sectors engage in trade, $p_{s,D}^{AUT}/p_{s,D}$.[54]Note that the price changes in domestic inputs $p_{s,D}^{AUT}/p_{s,D}$ vary across sectors due to heterogeneity in input-output linkages - as some sectors use relatively open sectors more intensively, they benefit more from input trade as their upstream suppliers experience larger unit cost reductions. Table 6 shows a large degree of heterogeneity across sectors. While e.g. the textile sector experiences a direct unit cost reduction of 31%, these are much more modest for producers in the wood industry (8%). This heterogeneity stems mostly from the observed import behavior in that textile producers source a larger share of their input abroad.[55] This is also true for the total changes in the sector prices, reported in column 3. While prices for

[51]These represent biases in the estimates of the gains from trade of 10% and 14% for the entire economy and manufacturing sector, respectively.

[52]We explain the details of the bootstrap procedure in Section 9.5 of the Appendix. A sketch of the procedure is as follows. For each bootstrap iteration we construct a new sample of the French manufacturing sector by drawing firms from the empirical distribution with replacement. We then simply redo the entire analysis of Sections 4.2 and 4.3, i.e. we reestimate our micro-specifications to recover the structural parameters. Finally for each iteration we recalculate the gains from Proposition 3. This provides us an empirical distribution of every statistic of interest.

[53]Given the large sample size, most of the uncertainty stems from the variation in the structural parameters and not from the resampling of firms. We display the bootstrap distribution of both the underlying parameters and the resulting sectoral productivity gains in Figure 7 in the Appendix.

[54]Note that these objects are related via

$$ln\left(\frac{P_s^{AUT}}{P_s}\right) = \gamma_s ln\left(\frac{p_{s,D}^{AUT}}{p_{s,D}}\right) + \Lambda_s,$$

as shown in (58) in the Appendix.

[55]Note that the share of intermediate inputs is almost the same for textile and wood producers (see Table 1) and that they - by construction - share the same elasticity of substitution ε.

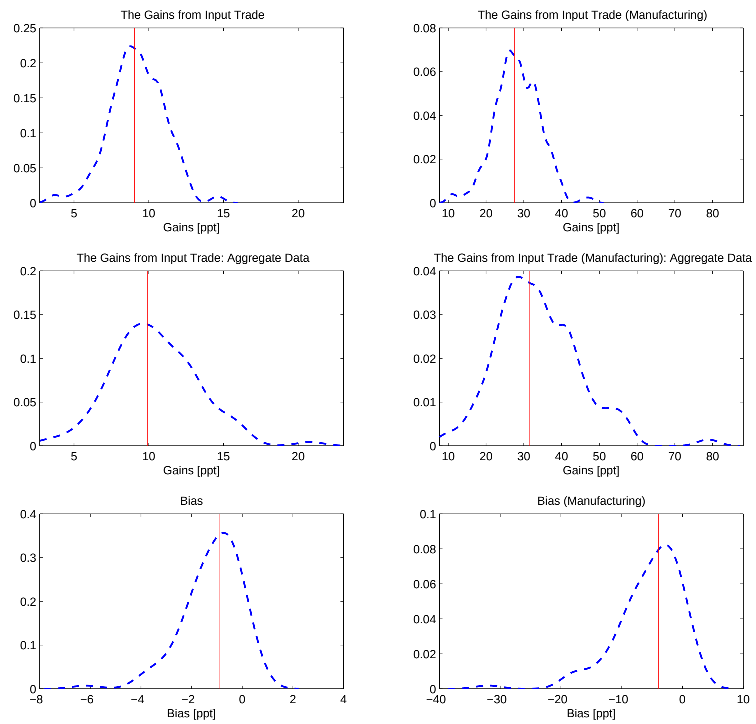

Figure 4: Bootstrap Distribution of Gains and Bias

Industry	ISIC	Λ		Input Prices		Sectoral Price Gains		$\frac{\gamma_s}{1-\varepsilon_s} ln(\lambda_s)$	
Mining	10-14	3.0	*[1.8,4.2]*	14.9	*[11.1,19.2]*	7.8	*[5.2,10.3]*	2.5	*[1.6,3.6]*
Food, tobacco, beverages	15-16	11.1	*[7.5,14.6]*	8.4	*[6.2,10.6]*	17.8	*[12.4,23.4]*	12.6	*[7.8,18.2]*
Textiles and leather	17-19	31.1	*[24.2,39.9]*	31.4	*[24.3,40.3]*	55.6	*[42.4,74]*	31.9	*[22.4,46.9]*
Wood and wood products	20	8.2	*[6.4,10.5]*	9.6	*[7.4,12.1]*	14.4	*[11.1,18.2]*	9.6	*[6.7,13.7]*
Paper, printing, publishing	21-22	12.2	*[9,16]*	14.5	*[10.9,18.7]*	20.1	*[14.7,26.5]*	11.0	*[7.7,15.4]*
Chemicals	24	27.2	*[20.1,36.4]*	21.6	*[16.1,28.2]*	45.1	*[32.7,60.7]*	28.1	*[18.7,41.8]*
Rubber and plastics products	25	20.1	*[14.3,26.5]*	27.3	*[20.2,36]*	38.4	*[27.5,50.9]*	21.5	*[13.9,31]*
Non-metallic mineral products	26	13.4	*[9.6,17.9]*	12.7	*[9.7,16.3]*	20.8	*[15.3,27.4]*	13.3	*[9,19]*
Basic metals	27	21.8	*[16.3,27.7]*	21.5	*[16.4,27.3]*	38.9	*[28.2,50.2]*	28.8	*[19,41.6]*
Metal products (ex machinery and equipment)	28	8.2	*[6.2,10.5]*	20.5	*[15.5,26.2]*	18.3	*[13.8,23.5]*	7.7	*[5.5,10.8]*
Machinery and equipment	29	17.6	*[12.8,23.2]*	20.0	*[15,25.7]*	31.7	*[23,41.6]*	18.2	*[12.2,26.2]*
Office and computing machinery	30	20.4	*[15.4,25.5]*	25.2	*[18.3,32.1]*	44.6	*[31.9,57]*	37.0	*[22.4,60.3]*
Electrical machinery	31	19.8	*[14.6,25.6]*	23.9	*[17.7,30.6]*	36.1	*[26.4,46.6]*	21.6	*[14.8,30.7]*
Radio and communication	32	21.5	*[13.1,31.1]*	23.3	*[16.6,30.5]*	38.5	*[23.5,54.8]*	22.1	*[12.5,36.1]*
Medical and optical instruments	33	17.9	*[12.8,23.4]*	20.4	*[15.1,26.2]*	29.2	*[21.1,38.3]*	15.9	*[10.7,22.5]*
Motor vehicles, trailers	34	6.2	*[3.2,16.4]*	21.7	*[17,29.3]*	23.3	*[17.4,39]*	11.2	*[6.1,24.3]*
Transport equipment	35	15.3	*[10.5,22]*	19.9	*[14.5,27.2]*	22.9	*[16,33.2]*	11.8	*[7.9,18.2]*
Manufacturing, recycling	36-37	12.9	*[9.7,16.3]*	19.0	*[14.5,24]*	26.0	*[19.2,33.4]*	14.1	*[9.5,20.4]*
Non-manufacturing		0.0	*[0,0]*	7.5	*[5.7,9.4]*	3.0	*[2.3,3.8]*	0.0	*[0,0]*

Notes: Λ, the direct unit costs reductions from international sourcing relative to autarky, is calculated according to (19). The reductions in the price of domestically sourced intermediary inputs, $\frac{p_{s,D}^{AUT}}{p_{s,D}}$, is contained in column 2. Column 3 contains the full change in sector prices relative to autarky. Column 4 considers the case of an aggregative approach, whereby the direct unit costs reductions from international sourcing relative to autarky would be given by $\frac{\gamma_s}{1-\varepsilon_s} ln(\lambda_s)$, where λ_s is the aggregate domestic share in sector s (see (24)).

Table 6: The Gains From Input Trade across Sectors

textile products would be 56% higher if textile producers were not allowed to source their inputs from abroad, the effects would be more modest for other industries in France.

The last column of Table 6 contains the sector-specific price reductions that arise from a representative firm model.[56] In line with the results of Table 5, in 12 of the 18 manufacturing sectors the gains based on aggregate data are higher. The reason for this sectoral asymmetry is contained in Proposition 4: given the empirically observed σ_s, ε and γ_s in Table 1, the bias in the aggregate data is positive in most sectors. Note also that the bias can be quite substantial. Consider for example the office and computing machinery sector, where an analysis based on an aggregative model would imply price changes of 37%, while the exact firm-based formula tells us that the direct unit cost reduction due to trade amounts to only 20%.

Importantly, there is a second source of bias that arises when using an aggregative approach which pertains to the "correct" elasticity of substitution ε.[57] While we explicitly treat ε as a production function parameter and estimate it from micro-data, aggregative models often estimate ε as a trade elasticity from observed trade flows. While there is a large literature concerning this particular parameter[58], most aggregative approaches find estimates that are larger than our preferred estimate of 2.37.[59] Costinot and Rodriguez-Clare (2014) for example use a trade elasticity of four as their benchmark value. As the implied gains from trade are decreasing in the elasticity of substitution, such choice would lead to substantially smaller gains from trade in an aggregative environment. In Section 9.6 of the Online Appendix, we redo the sectoral analysis of Table 6 for a range of values of ε spanning the estimates from the aggregative literature. Moving to $\varepsilon = 4$, for example, tends to reduce the gains from trade of the aggregative approach by 50%. We conclude that the bias in the estimates of the gains from trade arising from the use of an aggregative approach can be substantial in magnitude.[60]

Table 6 is also helpful to understand why the economy-wide gains are only a third of the gains in the manufacturing sector. The last row in Table 6 shows that aggregate non-manufacturing prices fall only by 3%. This relatively modest cost reduction in the non-manufacturing sector stems from two sources. Not only does this sector lack any *direct* cost reduction, as it does not engage in input trade by assumption, but it also relies little on the manufacturing sectors as its share of intermediate inputs γ^{NM} is only 40%. Hence, it sees its input prices decrease by only 7.5%.

Finally, to assess the importance of interconnections between sectors, we consider the case with no cross-industry input-output linkages where each sector uses only its own products as inputs.[61] In

[56] More precisely, this last column contains the analog to column 2 within an aggregative model, i.e. $\frac{\gamma_s}{1-\varepsilon_s} ln\left(\lambda_s\right)$ where λ_s is the aggregate domestic share in sector s - see (25) in Section 3.

[57] We thank Pol Antràs for suggesting this exercise to us.

[58] See e.g. Simonovska and Waugh (2013, 2014).

[59] Recall that our benchmark was chosen conservatively, in that all other estimates of ε in Table 2 were smaller.

[60] See also Ramanarayanan (2014), who compares a firm-based model with an aggregative economy in calibration exercise.

[61] In this case, the matrix of input-output linkages is given by $\zeta_j^s = 0$ for $j \neq s$ and $\zeta_j^j = 1$.

this case, we find a point estimate for the aggregate gains from trade of[62]

$$G = \sum_{s=1}^{S} \alpha_s \frac{\Lambda_s}{1-\gamma_s} = 12\%.$$

That is, shutting down input trade would reduce the aggregate price index by 12%. Compared to the actual gains of 9%, the economy without interlinkages over-estimates the aggregate gains by 33%. The reason is that the non-manufacturing sector is not only important for final consumers but also as a provider of inputs to other manufacturing firms. And as the non-manufacturing sector is not a direct beneficiary of input trade, such linkages actually dampen the aggregate gains from input trade.

6 The Effect of Input Trade on Welfare

So far we have considered gains from trade that arise from the reduction in prices faced by consumers. Under the general model of Sections 2-3 we were able to measure these gains from input trade without specifying a particular extensive margin mechanism. However, while these price-index gains are an important component of welfare, they do not take into account the resources spent by firms to attain their equilibrium sourcing strategies. Measuring the full effect of input trade on welfare therefore requires taking a stand on how the extensive margin of trade is determined. To this end, we consider a model where foreign sourcing is limited by the presence of fixed costs.

We start by deriving an expression for full consumer welfare in the context of the general model of Section 3, and hence valid for any extensive margin mechanism. We then impose additional structure in Section 6.2 by assuming that foreign sourcing is subject to fixed costs. In Section 6.3 we calibrate the structural model to the French micro data. We use the calibrated model to measure the full welfare gains from input trade, as well as to assess the importance of the micro data. In particular, we quantify the biases in the estimates of the gains from trade that arise when domestic share data is not available.[63]

6.1 A Formula for Welfare

The following proposition characterizes full consumer welfare in the context of the macroeconomic environment of Section 3. For expositional simplicity, we consider a one-sector version of the model and leave the analysis with multiple sectors to Section 9.7 in the Online Appendix. We assume that the costs associated with firms' extensive margin choices are in units of labor.

Proposition 5. *Consider a single-sector version of the above setup and let* $[l_{\Sigma i}]_i$ *denote the resource loss associated with firms' equilibrium sourcing strategies* $[\Sigma_i]_i$*. Let* W *and* W^{Aut} *denote the total*

[62]Note that (20) implies that $\Psi_s = \Lambda_s + \gamma_s \Psi_s$, so that (18) yields $G = \sum_s \alpha_s \frac{\gamma_s}{1-\gamma_s}\Lambda_s + \sum_s \alpha_s \Lambda_s = \sum_s \alpha_s \frac{\Lambda_s}{1-\gamma_s}$.

[63]See Gopinath and Neiman (2014) for an example where micro data on domestic expenditure shares is not available.

welfare in the trade equilibrium and autarky respectively. Then

$$\frac{W}{W^{Aut}} = \frac{P^{Aut}}{P} \times \left(\frac{L - \int_i^N l_{\Sigma_i} di}{L} \right). \tag{37}$$

Proof. See Section 9.7 in the Online Appendix, where we also prove Proposition 5 for the general multi-sector case. □

Proposition 5 provides an intuitive expression for welfare in our general environment. *Regardless* of the precise microstructure of the trading environment or the particulars of the extensive margin, the full welfare gains from input trade are given in (37). We see that the welfare gains from trade have two components. First, there is the reduction in consumer prices associated with input trade, which is captured by the term $P^{Aut}/P > 1$ and was the focus of Sections 3-5 above. Second, there is the resource loss associated with attaining the equilibrium sourcing strategies, captured by the second term in expression (37). As $l_{\Sigma i}$ are unobservable, quantifying the full welfare consequences of input trade requires imposing more structure into the model of Section 3 by taking a stand on how firms acquire their foreign inputs. This is where we turn now.

6.2 A Model Of Fixed Costs

To operationalize Proposition 5, we need to commit to a particular mechanism by which firms decide their sourcing strategy. In this section, we go back to our simple example of Section 2.1 and make the canonical assumption that engaging in international trade entails a payment of fixed costs. In particular, we assume that sourcing an input from country c entails a fixed cost f_c which can be firm-specific and is denominated in units of labor. There is also a fixed cost f_I to start importing. The firm's maximization problem is therefore given by (5) and (6).

As discussed above, the choice of the optimal sourcing strategy in the model with fixed costs can be computationally challenging. Hence, we seek a structure, which is tractable, but sufficiently rich to capture important features of the micro-data. Concerning tractability, we assume that fixed costs do not vary by country, $f_c = f$. This condition bypasses the complexity of firms' extensive margin problem, as sourcing countries can be ordered by their price-adjusted quality.[64] As for the potential to match the firm-level data, we allow fixed costs to vary by firm. If productivity was the single source of heterogeneity, the model would generate a one-to-one assignment between firm size and import intensity - a feature that is counterfactual in the French data as shown in Figure 1 above. The following Assumption contains these important conditions as well as other assumptions including parametrizations of the distribution of foreign quality and prices.

Assumption 1. *Consider the environment above and assume the following:*

1. *The fixed cost of sourcing is constant across sourcing countries but it can vary by firm, i.e.* $f_{ci} = f_i$ *where* f_i *and* φ_i *follow some joint distribution.*

[64] This assumption is also made by Gopinath and Neiman (2014) and Halpern et al. (2011). See also the discussion in Antràs et al. (2014), who characterize import behavior without this restriction.

2. *The fixed cost of importing f^I is constant across firms.*

3. *The distribution of quality is Pareto, i.e.* $G(q) = \Pr(q_c \leq q) = 1 - \left(\frac{q_{\min}}{q}\right)^\theta$ *where θ satisfies* $(\rho - 1)(1 - \nu) < \theta$.

4. *Foreign prices p_c are dependent on quality and are given by $p_c = \alpha q_c^\nu$.*

5. *Import demand is homothetic, i.e.* $\eta(q, \varphi) = q$.

Assumption 1 contains sufficient restrictions for a tight characterization of firms' extensive margin. The essential assumption is the first one. When fixed costs do not vary by country, the firm selects its sourcing countries purely based on price-adjusted quality, for a given mass of countries sourced. More precisely, if country c with price-adjusted quality q_c/p_c is an element of Σ so are all countries c' with $q_{c'}/p_{c'} > q_c/p_c$. Thus the firm's sourcing strategy reduces from an entire set to a scalar: Σ can be summarized by a quality cutoff $\overline{q}$.[65] See also Blaum et al. (2013) where this point is discussed in more detail and Antràs et al. (2014) who offer a very insightful analysis in a related model.[66]

Under a Pareto distribution for quality, the firm's import price index, given by equation (7) above, takes a convenient power form:

$$A(n) = zn^\eta, \tag{38}$$

where n is the share of countries the firm sources foreign inputs from and z and η are "auxiliary" parameters which depend on the underlying parameters governing import prices (α, ν), the distribution of quality $(q_{\min}, \theta)$ and the elasticity of substitution of foreign varieties ρ.[67] Thus, the underlying structure of the import environment matters for the firm's problem *only* through (z, η). This property is useful for quantifying the model, as it implies that knowledge of the deep parameters $(\alpha, q_{\min}, \theta, \rho, \nu)$ is irrelevant for all aggregate outcomes as long as (z, η) are known.[68]

[65] While in principle the cutoff is in terms of price-adjusted quality q/p, item (4) in Assumption 1 allows us to express it in terms of quality q. As long as $\nu < 1$, we have that $q/p = \frac{1}{\alpha} q^{1-\nu}$ is increasing in quality. Otherwise, one can reorder countries so that low-quality countries have the highest price-adjusted quality.

[66] The remaining four restrictions in Assumption 1 could be dispensed with and replaced by alternative ones. While we could allow for firm-heterogeneity in f^I, this extra degree of freedom in addition to the firm heterogeneity in f does not seem to be essential. Our restriction on prices is a tractable way to link the distribution of qualities to the distribution of price-adjusted qualities (which is what importing firms care about). That qualities are distributed Pareto is a convenient parametrization - but we could have considered other specifications. Finally, we consider the homothetic model. We do so not only because this is the benchmark model in the literature, but also because disciplining the functional form for $\eta(q, \varphi)$ is not straight-forward.

[67] See Section 9.8 in the Appendix for details and a derivation of (38).

[68] This however, does not mean that the degree of quality heterogeneity (θ) or the substitutability of inputs (ρ) do not have a well defined role in shaping firms' import demand. For example, it can be easily shown that diversity and substitutability increase import productivity z (i.e. $\frac{\partial z}{\partial \theta} < 0$ and $\frac{\partial z}{\partial \rho} > 0$) and that diversity and substitutability are complements (i.e. $\frac{\partial^2 z(E[q], \theta, \rho)}{\partial \theta \partial \rho} < 0$), if and only if $(\rho - 1)(1 - \nu) > 1$. When $(\rho - 1)(1 - \nu) > 1$, the firm is effectively *risk loving* and values diversity. As importing is effectively an option (as only the best countries are selected as sourcing countries), more variance in the unconditional distribution will increase the benefit of importing. Such gains from diversity however are only available, if inputs are sufficiently substitutable, i.e. the higher ρ the more can firms leverage such quality differences. Hence, import quality or the technology, which with imports are combined, do affect import demand. However, for given estimates of η and z, they do not change the researchers' conclusion on firms' import demand or the aggregate gains from trade. See also Melitz and Redding (2012).

The firm's optimal domestic share, given by (10) above, becomes under (38):

$$s_D(n) = \frac{1}{1+\left(\frac{p_D}{q_D} z n^{\eta}\right)^{\varepsilon-1}} \tag{39}$$

In this context, firms' extensive margin of importing can be easily characterized. As importing is subject to fixed costs, firms need to trade-off the import-induced reduction in variable costs vs payment of fixed costs. As innate productivity φ and the endogenous unit costs reductions through input trade $s_D^{\frac{\gamma}{1-\varepsilon}}$ are complements, there is a fixed-cost specific productivity cutoff $\overline{\varphi}(f)$ above which firms select into importing. Conditional on importing, the optimal share of countries firms import from is given by the solution to a first order condition associated to the profit function (5). Proposition 7 in Section 9.8 of the Online Appendix gives the full characterization of the extensive margin. We now have all the ingredients in place to estimate the full welfare effects of input trade along the lines of Proposition 5.

6.3 Calibration And Welfare Gains

A natural requirement for the structural model is for it to match the price-index gains from trade which, as argued above, can be read off directly from the micro data. It follows from Proposition 3 that this is ensured by disciplining the model with the joint distribution of value added and domestic expenditure shares. Our strategy is as follows. We consider a one-sector version of the model. While we do this for simplicity, we note that analysis that follows can be extended to a multi-sector setting. We estimate a number of important parameters directly from the micro data. Note that we have already estimated ε,γ and σ in Sections 4.2 and 4.3 above.[69] Thus, the only remaining parameter that needs to be estimated is η, which determines the import-price index (38) and hence the demand for foreign varieties.[70] We estimate η using the cross-sectional relationship between firms' extensive margin of trade and their domestic shares given by (39). Next, we parametrize the distribution of innate productivity and fixed costs as a joint log-normal distribution

$$\begin{pmatrix} ln(\varphi) \\ ln(f) \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \mu_\varphi \\ \mu_f \end{pmatrix}, \begin{pmatrix} \sigma_\varphi^2 & \rho\sigma_\varphi\sigma_f \\ \rho\sigma_\varphi\sigma_f & \sigma_f^2 \end{pmatrix}\right), \tag{40}$$

where ρ controls the correlation between productivity and fixed costs. We normalize μ_φ and calibrate the rest of the parameters in (40) to match salient features of the joint distribution of sales and domestic expenditure shares, namely the dispersion of each of these variables as well as their correlation. Finally, we choose the fixed cost of being an importer f_I to match the share of importers in the French population.

[69] Sections 4.2-4.3 provide estimates of σ and γ by sector. In this Section we use value-added weighted averages of these sectoral estimates.

[70] While the import-price index (38) also depends on z, it turns out that this parameter is not required for the calibration of the model.

Estimation of η. As can be seen in expression (39), η controls the sensitivity of the domestic share to changes in the extensive margin of trade (n). We use this property to identify η, for a given value of ε. At this point, we need to take a stand on what the counterpart of n is in the data. We focus on firms' choices of the number of import countries within products, which we displayed in Figure 1. We do so for a number of reasons[71] but we note that the analysis that follows can be done under alternative interpretations of n.

Equation (39) predicts a log-linear relationship between n and the term $(1 - s_D)/s_D$, with a slope given by η. We therefore run the following regression:

$$ln\left(\frac{1 - s_{D,ist}}{s_{D,ist}}\right) = \delta_s + \delta_t + \delta_{NK} + \eta\,(\varepsilon - 1)\,ln\,(n_{ist}) + u_{ist}, \tag{41}$$

where n_{ist} denotes the average number of countries per product sourced and δ_{NK} contains a set of fixed effects for the number of products sourced.[72] We measure products at the 8-digit level. Section 9.9 in the Online Appendix contains the results for a variety of specifications. Our preferred specification yields a value of η is 0.382 that is precisely estimated.

Calibration. To calibrate the five remaining structural parameters $(\mu_f, \sigma_f, \sigma_\varphi, \rho, f_I)$ we target the following five moments: (i) the aggregate domestic share of the French manufacturing sector[73], (ii) the share of importing firms, (iii) the standard deviation of log domestic shares, (iv) the standard deviation of log sales and (v) the correlation between log sales and log domestic shares. While all parameters are calibrated jointly, the average level of fixed costs (μ_f) controls mostly the aggregate domestic share, the fixed cost of importing (f_I) is mostly identified from the share of importers and the dispersion in fixed costs (σ_f) and productivity (σ_φ) from the dispersion in domestic shares and sales, respectively. Finally, the correlation between productivity and fixed costs (ρ) is disciplined by the correlation between sales and domestic spending. By explicitly targeting the economy's aggregate domestic share, we can compare our results to those of an aggregative approach where the moments from the micro-data are not used.

Table 7 contains the results of the calibration[74] and a summary of the parameters previously

[71] First, the choice of the number of products sourced may be determined to a large degree by technological considerations, while the demand for multiple supplier countries within a given product category might plausibly stem from love-for-variety effects as modeled in Section 2. Second, while most firms source their products from a limited number of countries, the subset of firms that engages intensively in multiple country sourcing accounts for a large share of aggregate imports. For example, Figure 1 above and Blaum et al. (2013) show that about 30% of the firms source their inputs from at least two countries (on average per product), but these firms account for more than 80% of aggregate imports. Finally, we note that this notion of "varieties" is widely used in the literature - see e.g. Broda and Weinstein (2006) and Goldberg et al. (2010).

[72] By including product fixed effects, we identify η solely from the variation stemming from the number of countries *per product*. That is, from firms sourcing the same number of products from a different number of supplier countries.

[73] This statistic is computed as the share of aggregate spending in materials that is accounted by domestic inputs.

[74] Note that solving the model in principle entails finding a fixed point for firms' optimal sourcing strategies, as the general equilibrium variables depend on the domestic shares of all other firms which in turn depend on the general equilibrium variables. The structure of the economy, however, suggests a calibration approach which bypasses this computation and which we describe in detail in the Appendix. Intuitively: we can calibrate a normalized version of both fixed costs and innate productivity, where these are scaled by an appropriate transformation of the general equilibrium variables. Because the general equilibrium variables depend on firms' import behavior only via domestic

estimated.[75] As can be seen, the model can be calibrated to match the data accurately. Note in particular that the correlation between productivity and fixed costs is calibrated to be positive. This is necessary to match the far from perfect negative correlation between sales and domestic shares in the data. We show in the Appendix that the calibrated model is also able to match a number of non-targeted moments relatively well.[76]

Estimated				
Parameter	Value	Identified from		
σ	3.83	Revenue/Cost Data, Section 4.2		
ε	2.38	Prod. Function Estimation, Section 4.3		
γ	0.61	Prod. Function Estimation, Section 4.3		
η	0.38	Dom Share and Ext. Margin, Section 6.3		
Calibrated				
		Moment	Data	Model
μ_f	4.53	Aggregate Dom Share	0.71	0.71
f_I	0.19	Share of Importer	0.20	0.20
σ_φ	0.54	Dispersion in log Sales	1.62	1.62
σ_f	2.28	Dispersion in log Dom Shares	0.36	0.36
ρ	0.77	Correlation log Sales - log Dom Shares	-0.35	-0.35

Notes: This table contains all the estimates and calibrated structural parameters. While $(\sigma, \varepsilon, \gamma, \eta)$ are estimated directly from the micro-data, $\left(\mu_f, f_I, \sigma_\varphi, \sigma_f, \rho\right)$ are calibrated to match the 5 moments listed in column 3. All moments are calibrated simultaneously to match the set of moments.

Table 7: Structural Parameters

Welfare Gains from Input Trade. With the calibrated model at hand, we now compute the full welfare gains from trade. Table 8 contains the results. The first column shows that the model-predicted price-index gains are very close to those measured in the data.[77] This should not come as a surprise since such gains are a function of the joint distribution of sales and domestic shares which is a direct calibration target. Column two contains the main result of this section: the full welfare gains from input trade between the current trade equilibrium and autarky are predicted to be 18.41%. Thus, we see that only about half of the price-index gains translate into welfare gains once the resources spent in fixed costs are taken into account. The reason is that, as seen in column

shares, which are itself a calibration target, we can compute all prices after the normalized calibration and thus back out the underlying true fixed costs. This not only reduces the computational burden substantially (as we do not have to solve for a fixed point), but also implies that the parameter z is not required for the calibration.

[75]Note that these parameters satisfy the sufficient conditions of Proposition 7 as $\eta(\varepsilon - 1) = 0.463 < 1$ and $\eta\gamma(\sigma - 1) = 0.3188 < 1$.

[76]Table 11 reports the average domestic expenditure share, both for importers and the full population, as well as features of the joint distribution of sales and domestic shares for importers. Figure 5 reports the marginal distributions of domestic shares and log sales for importers both in the model and in the French data. Figure 6 reports average domestic shares by sales quintile for the sample of importers.

[77]The number reported in Table 8 for the French data comes from applying the one-sector price-index gains formula (22) to the Manufacturing sector pooling all sub-sectors together. We also leave the non-manufacturing sectors out of the analysis. For these reasons, the price-index gains reported in 8 do not coincide with those reported for the Manufacturing sector in Table 5 above.

three, a move to autarky would free up 15% of the labor force, which would counteract the increase in prices.

	Price-Index Gains (in %)	Welfare Gains (in %)	% of Labor in Fixed Cost Production
Model	39.86	18.41	15.34
French Data	41.53	-	-

Notes: This table contains both the price-index gains $\left(\frac{P^{AUT}}{P}\right)$ and full welfare gains $\left(\frac{W}{W^{AUT}}\right)$ within our calibrated model. See the main text for details.

Table 8: Welfare Gains from Input Trade

6.4 The Importance of Domestic Shares

In this subsection, we assess the value of the micro data on domestic shares for estimating the aggregate gains from trade. Note that, in absence of firm-level data on domestic shares, most of the analysis performed above is no longer applicable. This is because domestic share data is required for both the methodology to measure the price-index gains from trade developed in Section (3), as well as for the calibration strategy employed above to quantify the full welfare gains. Without such data, the study of the aggregate effects of input trade, be it on the price-index or on welfare, requires the calibration of a structural model.

We consider a simple exercise to assess the value of the micro data on domestic expenditure shares. We re-calibrate the model of the previous section dropping the two moments associated with domestic expenditure shares, namely the dispersion of domestic shares and their correlation with sales. That is, we perform an alternative calibration where we match: (i) the share of importers, (ii) the aggregate import share and (iii) the dispersion in the distribution of sales. Accordingly, we set the dispersion in fixed costs and their correlation with physical productivity both to zero, i.e. $\sigma_f = \rho = 0$.

We report the results in Table 12 in the Appendix, where the baseline calibration is also displayed for comparison. The calibrated parameters in the model without data on domestic shares - henceforth NSD - imply aggregate gains from trade that are biased upwards relative to those of the baseline. In particular, the NSD model over-predicts both the price-index and the welfare gains from trade relative to the baseline. This is intuitive. By relying on productivity as the single source of heterogeneity, the NSD model generates a perfectly negative correlation between physical productivity and the domestic share. Since such share captures the reduction in unit cost associated with trade, this means that firms with higher innate productivity experience larger reductions in their unit costs, a feature that tends to make input trade more attractive. Notice that this model generates a counterfactually strong negative correlation between sales and domestic shares. Finally, note that the biases in the estimates of the gains from trade can be quantitatively meaningful, of about 11% for the price-index gains and 22% for welfare.

7 Conclusion

Firms around the world routinely engage in input trade. By accessing cheaper, better or novel inputs from abroad, they reduce their costs of production. Quantifying the aggregate consequences of input trade, however, has been limited by an inherent difficulty. On the one hand, firms differ vastly in the intensity with which they participate in international markets. This makes aggregative trade models inapplicable to measure the gains from input trade. On the other hand, fully-specified firm-based models of import behavior are challenging, at least as long as they are sufficiently rich to match salient features of the micro-data. Not only is firms' extensive margin of input trade non-trivial to characterize, but - more importantly - such an approach requires researchers to fully specify (and estimate) the entire environment including production technologies, the heterogeneity across firms and sourcing countries and the structure of output markets.

In this paper, we developed a methodology that bypasses these concerns. In particular, we show how one can use readily available micro-data to easily quantify the gains from input trade at both the firm and aggregate level for a wide class of models, which are central for applied work. Our first main result showed that firms' domestic expenditure share in material spending is a sufficient statistic for their unit costs, as long as domestic and foreign inputs are combined with a constant elasticity of substitution. Importantly, this result is remarkably general in that one does not have to impose any restrictions on firms' import environment. One can allow for arbitrary distributions of qualities and prices across potential sourcing countries, one does not have to assume that firms share the same production function for imported inputs and crucially not have to take a stand on how firms end up with their set of trading partners. Hence, regardless of the microstructure of firms' trading environment, all models (within the CES class) will imply exactly the same unit-cost reductions as long as they are consistent with the micro-data and share the same estimate of the elasticity of substitution. Our second main result concerns the aggregate gains from input trade. Within the context of a multi-sector, general-equilibrium trade model, with a rich input-output structure, we show that the observable joint distribution of value added and domestic expenditure shares fully determines the aggregate gains from input trade. Hence, despite the fact that the economy is non-aggregative, the micro-data on value added and domestic shares contains sufficient information to perform the correct aggregation.

We apply our methodology to the French economy. Using micro-data on the population of French manufacturing firms we first propose a new method to estimate the elasticity of substitution between domestic and imported varieties. Using the co-movement between changes in domestic spending and changes in revenue productivity at the firm-level and exploiting changes in aggregate trade flows as an instrument, we estimate a "firm-level trade elasticity" of about two and hence substantially smaller than elasticities estimated from aggregate data. We then focus on the normative implications of input trade. We first show that there is substantial variation in the benefits from international sourcing. While some firms manage to reduce their unit costs by 50% holding aggregate prices fixed, the median importer were to see its costs only increase by 10% if the French economy were to put into input autarky. At the aggregate level, manufacturing prices would be 27% higher and the economy-wide consumer price level were to increase by 9% if input trade was abolished. We also demonstrate that an aggregative trade model had led the researcher to biased conclusions and that such bias can be substantial depending on the choice of the elasticity. By calibrating a model with fixed costs, we finally show that taking into account the resource loss of engaging in international reduces these gains by about 50%.

References

Amiti, M. and J. Konings: 2007, 'Trade Liberalization, Intermediate Inputs and Productivity'. *American Economic Review* **97**(5), 1611–1638.

Antràs, P., T. C. Fort, and F. Tintelnot: 2014, 'The Margins of Global Sourcing: Theory and Evidence from U.S. Firms'. Working Paper.

Arkolakis, C.: 2010, 'Market Penetration Costs and the New Consumer Margin in International Trade'. *Journal of Political Economy* **118**(6), 1151–1199.

Arkolakis, C., A. Costinot, and A. Rodriguez-Clare: 2012, 'New Trade Models, Same Old Gains?'. *American Economic Review* **102(1)**, 94–130.

Arkolakis, C. and M.-A. Muendler: 2011, 'The Extensive Margin of Exporting Products: A Firm-level Analysis'. Working Paper.

Bernard, A. B., J. B. Jensen, S. Redding, and P. Schott: 2007, 'Firms in International Trade'. *Journal of Economic Perspectives* **3**(21), 105–130.

Bernard, A. B., J. B. Jensen, S. J. Redding, and P. K. Schott: 2012, 'The Empirics of Firm Heterogeneity and International Trade'. *Annual Review of Economics* **4**, 283–313.

Bernard, A. B., S. J. Redding, and P. K. Schott: 2011, 'Multi-Product Firms and Trade Liberalization'. *Quarterly Journal of Economics* **126**(3), 1271–1318.

Blaum, J., C. Lelarge, and M. Peters: 2013, 'Non-Homothetic Import Demand: Firm Productivity and Quality Bias'. Working Paper.

Broda, C. and D. E. Weinstein: 2006, 'Globalization and the Gains from Varieties'. *Quarterly Journal of Economics* **121**(2), 541–585.

Caliendo, L. and F. Parro: forthcoming, 'Estimates of the Trade and Welfare Effects of NAFTA'. *The Review of Economic Studies*.

Chaney, T.: 2013, 'The Network Structure of International Trade'.

Costinot, A. and A. Rodriguez-Clare: 2014, 'Trade Theory with Numbers: Quantifying the Consequences of Globalization'. In: *Handbook of International Economics*.

De Loecker, J. and P. Goldberg: 2013, 'Firm Performance in a Global Market'.

De Loecker, J., P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik: 2012, 'Prices, Markups and Trade Reform'. Working Paper.

Dhingra, S. and J. Morrow: 2012, 'Monopolistic Competition and Optimum Product Diversity Under Firm Heterogeneity'. Working Paper.

Eaton, J. and S. Kortum: 2002, 'Technology, Geography, and Trade'. *Econometrica* **70**(5), 1741–1779.

Eaton, J., S. Kortum, and F. Kramarz: 2004, 'Dissecting Trade: Firms, Industries, and Export Destinations'. *The American Economic Review* **Vol. 94**(2), 150–154.

Eaton, J., S. Kortum, and F. Kramarz: 2011, 'An Anatomy of International Trade: Evidence from French Firms'. *Econometrica* **79**(5), 1453–1498.

Eslava, M., A. C. Fieler, and D. Y. Xu: 2014, 'Trade, Skills, and Quality Upgrading: A Theory with Evidence from Colombia'. Working Paper.

Fabinger, M. and G. Weyl: 2012, '"Pass-Through and Demand Forms"'. Working Paper.

Goldberg, P., A. Khandelwal, N. Pavcnik, and P. Topalova: 2010, 'Imported Intermediate Inputs and Domestic Product Growth: Evidence from India'. *Quarterly Journal of Economics*.

Gopinath, G. and B. Neiman: 2014, 'Trade Adjustment and Productivity in Large Crises'. *American Economic Review* **104**(3).

Halpern, L., M. Koren, and A. Szeidl: 2011, 'Imported Inputs and Productivity'. Working Paper.

Hummels, D., R. Jorgenson, J. R. Munch, and C. Xiang: 2011, 'The Wage Effect of Offshoring: Evidence from Danish Matched Worker-Firm Data'. Working Paper.

Jia, P.: 2008, 'What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry"'. *Econometrica* **76**(6), 1263–1316.

Jones, C.: 2011, 'Intermediate Goods and Weak Links in the Theory of Economic Development'. *American Economic Journal: Macroeconomics*.

Kasahara, H. and J. Rodrigue: 2008, 'Does the use of imported intermediates increase productivity? Plant-level evidence'. *Journal of Development Economics* **87**, 106–118.

Khandelwal, A. and P. Topalova: 2011, 'Trade Liberalization and Firm Productivity: The Case of India'. *The Review of Economics and Statistics* **93**(3), 995–1009.

Kugler, M. and E. Verhoogen: 2011, 'Prices, Plant Size, and Product Quality'. *Review of Economic Studies* (14418), 1–33.

Levinsohn, J. and A. Petrin: 2003, 'Estimating Production Functions Using Inputs to Control for Unobservables'. *Review of Economic Studies* **70**, 317–341.

Levinsohn, J. and A. Petrin: 2012, 'Measuring aggregate productivity growth using plant-level data'. *The RAND Journal of Economics* **43**(4), 705–725.

Mayer, T., M. Melitz, and G. Ottaviano: 2010, 'Market size, Competition, and the Product Mix of Exporters'. Working Paper.

Melitz, M. J. and S. J. Redding: 2012, 'Heterogeneous Firms and Trade'. NBER Working Paper 18652.

Oberfield, E.: 2013, 'Business Networks, Production Chains, and Productivity: A Theory of Input-Output Architecture'. Working Paper.

Oberfield, E. and D. Raval: 2014, 'Micro Data and Macro Technology'. Working Paper.

Olley, S. and A. Pakes: 1996, 'The Dynamics of Productivity in the Telecommunications Equipment Industry'. *Econometrica* **64**(6), 1263–97.

Pavcnik, N.: 2002, 'Trade Liberalization, Exit and Productivity Improvements: Evidence from Chilean Plants'. *Review of Economic Studies* **69**(1), 245–276.

Ramanarayanan, A.: 2014, 'Imported Inputs and the Gains from Trade'. Working Paper.

Simonovska, I. and M. E. Waugh: 2013, 'The Elasticity of Trade: Estimates and Evidence'. *Journal of International Economics* **92**(1), 34–50.

Simonovska, I. and M. E. Waugh: 2014, 'TradeModels, Trade Elasticities, and the Gains from Trade'.

Wooldridge, J.: 2009, 'On estimating firm-level production functions using proxy variables to control for unobservables'. *Ecomics Letters* **104**, 112–114.

8 Appendix

8.1 Proof of Proposition 2

Consider firm i with efficiency φ_i and an extensive margin Σ_i, facing prices $[p_c^i]_{c\in\Sigma}$. The profit-maximizing demand for imported varieties and domestic inputs also has to solve the dual problem, i.e. the cost-minimization problem (for simplicity we drop the subscript i from now on)

$$\Gamma(y,\varphi,[p_c],\Sigma) \equiv \min_z \left\{ p_D z_D + \sum_{c\in\Sigma} p_c z_c \text{ s.t. } x \geq \left(\frac{y}{\varphi l^\gamma}\right)^{\frac{1}{1-\gamma}} \right\},$$

where

$$x = \left(\beta_i (q_D z_D)^{\frac{\varepsilon-1}{\varepsilon}} + (1-\beta_i)\left(h_i\left([\eta(q_c,\varphi) z_c]_{c\in\Sigma}\right)\right)^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{\varepsilon}{\varepsilon-1}}.$$

We can first solve the import problem

$$\lambda(x_I) = \min_z \left\{ \sum_{c\in\Sigma} p_c z_c \text{ s.t. } h_i\left([\eta(q_c,\varphi) z_c]_{c\in\Sigma}\right) \geq x_I \right\}.$$

Because $h(.)$ has constant-returns-to scale, we get that

$$\lambda(x_I) = \min_{\tilde{z}} \left\{ x_I \sum_{c\in\Sigma} p_c \tilde{z}_c \text{ s.t. } h_i\left([\eta(q_c,\varphi)\tilde{z}_c]_{c\in\Sigma}\right) \geq 1 \right\} = x_I \times \lambda(1) = x_I \lambda. \tag{42}$$

Hence, $\lambda = \lambda(\Sigma,\varphi,[p_{ci}],G_q)$ is the price index for firm i, which depends on the underlying heterogeneity in quality G_q, firm productivity φ, the set of prices and the production functions. Crucially: from the point of view of the firm, it is constant for given Σ. The cost-minimization problem is hence given by

$$\min_z \left\{ p_D z_D + \lambda(\Sigma,\theta) x_I \text{ s.t. } x \geq \left(\frac{y}{\varphi l^\gamma}\right)^{\frac{1}{1-\gamma}} \right\}, \tag{43}$$

where

$$x = \left(\beta_i (q_D z_D)^{\frac{\varepsilon-1}{\varepsilon}} + (1-\beta_i) x_I^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{\varepsilon}{\varepsilon-1}}.$$

Hence

$$\begin{aligned} p_D - \mu\frac{\varepsilon}{\varepsilon-1} x^{1/\varepsilon} \frac{\varepsilon-1}{\varepsilon} \beta q_D^{\frac{\varepsilon-1}{\varepsilon}} z_D^{\frac{\varepsilon-1}{\varepsilon}} z_D^{-1} &= 0 \\ \lambda(\Sigma,\theta) - \mu\frac{\varepsilon}{\varepsilon-1} x^{1/\varepsilon} \frac{\varepsilon-1}{\varepsilon} (1-\beta) x_I^{\frac{\varepsilon-1}{\varepsilon}} x_I^{-1} &= 0 \end{aligned}$$

This implies

$$p_D z_D + \lambda(\Sigma,\theta) x_I = Q x^{1/\varepsilon} \left(\beta q_D^{\frac{\varepsilon-1}{\varepsilon}} z_D^{\frac{\varepsilon-1}{\varepsilon}} + (1-\beta) x_I^{\frac{\varepsilon-1}{\varepsilon}}\right) = Qx.$$

Hence, the correct price index is Q. As usual

$$\begin{aligned} z_D^{\frac{1}{\varepsilon}} &= Qx^{1/\varepsilon}\beta q_D^{\frac{\varepsilon-1}{\varepsilon}}\frac{1}{p_D} \\ (q_D z_D)^{\frac{1}{\varepsilon}} &= Qx^{1/\varepsilon}\beta\frac{q_D}{p_D} \\ \beta\left(q_D z_D\right)^{\frac{\varepsilon-1}{\varepsilon}} &= Q^{\varepsilon-1}x^{(\varepsilon-1)/\varepsilon}\beta^{\varepsilon}\left(\frac{q_D}{p_D}\right)^{\varepsilon-1} \end{aligned} \tag{44}$$

Similarly

$$(1-\beta)\left(x_I\right)^{\frac{\varepsilon-1}{\varepsilon}} = Q^{\varepsilon-1}x^{(\varepsilon-1)/\varepsilon}\left(1-\beta\right)^{\varepsilon}\left(\frac{1}{\lambda\left(\Sigma,\theta\right)}\right)^{\varepsilon-1}$$

so that

$$x^{\frac{\varepsilon-1}{\varepsilon}} = Q^{\varepsilon-1}x^{(\varepsilon-1)/\varepsilon}\left(\left(1-\beta\right)^{\varepsilon}\left(\frac{1}{\lambda\left(\Sigma,\theta\right)}\right)^{\varepsilon-1}+\beta^{\varepsilon}\left(\frac{q_D}{p_D}\right)^{\varepsilon-1}\right)$$

and hence

$$\begin{aligned} Q^{1-\varepsilon} &= \left(\left(1-\beta\right)^{\varepsilon}\left(\frac{1}{\lambda\left(\Sigma,\theta\right)}\right)^{\varepsilon-1}+\beta^{\varepsilon}\left(\frac{q_D}{p_D}\right)^{\varepsilon-1}\right) \\ Q &= \left(\left(1-\beta\right)^{\varepsilon}\left(\lambda\left(\Sigma,\theta\right)\right)^{1-\varepsilon}+\beta^{\varepsilon}\left(\frac{p_D}{q_D}\right)^{1-\varepsilon}\right)^{\frac{1}{1-\varepsilon}}. \end{aligned}$$

But then we get from using (44)

$$\begin{aligned} s_D &= \frac{p_D z_D}{Qx} = \frac{Qx^{1/\varepsilon}\beta q_D^{\frac{\varepsilon-1}{\varepsilon}}z_D^{\frac{\varepsilon-1}{\varepsilon}}}{Qx} = \frac{\beta\left(q_D z_D\right)^{\frac{\varepsilon-1}{\varepsilon}}}{x^{\frac{\varepsilon-1}{\varepsilon}}} \\ &= Q^{\varepsilon-1}\beta^{\varepsilon}\left(\frac{q_D}{p_D}\right)^{\varepsilon-1}, \end{aligned}$$

so that

$$Q = s_D^{\frac{1}{\varepsilon-1}}\left(\frac{1}{\beta}\right)^{\frac{\varepsilon}{\varepsilon-1}}\left(\frac{p_D}{q_D}\right). \tag{45}$$

The unit costs are then

$$\begin{aligned} UC &= \frac{1}{\varphi}w^{1-\gamma}Q^{\gamma} \\ &= \frac{1}{\varphi}w^{1-\gamma}\left(s_D^{\frac{1}{\varepsilon-1}}\left(\frac{1}{\beta}\right)^{\frac{\varepsilon}{\varepsilon-1}}\left(\frac{p_D}{q_D}\right)\right)^{\gamma} \\ &= \frac{1}{\varphi}\left(\frac{1}{\beta}\right)^{\frac{\varepsilon\gamma}{\varepsilon-1}}s_D^{\frac{\gamma}{\varepsilon-1}}w^{1-\gamma}\left(\frac{p_D}{q_D}\right)^{\gamma}, \end{aligned}$$

which is the required expression given in (12).

8.2 Proof of Proposition 3

Consider sector s. Let $[s_{D,i}]$ be the distribution of domestic shares. The unit costs of firm i in sector s are hence given by (see (12))

$$UC_{i,s} = \frac{1}{\tilde{\varphi}_i} \times (s_{D,i})^{\gamma_s/(\varepsilon_s-1)} \times \left(\frac{p_{D,s}}{q_{D,s}}\right)^{\gamma_s}. \tag{46}$$

Note that we normalized $w = 1$. Monopolistic competition implies that $p_{i,s} = \mu_s UC_{i,s}$, where $\mu_s = \frac{\sigma_s}{\sigma_s-1}$ is the equilibrium mark-up. Let $[p_{D,s}]$ be given. The ideal price index for sector $s$output, P_s, is given by

$$P_s = \mu_s \left(\int_{i=0}^{N_s} UC_i^{1-\sigma_s} di\right)^{\frac{1}{1-\sigma_s}} = \mu_s \left(\frac{p_{D,s}}{q_{D,s}}\right)^{\gamma_s} \left(\int_{i=0}^{N_s} \left(\frac{1}{\tilde{\varphi}_i} (s_{D,i})^{\gamma_s/(\varepsilon_s-1)}\right)^{1-\sigma_s} di\right)^{\frac{1}{1-\sigma_s}} \tag{47}$$

$$\equiv \mu_s \left(\frac{p_{D,s}}{q_{D,s}}\right)^{\gamma_s} \Upsilon_s, \tag{48}$$

where

$$\Upsilon_s = \left(\int_{i=0}^{N_j} \left(\frac{1}{\tilde{\varphi}_i} (s_{D,i})^{\gamma_j/(\varepsilon_j-1)}\right)^{1-\sigma_j} di\right)^{\frac{1}{1-\sigma_j}}. \tag{49}$$

The costs of domestic intermediaries for firms in sector s are therefore given by

$$p_{D,s} = \zeta_s^* \prod_{j=1}^{S} P_j^{\zeta_j^s}, \tag{50}$$

where $\zeta_s^* = \prod_{j=1}^{S} \left(\zeta_j^s\right)^{-\zeta_j^s}$. Using (47) we get

$$\begin{aligned} p_{D,s} &= \zeta_s^* \left(\prod_{j=1}^{S} \mu_j^{\zeta_j^s}\right) \left(\prod_{j=1}^{S} \left(\frac{1}{q_{D,j}}\right)^{\gamma_j \zeta_j^s}\right) \left(\prod_{j=1}^{S} p_{D,j}^{\gamma_j \zeta_j^s}\right) \left(\prod_{j=1}^{S} \Upsilon_j^{\zeta_j^s}\right) \\ &= \zeta_s^* \mu_s^* q_s^* \left(\prod_{j=1}^{S} p_{D,j}^{\gamma_j \zeta_j^s}\right) \left(\prod_{j=1}^{S} \Upsilon_j^{\zeta_j^s}\right) \\ &= \left[\zeta_s^* \mu_s^* q_s^* \left(\prod_{j \neq s} p_{D,j}^{\gamma_j \zeta_j^s}\right) \left(\prod_{j=1}^{S} \Upsilon_j^{\zeta_j^s}\right)\right]^{1/(1-\gamma_s \zeta_s^s)}, \end{aligned} \tag{51}$$

where $\mu_s^* = \prod_{j=1}^{S} \mu_j^{\zeta_j^s}$ and $\frac{1}{q_s^*} = \left(\prod_{j=1}^{S} (1/q_{D,j})^{\gamma_j \zeta_j^s}\right)$. Given the technological constants $\zeta_s^* \mu_s^* q_s^*$ and the endogenous trade-induces productivity terms $\{\Upsilon_s\}_s$, (51) is a system of S equation in the S unknowns $[p_{D,s}]_{s=1}^{S}$, which can be easily solved.

Given $[p_{D,s}]_{s=1}^{S}$, the consumer price index is then given by

$$\begin{aligned} P &= \prod_{s=1}^{S}\left(\frac{P_s}{\alpha_s}\right)^{\alpha_s} = \prod_{s=1}^{S}\left(\frac{1}{\alpha_s}\right)^{\alpha_s}\prod_{s=1}^{S}P_s^{\alpha_s} \\ &= \alpha^*\mu^*\frac{1}{q^*}\prod_{s=1}^{S}p_{D,s}^{\gamma_s\alpha_s}\prod_{s=1}^{S}\Upsilon_s^{\alpha_s}, \end{aligned}$$

where P_s follows from (47) and $\alpha^* = \prod_{s=1}^{S}\alpha_s^{-\alpha_s}$, $\mu^* = \prod_{s=1}^{S}\mu_s^{\alpha_s}$, $\frac{1}{q^*} = \prod_{s=1}^{S}q_{D,s}^{-\gamma_s\alpha_s}$.

The Gains from Trade are then defined by

$$G = ln\left(\frac{P^{Aut}}{P}\right),$$

where P is recursively defined by

$$\begin{aligned} P &= \alpha^*\mu^*\frac{1}{q^*}\prod_{s=1}^{S}p_{D,s}^{\gamma_s\alpha_s}\prod_{s=1}^{S}\Upsilon_s^{\alpha_s} \\ p_{D,s} &= \left[\zeta_s^*\mu_s^*q_s^*\left(\prod_{j\neq s}p_{D,j}^{\gamma_j\zeta_j^s}\right)\left(\prod_{j=1}^{S}\Upsilon_j^{\zeta_j^s}\right)\right]^{1/(1-\gamma_s\zeta_s^s)} \end{aligned}$$

and P^{Aut} by

$$\begin{aligned} P^{Aut} &= \alpha^*\mu^*\frac{1}{q^*}\prod_{s=1}^{S}\left(p_{D,s}^{Aut}\right)^{\gamma_s\alpha_s}\prod_{s=1}^{S}\Psi_s^{\alpha_s} \\ p_{D,s}^{Aut} &= \left[\zeta_s^*\mu_s^*q_s^*\left(\prod_{j\neq s}\left(p_{D,j}^{Aut}\right)^{\gamma_j\zeta_j^s}\right)\left(\prod_{j=1}^{S}\Psi_j^{\zeta_j^s}\right)\right]^{1/(1-\gamma_s\zeta_s^s)}, \end{aligned}$$

where Ψ is defined akin to (49) with $s_{D,i} = 1$.

Hence

$$exp\left(G\right) = \frac{\alpha^*\mu^*\frac{1}{q^*}\prod_{s=1}^{S}\left(p_{D,s}^{Aut}\right)^{\gamma_s\alpha_s}\prod_{s=1}^{S}\Psi_s^{\alpha_s}}{\alpha^*\mu^*\frac{1}{q^*}\prod_{s=1}^{S}p_{D,s}^{\gamma_s\alpha_s}\prod_{s=1}^{S}\Upsilon_s^{\alpha_s}} = \prod_{s=1}^{S}\left(\frac{p_{D,s}^{Aut}}{p_{D,s}}\right)^{\gamma_s\alpha_s}\prod_{s=1}^{S}\left(\frac{\Psi_s}{\Upsilon_s}\right)^{\alpha_s}, \tag{52}$$

where

$$\frac{p_{D,s}^{Aut}}{p_{D,s}} = \left(\prod_{j\neq s}\left(\frac{p_{D,j}^{Aut}}{p_{D,j}}\right)^{\gamma_j\zeta_j^s}\prod_{j=1}^{S}\left(\frac{\Psi_j}{\Upsilon_j}\right)^{\zeta_j^s}\right)^{1/(1-\gamma_s\zeta_s^s)}. \tag{53}$$

Now note that

$$\frac{\Psi_s}{\Upsilon_s} = \left(\frac{\int_{i=0}^{N_s}\left(1/\tilde{\varphi}_i\right)^{1-\sigma_s}di}{\int_{i=0}^{N_s}\left(\frac{1}{\tilde{\varphi}_i}\left(s_{D,i}\right)^{\gamma_s/(\varepsilon_s-1)}\right)^{1-\sigma_s}di}\right)^{\frac{1}{1-\sigma_s}}.$$

In the trade equilibrium, firm i's value added is given by

$$va_i = \kappa \times \left(\tilde{\varphi}_i \left(s_{D,i}\right)^{\gamma_s/(1-\varepsilon_s)}\right)^{\sigma-1}.$$

Hence

$$\begin{aligned}
\frac{\Psi_s}{\Upsilon_s} &= \left(\frac{\int_{i=0}^{N_s} \left(\tilde{\varphi}_i\right)^{\sigma_s-1} di}{\int_{i=0}^{N_s} \left(\tilde{\varphi}_i \left(s_{D,i}\right)^{\gamma_s/(1-\varepsilon_s)}\right)^{\sigma_s-1} di}\right)^{\frac{1}{1=\sigma_s}} \\
&= \left(\frac{\int_{i=0}^{N_s} \left(\tilde{\varphi}_i \left(s_{D,i}\right)^{\gamma_s/(1-\varepsilon_s)}\right)^{\sigma_s-1} \left(s_{D,i}\right)^{\gamma_s(\sigma_s-1)/(\varepsilon_s-1)} di}{\int_{i=0}^{N_s} \left(va_i \times \kappa^{-1}\right) di}\right)^{\frac{1}{1=\sigma_s}} \\
&= \left(\frac{\int_{i=0}^{N_s} va_i \left(s_{D,i}\right)^{\gamma_s(\sigma_s-1)/(\varepsilon_s-1)} di}{\int_{i=0}^{N_s} va_i di}\right)^{\frac{1}{1=\sigma_s}} \\
&= \left(\int_{i=0}^{N_s} \frac{va_i}{\int_{i=0}^{N_s} va_i di} \left(s_{D,i}\right)^{\gamma_s(\sigma_s-1)/(\varepsilon_s-1)} di\right)^{\frac{1}{1=\sigma_s}} \\
&\equiv \Lambda_s.
\end{aligned}$$

where $\pi_s \equiv ln\left(\frac{p_{D,s}^{Aut}}{p_{D,s}}\right)$. (53) implies that

$$\left(\frac{p_{D,s}^{Aut}}{p_{D,s}}\right)^{1-\gamma_s\zeta_s^s} = \left(\prod_{j=1}^{S} \Lambda_j^{\zeta_j^s}\right) \prod_{j\neq s} \left(\frac{p_{D,j}^{Aut}}{p_{D,j}}\right)^{\gamma_j\zeta_j^s}. \tag{54}$$

$$\frac{p_{D,s}^{Aut}}{p_{D,s}} = \left(\prod_{j=1}^{S} \Lambda_j^{\zeta_j^s}\right) \prod_{j} \left(\frac{p_{D,j}^{Aut}}{p_{D,j}}\right)^{\gamma_j\zeta_j^s}. \tag{55}$$

Hence, (52) and (54) yield

$$G = \sum_{s=1}^{S} \gamma_s\alpha_s\pi_s + \sum_{s=1}^{S} \alpha_s\Lambda_s, \tag{56}$$

$$\pi_s = \sum_{j=1}^{S} \zeta_j^s\Lambda_j + \sum_{j=1}^{S} \gamma_j\zeta_j^s\pi_j. \tag{57}$$

As Λ_s is directly observable from (19), G can easily be evaluated from (56) and (57). This is a simple system of linear equations. In particular, we can use (57) to get $[\pi_s]_{s=1}^S$ given $[\Lambda_s]_{s=1}^S$ and then use (56) to derive G.

Note also that (47) implies that

$$\frac{P_s^{AUT}}{P_s} = \left(\frac{p_{D,s}^{Aut}}{p_{D,s}}\right)^{\gamma_s} \frac{\Psi_s}{\Upsilon_s},$$

so that

$$ln\left(\frac{P_s^{AUT}}{P_s}\right) = \gamma_s \pi_s + \Lambda_s. \tag{58}$$

To express (56) and (57) in matrix notation, note that (57) can be written as

$$\begin{bmatrix} \pi_1 \\ \pi_2 \\ ... \\ \pi_S \end{bmatrix} = \begin{bmatrix} \zeta_1^1 & \zeta_2^1 & .. & \zeta_S^1 \\ \zeta_1^2 & \zeta_2^2 & ... & ... \\ ... & & & \\ \zeta_1^S & \zeta_2^S & & \zeta_S^S \end{bmatrix} \begin{bmatrix} \Lambda_1 \\ \Lambda_2 \\ ... \\ \Lambda_S \end{bmatrix} + \begin{bmatrix} \zeta_1^1\gamma_1 & \zeta_2^1\gamma_2 & .. & \zeta_S^1\gamma_S \\ \zeta_1^2\gamma_1 & \zeta_2^2\gamma_2 & ... & ... \\ ... & & & \\ \zeta_1^S\gamma_1 & \zeta_2^S\gamma_2 & & \zeta_S^S\gamma_S \end{bmatrix} \begin{bmatrix} \pi_1 \\ \pi_2 \\ ... \\ \pi_S \end{bmatrix}$$

or in matrix form

$$\pi = \Xi \times \Lambda + \Xi \times \Gamma \times \pi,$$

where

$$\Gamma = \begin{bmatrix} \gamma_1 & 0 & & 0 \\ 0 & \gamma_2 & & \\ & & ... & \\ 0 & & & \gamma_S \end{bmatrix}.$$

Hence,

$$\left(\mathcal{I} - \Xi \times \Gamma\right) \pi = \Xi \times \Lambda,$$

so that

$$\pi = \left(\mathcal{I} - \Xi \times \Gamma\right)^{-1} \Xi \times \Lambda. \tag{59}$$

Given (59), we can then solve for the aggregate gains from (56) as

$$ln\,(G) \;\; = \;\; \alpha'\Gamma\pi + \alpha'\Lambda = \alpha'\Gamma\left(\mathcal{I} - \Xi \times \Gamma\right)^{-1} \Xi\Lambda + \alpha'\Lambda.$$

8.3 Data Description

Our main data set stems from the information system of the French custom administration (DGDDI) and contains the universe of import and export flows by French manufacturing firms.[78] The data is collected at the 8-digit (NC8) level and a firm located within the French metropolitan territory must report this detailed information as long as the following criteria are met. For imports from outside the EU, reporting is required from each firm and flow if the imported value to exceeds 1,000 Euros. For within EU imports, import flows had to be reported (between 2001 and 2006)[79] as long as the firm's annual trade value exceeds 100,000 Euros. However, some firms (ca. 10,000 firm year observations out of ca. 130,000) also report while they are below the threshold. Clearly, the existence of this administrative threshold induces a censoring for small EU importers. While

[78]For imports from outside the EU, all shipments must be reported to the custom administration. The conditions are more stringent for exports since all shipments (even within EU) must be reported to the custom administration.

[79]Between 1993 and 2001, the threshold was ca. 40,000 euros. After 2006, it was raised to 150,000 euros and to 460,000 euros after 2011.

it does not affect aggregate, value weighted statistics, it will generate some attenuation bias in our econometric analyses, which would bias our results against finding any result.

In spite of this limitation, the attractive feature of the French data is the presence of unique firm identifiers (the SIREN code), which is available in all French administrative files. Hence, various other datasets can be matched to the trade data at the firm level. To learn about the characteristics of the firms in our sample we employ fiscal files.[80] Sales are deflated using price indices of value added at the 3 digit level obtained from the French national accounts. To measure the expenditure on domestic inputs, we subtract the total import value from the total expenditure on wares and inputs reported in the fiscal files. Capital, used for the TFP estimation, is measured at book value (historical cost).

Finally we incorporate information on the ownership structure from the LIFI/DIANE (BvDEP) files. These files are constructed at INSEE using a yearly survey (LIFI) describing the structure of ownership of all of the French firms in the private sector whose financial investments in other firms (participation) are higher than 1.2 million Euros or having sales above 60 million Euros or more than 500 employees. This survey is complemented with the information about ownership structure available in the DIANE (BvDEP) files, which are constructed using the annual mandatory reports to commercial courts, and with the register of firms that are controlled by the State.

Using these datasets, we construct a non-balanced panel dataset spanning the period from 2001 to 2006. Some basic characteristics of importing and non-importing firms are contained in Table 9. For comparison, we also report the results for exporting firms. Expectedly, importers outperform domestic firm in essentially all dimensions we look at (see also Bernard et al. (2012)). Furthermore, import and export status are highly correlated.

8.4 Estimates of the Elasticity of Substitution ε

Table 10 below contains various robustness checks for our estimates for the elasticity of substitution using both factor shares and proxy methods to arrive at an estimate for γ.

8.5 Non-Targeted Moments and Model Comparison

Table 11 and Figures 5 and 6 report the fit between the model and additional non-targeted moments. Table 11 compares the model with the data at various dimensions. The model performs relatively well. That fact that we under predict the dispersion of sales might be due to the fact that the log-normal distribution of productivity has too thin tails. This might also be the reason why we under predict the share of sales by importers - there are simply too few very large firms in our model.

[80]The firm level accounting information is retrieved from two different files: the BRN ("Bénéfices Réels Normaux") and the RSI ("Régime Simplifié d'Imposition"). The BRN contains the balance sheet of all firms in the traded sectors with sales above 730,000 Euros. The RSI is the counterpart of the BRN for firms with sales below 730,000 Euros. Although the details of the reporting differs, for our purpose these two data sets contain essentially the same information. Their union covers nearly the entire universe of all French firms.

	Full sample	**Importers**	**Non importers**	**Exporters**	**Non exporters**
Employment	25	92	8	81	9
Sales	5,455	21,752	1,379	19,171	1,468
Sales per worker	126	208	105	196	105
Value added	1,515	5,972	400	5,294	416
Value added per worker	45	55	43	55	43
Capital	2,217	8,728	588	7,661	634
Capital per worker	44	64	40	61	40
Inputs	2,600	10,225	693	8,943	756
Domestic share	0.943	0.698	1	0.790	0.986
Share of importers	0.200	1	0	0.677	0.061
Share of exporters	0.225	0.762	0.091	1	0
Share of firms that are part of an international group	0.029	0.131	0.004	0.113	0.005
Productivity (factor shares)	39.173	65.450	32.989	63.858	32.359
Number of observations (firm * year)	650,401	130,135	520,266	146,496	503,905
Number of firms	172,244	38,240	148,619	44,648	146,423

Notes: Sales, wages, expenditures on imports or exports are all expressed in 2005 prices using a 3-digit industry level price deflator. Our capital measure is the book value reported in firms' balance sheets ("historical cost"). We measure employees by occupation. Skilled workers are engineers, technicians and managers, workers of intermediate skills are skilled blue and white collars and low skilled workers are members of unskilled occupations. A firm is member of an international group if at least one affiliate or the headquarter is located outside of France.

Table 9: Characteristics of importers, exporters and domestic firms

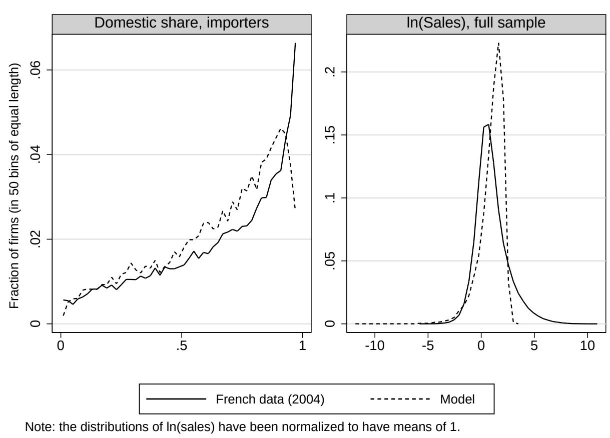

Notes: The left panel (right panel) shows the distribution of domestic shares (log sales) in the data (solid line) and in the model (dashed line).

Figure 5: Marginal Distributions: Model and Data

Figure 5 reports the marginal distributions of domestic shares and log sales for importers both in the model and in the French data. The model captures the marginal distribution of domestic shares quite well. It under predicts the density for very small importers, which is natural in a model of fixed costs - it not worth it pay the fixed costs to then import tiny amounts. When we compare the

			$\Delta \ln(WES)$	$\tilde{\gamma}_s \times \Delta \ln(s_D)$	ε	N
Factor shares	Full sample	First stage	-0.019***	-	-	526,687
Sample: 2002-2006			(0.003)			
	Importers only	First stage	-0.010***	-	-	65,799
			(0.004)			
Factor shares	Full sample	All weights	-	-0.726***	2.378***	526,687
(bootstrapped SE)				(0.197)	(0.523)	
Sample: 2002-2006		Pre-sample (2001) weights	-	-1.407***	1.711***	443,954
				(0.356)	(0.166)	
	Importers only	All weights	-	-0.756	2.322**	65,799
	(2 digit dummies)			(0.537)	(1.014)	
		Pre-sample (2001) weights	-	-1.121*	1.892***	54,604
				(0.632)	(0.541)	
2-step GMM	Full sample	First stage	-0.017***	-	-	331,421
Sample: 2004-2006			(0.003)			
	Importers only	First stage	-0.008**	-	-	53,349
			(0.003)			
2-step GMM	Full sample	All weights	-	-1.288***	1.776***	331,421
(bootstrapped SE)				(0.395)	(0.288)	
Sample: 2004-2006		Pre-sample (2001) weights	-	-2.152***	1.465***	258,957
				(0.652)	(0.148)	
	Importers only	All weights	-	-1.116	1.896**	53,349
	(2 digit dummies)			(1.203)	(0.850)	
		Pre-sample (2001) weights	-	-1.968	1.508***	43,393
				(1.910)	(0.402)	
2-step GMM, translog	Full sample	First stage	-0.016***	-	-	331,421
Sample: 2004-2006			(0.003)			
	Importers only	First stage	-0.008**	-	-	53,349
			(0.003)			
2-step GMM, translog	Full sample	All weights	-	-1.376***	1.727***	331,421
(bootstrapped SE)				(0.402)	(0.235)	
Sample: 2004-2006		Pre-sample (2001) weights	-	-2.371***	1.422***	258,957
				(0.693)	(0.119)	
	Importers only	All weights	-	-1.246	1.802**	53,349
	(2 digit dummies)			(1.071)	(0.735)	
		Pre-sample (2001) weights	-	-2.171	1.461***	43,393
				(1.835)	(0.332)	

Table 10: Estimating the Elasticity of Substitution ε

Non-Targeted Moments	French Data	Baseline
Avg Domestic Share (Importers)	0.70	0.65
Avg Domestic Share (Population)	0.94	0.93
Agg Domestic Share (Importers)	0.63	0.55
Dispersion log Sales (Importers)	1.66	0.99
Dispersion log Dom Shares (Importers)	0.69	0.62
Correlation log Sales - log Dom Shares (Importers)	-0.03	0.00
Share of Sales by Importers	0.80	0.63

Notes: This table reports some non-targeted moments in the micro-data (column 2) and in our baseline calibration (column 3). The calibrated parameters of the benchmark model are contained in Table 7.

Table 11: Non-Targeted Moments

distribution of sales between model and the data, we again see that the model generates too little dispersion.

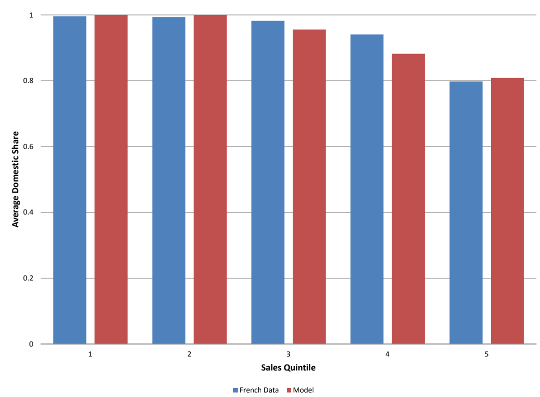

Notes: The graph depicts the average domestic expenditure share for different size groups in the French economy. We depict both the micro-data and the data from the calibrated model.

Figure 6: Correlation Structure: Model and Data

Figure 6 reports the average domestic share by sales quintile. The model fits that moment relatively well. In particular, the model captures that sales and import behavior are not perfectly aligned.

Table 12 finally reports the calibrated parameters and model fit of our alternative calibration strategy if no information domestic share was available.

	Baseline		No s_D Data	
	Model	Parameter	Model	Parameter
Aggregate Domestic Share	0.71	$\mu_f = 4.53$	0.71	$\mu_f = 3.08$
Share of Importers	0.20	$f_I = 0.19$	0.20	$f_I = 0.14$
Dispersion in log Sales	1.62	$\sigma_\varphi = 0.54$	1.62	$\sigma_\varphi = 0.55$
Dispersion in Domestic Shares	0.36	$\sigma_f = 2.28$	0.13	$\sigma_f = 0$
Correlation log Sales - Dom Shares	-0.35	$\rho = 0.77$	-0.70	$\rho = 0$
Price-Index Gains	39.86 %		44.26%	Bias:11.04 %
Welfare Gains	18.41%		22.50%	Bias:22.22 %

Notes: The table contains the calibrated parameters and moments of an alternative model, which does not use the information on domestic expenditure shares. This model is only calibrated to the first three moments. For completeness we also report our baseline calibration.

Table 12: Calibration Without Domestic Shares

While the model is calibrated to match the first three moments, the results in Table 12 show that it is unsuccessful to predict a distribution of import shares, which are consistent with the data. Not only is there too strong a negative correlation with firm sales, the model also under predicts the cross-sectional dispersion in domestic shares.