

NBER WORKING PAPER SERIES

THE GAINS FROM INPUT TRADE IN FIRM-BASED MODELS OF IMPORTING

Joaquin Blaum
Claire LeLarge
Michael Peters

Working Paper 21504
<http://www.nber.org/papers/w21504>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2015

We especially thank Arnaud Costinot for his feedback and discussion at the Cowles conference and Jonathan Eaton and Pablo Fajgelbaum for very helpful comments. We also thank Pol Antràs, Costas Arkolakis, Penny Goldberg, Sam Kortum, Jonathan Vogel and Daniel Xu. We are grateful to seminar participants at Brown, Columbia, Dartmouth, LSE, Penn State, Princeton, Stanford, UCLA and Yale. A previous version of this paper circulated under the title “Estimating the Productivity Gains from Importing”. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by Joaquin Blaum, Claire LeLarge, and Michael Peters. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Gains from Input Trade in Firm-Based Models of Importing
Joaquin Blaum, Claire LeLarge, and Michael Peters
NBER Working Paper No. 21504
August 2015
JEL No. D21,D22,F11,F12,F14,F62

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 how much they buy in foreign markets, results based on aggregate models do not apply. We develop a methodology to quantify the gains from input trade for a class of firm-based models of importing. We derive a sufficiency result: the change in consumer prices induced by input trade is fully determined from the joint distribution of value added and domestic expenditure shares in material spending across firms. We provide a simple formula that can be readily evaluated given the micro-data. In an application to French data, we find that consumer prices of manufacturing products would be 27% higher in the absence of input trade.

Joaquin Blaum
64 Waterman Street
Box B, Brown University
Providence, RI 02912
joaquin_blaum@brown.edu

Claire LeLarge
National Institute of Statistics
and Economic Studies
Direction des Études et
Synthèses Économiques
15 Boulevard Gabriel Peri
92244 Malakoff Cedex, France
and CREST
claire.lelarge@gmail.com

Michael Peters
Department of Economics
Yale University
28 Hillhouse Avenue
New Haven, CT 06511
and NBER
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 of how domestic consumers benefit from producers engaging in international sourcing. By providing access to novel, cheaper or higher quality inputs, input trade reduces firms' unit costs and lowers domestic prices, therefore increasing consumers' purchasing power. In this paper, we develop a methodology to quantify this channel and provide an application to France.

Quantifying the welfare consequences of input trade is not straightforward. Recent quantitative trade models that allow for trade in inputs feature the convenient property that welfare can be measured with aggregate data only - e.g. Eaton et al. (2011), Caliendo and Parro (2015) and Costinot and Rodríguez-Clare (2014). This property, however, relies on the assumption that firms' import intensities are equalized - a feature that is at odds with the data. In particular, importing firms differ substantially in the share of material spending they allocate to foreign inputs. In this paper, we show that accounting for this heterogeneity in import exposure, which requires resorting to firm-based models of importing, significantly affects the measurement of the gains from input trade.

We provide a sufficiency result that applies to a class of firm-based models of importing where firms' demand system between domestic and foreign inputs is CES.¹ In particular, we show that firm-level data on domestic shares of intermediate spending and value added is sufficient to compute the *consumer price gains from input trade*, i.e. the change in consumer prices relative to a situation of "input autarky" where firms can use only domestic inputs. Because this result does not rely on specific assumptions on firms' import environment nor on how firms' determine their trading partners, a variety of models implies the exact same consumer price gains given the micro data. We provide a closed-form expression that makes calculating the consumer price gains straightforward.

Our result builds on a simple insight. By inverting the demand system for intermediates, we can link the firm's unit cost to its spending pattern on domestic inputs. When such a demand system is CES, the unit cost reduction from importing, which we refer to as the *producer gains*, is fully determined by the domestic expenditure share and two structural parameters.² In particular, the producer gains are high when the domestic share is low. In a second step, we then show how these producer gains can be aggregated to compute the consumer price gains of input trade taking general equilibrium effects into account. In a multi-sector trade model with intersectoral linkages and monopolistic competition, such consumer price gains are akin to a value-added weighted average of the producer gains. In this way, the joint distribution of domestic shares and value added is sufficient to characterize the effect of input trade on consumer prices. Importantly, a key aspect of the data is how firm size and domestic shares correlate; if bigger firms feature higher trade shares, then the consumer gains will turn out to be large.

Our procedure places no restrictions on several components of the theory related to firms' import

¹This class nests several frameworks used in the literature, e.g. Halpern et al. (2011), Gopinath and Neiman (2014), Antràs et al. (2014) and Goldberg et al. (2010).

²These are the elasticity of firm output to intermediate inputs and the elasticity of substitution between domestic and international varieties.

environment. First, we do not require information on the prices and qualities of the foreign inputs, nor on how these are combined for production.³ While these elements are in principle required to compute the firm’s unit cost, they are fully summarized by the domestic expenditure share. Consider next the extensive margin of trade. Because our sufficiency result is derived purely from the cost minimization problem taking the set of trading partners as given, it holds regardless of how the firm finds its suppliers, e.g. whether importing is limited by the presence of fixed costs or a process of network formation. In this way, our approach bypasses data requirements as well as functional form and behavioral assumptions and therefore holds in a variety of settings.

An important parameter in our methodology is the elasticity of substitution between domestically sourced and imported inputs. Because firm-based models of importing do not generate a standard gravity equation, this parameter is not identified from aggregate data. We therefore devise a strategy to identify it from firm-level variation. By expressing firms’ output in terms of material spending, the domestic share appears as an additional input in the production function. Because the sensitivity of firm revenue to domestic spending depends on the elasticity of substitution, we can estimate this parameter with methods akin to production function estimation. To address the endogeneity concern that unobserved productivity shocks might lead to both lower domestic spending and higher revenue, we use changes in the world supply of particular varieties as an instrument for firms’ domestic spending.

We apply our methodology to the population of manufacturing firms in France. We estimate the distribution of trade-induced changes in unit costs across firms. We find substantial cross-sectional dispersion in these producer gains, which is induced by the observed variation in domestic expenditure shares. While the median unit cost reduction is 11%, it exceeds 80% for 10% of the firms. Moreover, bigger firms benefit more from input trade. We then aggregate the producer gains to compute the consumer price gains by relying on the joint distribution of domestic shares and value added. We find that input trade reduces consumer prices of manufacturing products by 27%.⁴ There are three reasons why the consumer gains exceed the median producer gains, which go back to the above-mentioned patterns. First, the dispersion in producer gains is valued by consumers given their elastic demand. Second, the positive relation between the producer gains and firm size is beneficial because the endogenous productivity gains from importing and firm efficiency are complements. Finally, there are important linkages between firms whereby non-importers buy intermediates from importing firms. This structure of round-about production amplifies the gains from input trade in general equilibrium.

We then consider the effect of input trade on a broader notion of welfare. While the consumer price gains are an important component of the welfare gains from input trade, they do not take into account any resources spent by firms to attain their equilibrium sourcing strategies. Because such

³We consider a production structure where foreign inputs are aggregated into an import bundle. We require that such an import bundle is combined with a bundle of domestic inputs in a CES fashion, but place no restrictions on the foreign input aggregator.

⁴When we include the non-manufacturing sector, the consumer price gains amount to 9%. Note that manufacturing accounts for 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.

resource loss cannot be read off the data, we need to commit to a particular model of the extensive margin of trade and fully calibrate it. We consider a model where participation in international markets is limited by fixed costs. We parametrize the distributions of qualities, prices and fixed costs, and discipline the model with moments of the French data. We target the joint distribution of domestic expenditure shares and value added, which as argued above contains important information about the gains from input trade. The main result of this exercise is that the full welfare gains are about half as large as the consumer price gains.

Because our methodology stresses the importance of micro-data, a natural question is: how do our estimates change when only aggregate data is used? Relying on aggregate data affects the estimates of the gains from input trade in two distinct ways. First, there is a bias that arises from ignoring the heterogeneity in firms' import shares for given parameters. While this bias can be positive or negative, we show that the sign depends only on parameters and not on the micro-data. A second type of bias is related to the estimation of the elasticity of substitution. Approaches that rely on a standard gravity equation to estimate this parameter may lead to different results than an analysis based on micro-data. In our application to the French data, the first bias leads to overestimating the consumer gains by about 10%, while the second one leads to underestimating them by 50%.⁵ Thus, the magnitude of the different errors from using aggregate data can be substantial.

Our paper contributes to a recent literature on quantitative models of input trade. On the one hand, there are aggregate trade models as Eaton et al. (2011), Caliendo and Parro (2015) and Costinot and Rodríguez-Clare (2014). These models have the convenient implication that the welfare consequences of input trade are fully determined from readily available aggregate data.⁶ This property, however, crucially relies on a theoretical structure where firms' import shares are equalized - an implication which is strongly at odds with the data. On the other hand, there is a literature on firm-based models of importing - see Halpern et al. (2011), Gopinath and Neiman (2014) or Ramanarayanan (2014). Our approach is different in two aspects. First, the existing contributions do not rely on firms' domestic expenditure shares to directly measure the unit cost reductions from importing at the firm level. Instead, they measure these producer gains indirectly by estimating or calibrating the entire structural model. Because the firm's extensive margin problem is tractable only under particular assumptions and the structure of output markets needs to be fully specified, their results rely on these restrictions. Secondly, the existing papers do not target the joint distribution of value added and domestic shares in their estimations, nor exploit the fact that such data is sufficient to characterize the effect of input trade on consumer prices. Ramanarayanan (2014) and Gopinath and Neiman (2014) for example consider a model that generates a perfect, and hence counterfactual, correlation between firm-size and domestic shares. We explicitly show that the consumer price gains in such type of model are too high.⁷

⁵Using the micro-data to estimate the elasticity of substitution turns out to be important as we obtain a value close to two. Estimation approaches that rely on aggregate data typically find values closer to four.

⁶Specifically, the welfare gains are summarized by the change in the aggregate domestic expenditure share and a trade elasticity, which can be estimated from aggregate trade flows.

⁷Specifically, models where physical efficiency is the single source of firm heterogeneity generate a perfect assignment between efficiency and the domestic share. As more efficient firms experience larger reductions in their unit cost, the

Our paper also builds on a recent literature which stresses that complementarities across inputs of production make the import problem different from the better known export problem. In particular, firms' extensive margin of trade is in general harder to characterize - see Blaum et al. (2013) and Antràs et al. (2014). On the export side, recent work has been able to quantitatively account for firms' entry behavior into different markets.⁸ In contrast, theories that can account for the pattern of entry into import markets are less developed. A notable exception is the recent contribution by Antràs et al. (2014), who study a firm-based model of importing and adapt the estimation procedure by Jia (2008) to match positive aspects of import behavior. In contrast, our paper focuses on normative aspects of input trade. Our main result stresses that, conditional on the micro-data, the effect of input trade on consumer prices does not depend on the mechanics of the extensive margin or other aspects of the import environment.

At a conceptual level, our paper is related to Feenstra (1994) and Arkolakis et al. (2012). As in Feenstra (1994), we express changes in unobserved unit costs in terms of observed expenditure shares. Relative to Arkolakis et al. (2012), our sufficient statistic for the firm's unit cost is related to their sufficient statistic for aggregate welfare. In particular, we show that, conditional on the micro-data on firms' domestic shares and a "trade elasticity", which in our setup corresponds to the elasticity of substitution of the firm's import demand system, a wide class of models will imply the exact same distribution of producer gains across firms.

Finally, a number of empirically oriented papers study trade liberalization episodes to provide evidence on the link between imported inputs and firm productivity - see e.g. Amiti and Konings (2007), Goldberg et al. (2010) or Khandelwal and Topalova (2011).⁹ Our results are complementary to this literature as we provide a structural interpretation of this empirical evidence. In particular, from the point of view of applied researchers, our sufficiency result provides a way to analyze episodes of trade liberalization, or other changes in firms' import environment, without having to fully specify and solve a structural model of importing. The observable change in the domestic expenditure shares correctly measures the effect of the policy on firms' unit costs, taking all adjustments into account. If micro-data on value added is also available, our formula for the consumer gains can be used to gauge the full effect of the policy on consumer prices in general equilibrium.

The remainder of the paper is structured as follows. In Section 2, we present direct evidence from the population of French firms for why firm-based models of importing are necessary to study the normative consequences of input trade. Section 3 lays out the class of models we consider and derives our sufficiency results for the producer and consumer gains from input trade. The empirical application to France is contained in Section 4. In Section 5, we calibrate a version of our model with a fully-specified extensive margin of importing to provide a full measure of welfare. Section 6 concludes.

aggregate gains from input trade turn out to be too large.

⁸See in particular Eaton et al. (2011), Arkolakis (2010), Arkolakis and Muendler (2011), and Bernard et al. (2012) for a recent survey.

⁹Kasahara and Rodrigue (2008) study the effect of imported intermediates on firm productivity through a production function estimation exercise. See also the recent survey in De Loecker and Goldberg (2013) for a more general empirical framework to study firm performance in international markets.

2 Why Firm-Based Models of Importing?

In this section, we present data on firms' heterogeneous import behavior that is informative about the aggregate consequences of input trade. We rely on data from the population of manufacturing firms in France.¹⁰ In Figure 1, we display the cross-sectional distribution of importers' domestic shares, i.e. the share of material spending allocated to domestic inputs. 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%. This heterogeneity in import intensities is at odds with aggregate models which presume that import shares are equalized across importers. To rationalize the data of Figure 1, we therefore have to resort to firm-based models of importing.

[Figure 1 here]

In this paper, we show that the dispersion in firms' import exposure documented in Figure 1 has aggregate implications. The intuition is simple. As a firm's domestic share measures the extent to which it benefits from foreign input sourcing, Figure 1 shows that the gains from input trade are heterogeneous at the micro-level. To correctly aggregate these producer gains, we have to know firms' relative importance in the economy. In particular, the consumer gains from input trade will be high whenever intense importers, i.e. firms with low domestic shares, are large. Figure 2 displays the extent to which this is the case in France. In the left panel, we depict the distribution of value added by import status. While importers are significantly larger than non-importers, there is ample overlap in their distribution of value added. In the right panel, we focus on the population of importers and show the distribution of domestic shares for different value added quantiles. The relationship between firms' import intensity and size is essentially flat and there is substantial dispersion in import shares conditional on size.

[Figure 2 here]

These patterns are important for our understanding of input trade. Holding the marginal distribution of domestic spending displayed in Figure 1 fixed, the gains from input trade would be higher if import intensity and firm size were more tightly linked. The joint distribution of domestic shares and value added therefore contains important information about the normative implications of input trade. In the next section, we make these statements precise and derive a simple formula to quantify the effect of input trade on consumer prices that only relies on the data displayed in Figures 1 and 2.

3 Theory

In this section, we lay out the theoretical framework of importing that we use to quantify the gains from input trade. In Section 3.1, we study the firm's import problem and formally show our unit

¹⁰We describe the dataset in more detail in Section 4.1 below.

cost sufficiency result. In Section 3.2, we embed the firm problem into a general equilibrium trade model with input-output linkages to quantify the effect of input trade on consumer prices.

3.1 The Producer Gains from Input Trade

Consider the problem of a firm, which we label as i , that uses local and foreign inputs according to the following production structure:

$$y = \varphi_i f(l, x) = \varphi_i l^{1-\gamma} x^\gamma \quad (1)$$

$$x = \left(\beta_i (q_D z_D)^{\frac{\varepsilon-1}{\varepsilon}} + (1 - \beta_i) x_I^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}} \quad (2)$$

$$x_I = h_i \left([q_{ci} z_c]_{c \in \Sigma_i} \right). \quad (3)$$

where $\gamma, \beta_i \in (0, 1)$ and $\varepsilon > 1$.¹¹ The firm combines intermediate inputs x with primary factors l , which we for simplicity refer to as labor, in a Cobb-Douglas fashion with efficiency φ_i .¹² Intermediate inputs are a CES composite of a domestic variety, with quantity z_D and quality q_D , and a foreign input bundle x_I , with relative efficiency for domestic inputs given by β_i . The firm has access to foreign inputs from multiple countries, whose quantity is denoted by $[z_c]$, which may differ in their quality $[q_{ci}]$, where c is a country index.¹³ Foreign inputs are aggregated according a constant returns to scale production function $h_i(\cdot)$.¹⁴ An important endogenous object in the production structure is the set of foreign countries the firm sources from, which we denote by Σ_i and henceforth refer to as the sourcing strategy. We do not impose any restrictions on how Σ_i is determined.

As far as the market structure is concerned, we assume that the firm faces prices of domestic and foreign inputs ($p_D, [p_{ci}]$) as parametric, i.e. it can buy any quantity at given prices. Similarly, we assume that labor can be hired frictionlessly at a given wage w . On the output side, we do not impose any restrictions, i.e. we do not specify whether firms produce a homogeneous or differentiated final good and how they compete.

The setup above describes a class of firm-based models that have been used in the literature. In particular, it for example nests the contributions by Gopinath and Neiman (2014), Halpern et al. (2011), Antràs et al. (2014), Kasahara and Rodrigue (2008), Amiti et al. (2014) and Goldberg et al. (2010).¹⁵ In this class of models, firms engage in input trade because it lowers their unit cost of

¹¹While the case of $\varepsilon \leq 1$ can also be accommodated by the theory, it implies that all firms are importers - a feature that is inconsistent with the data.

¹²We consider a single primary factor for notational simplicity. It will be clear below that our results apply to $l = g(l_1, l_2, \dots, l_T)$, where $g(\cdot)$ is a constant returns to scale production function and l_j are primary factors of different types. In the empirical application of Section 4, we consider labor and capital.

¹³We discuss below how to generalize the results of this section when the Cobb-Douglas and CES functional forms in (1)-(2) are not satisfied. We also consider the case where firms can source multiple products from different countries.

¹⁴Note that this setup nests the canonical Armington structure where all countries enter symmetrically in the production function. Additionally, this setup allows for an interaction between quality flows and the firm's efficiency, i.e. a form of non-homothetic import demand that is consistent with the findings in Kugler and Verhoogen (2011) and Blaum et al. (2013).

¹⁵While Antràs et al. (2014) consider a model of importing in the spirit of Eaton and Kortum (2002) instead of a variety-type model, the Fréchet assumption implies that these models are isomorphic.

production via love of variety and quality channels. Crucially, the assumptions made above, most importantly parametric prices and constant returns to scale, guarantee that the unit cost is constant *given the sourcing strategy* Σ . This property allows us to characterize the unit cost without solving for the extensive margin. Formally, the unit cost is given by

$$u(\Sigma_i; \varphi_i, \beta_i, [q_{ci}], [p_{ci}], h_i) \equiv \min_{z,l} \left\{ wl + p_D z_D + \sum_{c \in \Sigma_i} p_{ci} z_c \text{ s.t. } \varphi_i l^{1-\gamma} x^\gamma \geq 1 \right\}, \quad (4)$$

subject to (2)-(3). For simplicity, we refer to the unit cost as u_i . Standard calculations imply that there is an *import price index* given by

$$A(\Sigma_i, [q_{ci}], [p_{ci}], h_i) \equiv \frac{m_I}{x_I}, \quad (5)$$

where m_I denotes import spending and x_I is the foreign import bundle defined in (3). Importantly, conditional on Σ_i , this price-index is exogenous from the point of view of the firm and we henceforth denote it by $A_i(\Sigma_i)$. Next, given the CES production structure between domestic and foreign inputs, the price index for intermediate inputs is given by

$$Q_i(\Sigma_i) = \left(\beta_i^\varepsilon (p_D/q_D)^{1-\varepsilon} + (1-\beta_i)^\varepsilon A_i(\Sigma_i)^{1-\varepsilon} \right)^{\frac{1}{1-\varepsilon}}, \quad (6)$$

so that intermediate inputs $x = m/Q_i(\Sigma_i)$ where m denotes total spending in materials. It follows that the firm's unit cost is given by¹⁶

$$u_i = \frac{1}{\varphi_i} w^{1-\gamma} Q_i(\Sigma_i)^\gamma. \quad (7)$$

We see that input trade affects the unit cost through the price index for intermediate inputs. This price index, however, depends on a number of unobserved parameters related to the trading environment, e.g. the prices and qualities of the foreign inputs. We use the fact that the unobserved price index $Q_i(\Sigma_i)$ is related to the observed expenditure share on domestic inputs s_{Di} via

$$s_{Di} = Q_i(\Sigma_i)^{\varepsilon-1} \beta_i^\varepsilon \left(\frac{q_D}{p_D} \right)^{\varepsilon-1}. \quad (8)$$

Substituting (8) into (7) yields

$$u_i = \frac{1}{\tilde{\varphi}_i} \times (s_{Di})^{\frac{\gamma}{\varepsilon-1}} \times \left(\frac{p_D}{q_D} \right)^\gamma w^{1-\gamma}, \quad (9)$$

where $\tilde{\varphi}_i \equiv \varphi_i \beta_i^{\frac{\varepsilon\gamma}{\varepsilon-1}}$. (9) is a sufficiency result: conditional on the firm's domestic expenditure share s_{Di} , no aspects of the import environment, including the sourcing strategy Σ_i , the prices p_{ci} , the qualities q_{ci} or the technology h_i , affect the firm's unit cost. With (9) at hand, we can derive the

¹⁶With 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 (7).

effect of input trade on the firm’s unit cost, which is sometimes referred to as the productivity gains from importing.

Proposition 1. *Consider the model above. We define the producer gains from input trade as the reduction in unit cost relative to autarky holding prices fixed, i.e. $G_i \equiv \ln \left(\frac{u_i^{Aut}}{u_i} \right) \Big|_{p_D, w}$. Then*

$$G_i = \frac{\gamma}{1 - \varepsilon} \ln(s_{Di}). \quad (10)$$

Proof. Follows directly from (9) and the fact that the domestic share in autarky is unity. \square

Proposition 1 shows that the effect of participating in international input markets on the firm’s unit cost is observable given data on its domestic share and values of the elasticities γ and ε .¹⁷ More precisely, the increase in production costs that firm i would experience if it (and only it) was excluded from international markets can be recovered from the firm’s domestic expenditure share. Intuitively, input trade benefits the firm by reducing the price index of intermediate inputs Q_i . Conditional on an import demand system, we can invert the change in this price index from the change in the domestic expenditure share - see (8).¹⁸ Because in general p_D and w may change when the economy moves to input autarky, Proposition 1 is a partial equilibrium result. We explicitly allow for general equilibrium effects in Section 3.2 below. We note, however, that (10) identifies the dispersion of the producer gains across firms in general equilibrium.¹⁹ In this way, we can assess the distributional effects of input trade and determine whether particular firm characteristics are associated with larger gains.

The sufficiency result in Proposition 1 allows us to measure the change in the unit cost without specifying several components of the theory. As equations (5)-(7) show, the firm’s unit cost depends on the import environment parameters $[p_{ci}, q_{ci}, h_i, \beta_i]$. The domestic expenditure share conveniently encapsulates all the information from these parameters that is relevant for the unit cost - see (9). Instead, the standard approach in the literature consists of estimating these parameters in the context of a fully-specified model of importing. This approach requires researchers to specify the entire import environment, including the structure of output markets, and to solve for firms’ optimal sourcing strategies which, as discussed below, can be a non-trivial problem. Hence Proposition 1 is useful because it allows us to bypass the challenges of firm-based models of importing and quantify the producer gains in a wide class of models.

Finally, we note that Proposition 1 can be used to analyze counterfactuals other than input autarky. Consider for concreteness an episode of trade liberalization (e.g. Chile in 1980s (Pavcnik,

¹⁷In Section 8.1 of the Appendix, we generalize Proposition 1 in three ways. First, we derive a local version of (9) for the case where domestic and foreign inputs are not combined in a CES fashion. Second, we consider the case where the output elasticity of material inputs is not constant. Finally, we allow firms to source multiple products from different countries. We also discuss what additional information, relative to Proposition 1, is required to perform counterfactual analysis.

¹⁸Hence, Proposition 1 is akin to a firm-level analogue of Arkolakis et al. (2012). In the same vein as consumers gain purchasing power by sourcing cheaper or complementary products abroad, firms can lower the effective price of material services by tapping into foreign input markets.

¹⁹This follows from the fact that the relative unit cost u_i/u_j does not depend on prices p_D, w - see (9).

2002), Indonesia in the late 1980s and early 1990s (Amiti and Konings, 2007) or India in the 1990s (De Loecker et al., 2012)). The associated change in the unit cost holding prices fixed is then given by

$$\ln \left(\frac{u'_i}{u_i} \right) \Big|_{p_D, w} = \frac{\gamma}{1 - \varepsilon} \ln \left(\frac{s_{Di}}{s'_{Di}} \right), \quad (11)$$

where s'_{Di} denotes firm i 's domestic share after the shock. (11) can be used for a structural evaluation of changes in trade policy, as long as data on firms' domestic shares before and after is available.^{20,21} In particular, (11) identifies the dispersion in the policy-induced changes in unit costs across firms, that is, the distributional effects of the policy. Note that (11) contains both the exogenous change in foreign prices due to lower trade barriers as well as the endogenous change from adjustments in the sourcing pattern.

An Example with Fixed Costs. To compare our approach with the existing literature, consider the following example of a static economy where international sourcing is limited by the presence of fixed costs. In particular, suppose that sourcing an input from country c entails paying a fixed cost f_{ci} in units of labor. The profit maximization problem is then given by

$$\pi_i \equiv \max_{\Sigma, y} \left\{ (p(y) - u_i) y - w \sum_{c \in \Sigma} f_{ci} \right\}, \quad (12)$$

where the unit cost is

$$u_i = \frac{1}{\varphi_i} w^{1-\gamma} \left[\beta_i^\varepsilon (p_D/q_D)^{1-\varepsilon} + (1 - \beta_i)^\varepsilon A_i (\Sigma_i)^{1-\varepsilon} \right]^{\frac{\gamma}{1-\varepsilon}}, \quad (13)$$

and $p(y)$ denotes the demand function. Firms choose their size y and set of imported varieties Σ to maximize profits.

Albeit conceptually easy, solving this profit maximization problem presents us with two practical challenges. First, one has to specify the entire set of structural primitives of the model, including the distribution of prices, qualities and fixed costs across countries, the demand function and the structure of output markets. Second, even after making such assumptions, the choice of the optimal sourcing strategy can be computationally difficult.²² The reason is the interdependence between

²⁰This 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 any empirical analysis trying to infer the causal effect of trade liberalization.

²¹Opening up to trade might induce firms' to engage in productivity enhancing activities that directly increase efficiency φ , such as R&D. Such increases in complementary investments are not encapsulated in (11) nor in Proposition 1, which only measure the static gains from trade holding efficiency fixed. To disentangle the dynamic from the static gains from trade, more structure and data is required - see for example Eslava et al. (2014).

²²Note that the extensive margin problem cannot be 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. The reason is that, to evaluate (13), one needs to know the full set of structural parameters. While these parameters can be in principle estimated, such estimation would typically entail solving for the optimal sourcing set in (12).

entry decisions in different import markets.²³ When imported varieties are imperfect substitutes, the cost reduction associated with entering a particular foreign market depends on the quantities sourced from all other markets - see (5). If foreign inputs differ in both quality and fixed costs, the profit maximization problem in (12) is in general non-convex and the choice of the optimal sourcing set requires evaluating all possible sourcing strategies, entailing substantial computational burden - see Antràs et al. (2014) for conditions under which this issue can be sidestepped and a solution algorithm.

The benefit of Proposition 1 is that these challenges can be bypassed for certain normative questions. With micro-data on domestic expenditure shares and the two structural parameters γ and ε , we can directly measure the endogenous reduction in unit cost arising from input trade at the firm-level. Not only is the calculation straight-forward but it does not rely on any assumptions made to make the solution to (12) feasible.

3.2 The Consumer Gains from Input Trade

In this section, we embed the model of firm behavior of Section 3.1 in a macroeconomic environment and study the aggregate effect of input trade. We focus on the change in consumer prices, i.e. how much more would domestic consumers pay for the locally produced goods if firms were not allowed to source their inputs from abroad.²⁴ To isolate the effect of input trade, we abstract from trade in final goods. That is, we consider an environment where domestic consumers solely benefit from trade openness indirectly through firms' cost reductions. The micro result in Proposition 1 above is crucial as it allows us to measure such firm-level unit cost reductions in the data. To aggregate these producer gains, we need to take a stand on two aspects of the macroeconomic environment: (i) the nature of input-output linkages across firms and (ii) the degree of pass-through, which depends on consumers' demand system and output market structure. While the former determines the effect of trade on the price of domestic inputs p_D , the latter determines how much of the trade-induced cost reductions actually benefit consumers.

We consider the following multi-sector CES monopolistic competition environment, which is for example also used in Caliendo and Parro (2015).²⁵ There are S sectors, each comprised of a measure N_s of firms which we treat as fixed. There is a unit measure of consumers who supply L units of

²³This interdependence 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, and has the implication that more productive firms need not source their inputs from more countries, unless more restrictions are imposed.

²⁴Note that the change in consumer prices does not capture the full effect of input trade on welfare. For the latter, we need to take into account the resources spent by firms to attain their sourcing strategies and we do so in Section 5 below.

²⁵We adopt a model with perfect pass-through because we lack data on firm-specific prices, which would be necessary to discipline the extent of pass-through in a more general framework.

labor inelastically and whose preferences are given by

$$U = \prod_{s=1}^S C_s^{\alpha_s} \quad (14)$$

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

where $\alpha_s \in (0, 1)$, $\sum_s \alpha_s = 1$ and $\sigma_s > 1$. Firm i in sector $s = 1, \dots, S-1$ produces according to the production technology given by (1)-(3) in Section 3.1 above, where the structural parameters ε and γ are allowed to be sector-specific. As before, we do not assume any particular mechanism of how the extensive margin of trade is determined nor impose any restrictions on $[p_{ci}, q_{ci}, h_i, \beta_i]$. That is, the distribution of prices and qualities across countries and the aggregator of foreign inputs can take any form. Additionally, these parameters can vary across firms in any way. We assume sector S to be comprised of firms that do not trade inputs and refer to it as the non-manufacturing sector.²⁶

We assume the following structure of roundabout production. Firms use a sector-specific domestic input that is produced using the output of all other firms in the economy according to

$$z_{Ds} = \prod_{j=1}^S Y_j^{\zeta_j^s} \text{ and } Y_j = \left(\int_0^{N_j} y_{\nu js}^{\frac{\sigma_j-1}{\sigma_j}} d\nu \right)^{\frac{\sigma_j}{\sigma_j-1}}, \quad (16)$$

where z_{Ds} denotes the bundle of domestic inputs, ζ_j^s is a matrix of input-output linkages with $\zeta_j^s \in [0, 1]$ for all s and j and $\sum_{j=1}^S \zeta_j^s = 1$ for all s , and $y_{\nu js}$ is the output of firm ν in sector j demanded by a firm in sector s . In this setting, the price of the domestic input p_{Ds} is endogenous so that domestic firms are affected by trade policy via their purchases of intermediate inputs from importers.

Building on our result from Section 3.1, we now show that the consumer price index associated with (14)-(15) can be expressed in terms of observables. Given the CES demand and monopolistic competition structure, the consumer price index for sector s is given by

$$P_s = \mu_s \left(\int_0^{N_s} u_i^{1-\sigma_s} di \right)^{\frac{1}{1-\sigma_s}} = \mu_s \left(\frac{p_{Ds}}{q_{Ds}} \right)^{\gamma_s} \times \left(\int_0^{N_s} \left(\frac{1}{\tilde{\varphi}_i} (s_{Di})^{\gamma_s/(\varepsilon_s-1)} \right)^{1-\sigma_s} di \right)^{\frac{1}{1-\sigma_s}}, \quad (17)$$

where $\mu_s \equiv \sigma_s/(\sigma_s - 1)$ is the mark-up in sector s and we treat labor as the numeraire. The second equality follows from (9) above which allows us to express firms' unit costs in terms of their domestic expenditure shares (s_{Di}) and efficiency ($\tilde{\varphi}_i$). (17) shows that, holding domestic input prices fixed, the effect of input trade on consumers' purchasing power is an efficiency-weighted average of the

²⁶We introduce this sector for empirical reasons. In the next section, we consider an application to France where we do not have data on firm-level imports outside of the manufacturing sector. To make aggregate statements about input trade, we take the non-manufacturing sector into account. See Section 4 for details.

firm-level gains. While firm efficiency $\tilde{\varphi}_i$ is not observed, it can be recovered up so scale from data on value added and domestic spending as²⁷

$$va_i \propto \left(\tilde{\varphi}_i (s_{Di})^{\gamma_s / (1 - \varepsilon_s)} \right)^{\sigma_s - 1}. \quad (18)$$

Combining (17) and (18), the change in the sectoral consumer price index relative to autarky is given by

$$\ln \left(P_s^{Aut} / P_s \right) = \gamma_s \ln \left(p_{Ds}^{Aut} / p_{Ds} \right) + \Lambda_s, \quad (19)$$

where

$$\Lambda_s = \frac{1}{1 - \sigma_s} \ln \left(\int_0^{N_s} \omega_i s_{Di}^{\frac{\gamma_s}{1 - \varepsilon_s} (1 - \sigma_s)} di \right) \quad (20)$$

and ω_i denotes firm i 's share in value added. (19) shows that input trade affects consumer prices through two channels. First, there is a direct effect stemming from firms in sector s sourcing inputs internationally, Λ_s . Second, there is an indirect effect as the price of domestic inputs changes because of input-output linkages, p_{Ds}^{Aut} / p_{Ds} .

(19) and (20) contain a sufficiency result for the change in consumer prices. Note first that the direct price reduction Λ_s can be computed with data on value added and domestic shares. Next, because of the structure of roundabout production in (16), the change in domestic input prices p_{Ds}^{Aut} / p_{Ds} is a function of the Λ_s of all sectors. Hence, the consumer price gains from input trade can be expressed in terms of observables.

Proposition 2. *Let P and P^{Aut} be the consumer price indices in the trade equilibrium and autarky. We define the consumer price gains from input trade as the reduction in the consumer price index relative to autarky, i.e. $G \equiv \ln \left(P^{Aut} / P \right)$. Then,*

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

where $\Lambda = [\Lambda_1, \Lambda_2, \dots, \Lambda_S]$, $\Xi = [\zeta_j^s]$ 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 $\Gamma = \text{diag}(\gamma)$, where γ is the $S \times 1$ vector of input intensities.

Proof. See Section 8.2 in the Appendix. □

Proposition 2 is the main result of the paper. It shows that the information contained in the micro-data on domestic spending and value added is sufficient to characterize the consumer price gains from input trade relative to autarky in the class of models considered in this section. Thus, the consumer price gains can essentially be read off the micro-data given the parameters for consumer demand and production. Information about firms' import environment or firms' endogenous choice of their extensive margin of importing is not required.

²⁷This assumes that the data on value added does not record firms' expenses to attain their sourcing strategies. If it did, one could express (18) in terms of sales or employment.

To understand Proposition 2, it is instructive to consider the case of a single sector economy. Expression (21) then becomes

$$G = \frac{\Lambda}{1 - \gamma}, \quad (22)$$

that is, the consumer price gains are simply given by the direct price reduction Λ , inflated by $1/(1 - \gamma)$ to capture the presence of roundabout production.

Finally, Proposition 2 can be generalized to study counterfactuals beyond input autarky. In particular, consider a policy that changes firms' domestic expenditure shares from $[s_{Di}]$ to $[s'_{Di}]$. The effect of such policy on consumer prices is given by the same expression as in Proposition 2, except that Λ_s is now given by

$$\Lambda_s^* = \frac{1}{1 - \sigma_s} \ln \left(\int_0^{N_s} \omega_i \left(s_{Di}/s'_{Di} \right)^{\frac{\gamma_s}{1 - \varepsilon_s} (1 - \sigma_s)} di \right). \quad (23)$$

In the case of observed policy experiments, the consumer gains can be easily computed as long as data on domestic shares before and after the change is available. (23) is also useful for unobserved counterfactuals. In particular, all models within our class will have the exact same normative implications as long as they generate the same counterfactual distribution of domestic shares. While the underlying import environment matters for the predicted domestic shares, conditional on such predictions the implied consumer gains are same.

The Bias of Models of Importing

Proposition 2 is a useful organizing tool for the existing models of importing. It shows that, in terms of their normative implications, existing models differ only in their implied distribution of domestic shares and value added which translate into different price reductions Λ .

Aggregate Models. Consider first the aggregate models of importing where firms' domestic expenditure shares are equalized - see Eaton et al. (2011), Caliendo and Parro (2015) and Costinot and Rodríguez-Clare (2014).²⁸ In these models, the direct price reductions relative to autarky are given by

$$\Lambda_s^{Agg} = \frac{\gamma_s}{1 - \varepsilon_s} \ln \left(s_{Ds}^{Agg} \right) = \frac{\gamma_s}{1 - \varepsilon_s} \ln \left(\int_0^{N_s} \omega_i s_{Di} di \right), \quad (24)$$

where s_{Ds}^{Agg} is the aggregate domestic expenditure share in sector s .²⁹ While these frameworks have the benefit of only requiring aggregate data, Figure 1 in Section 2 shows that their implication of equalized domestic shares is rejected in the micro-data, and Proposition 2 shows that such deviation has aggregate consequences. In particular, (20) and (24) imply that the bias from measuring the

²⁸In our setup, domestic shares are equalized if for example firms have the same sourcing strategy (e.g. in the absence of fixed costs) and prices, qualities and technology do not vary across firms.

²⁹Note that, because of Cobb-Douglas production, firm value added is proportional to material spending, so that s_{Ds}^{Agg} is indeed equal to the aggregate share of material spending allocated towards domestic producers.

price reduction in sector s through the lens of an aggregate model is given by

$$Bias_s \equiv \Lambda_s^{Agg} - \Lambda_s = \frac{\gamma_s}{\varepsilon_s - 1} \times \ln \left[\frac{\left(\int_0^{N_s} \omega_i s_{Di}^{\chi_s} di \right)^{1/\chi_s}}{\int_0^{N_s} \omega_i s_{Di} di} \right], \quad (25)$$

where $\chi_s = \frac{\gamma_s(\sigma_s - 1)}{\varepsilon_s - 1}$. Heterogeneity in import shares induces a bias in the estimates of the gains from trade of aggregate models, as long as $\chi_s \neq 1$. The magnitude of the bias depends on the underlying dispersion in domestic shares and their correlation with firm size - we quantify it in our empirical application below. The sign of the bias, however, depends only on parameters and not on the underlying micro-data. In particular, (25) together with Jensen's inequality directly imply that

$$Bias_s > 0 \text{ if and only if } \chi_s = \frac{\gamma_s(\sigma_s - 1)}{\varepsilon_s - 1} > 1. \quad (26)$$

The economic intuition behind (26) is as follows. Because the current trade equilibrium is observed in the data, quantifying the gains from trade boils down to predicting consumer prices in the counterfactual autarky allocation - see (17) and (18). Such prices are fully determined from producers' efficiencies, i.e. $\tilde{\varphi}_i^{\sigma-1}$. As these are unobserved, they are inferred from data on value added and domestic shares. More specifically, given the data on value added, (18) shows that $\tilde{\varphi}_i^{\sigma-1}$ is proportional to s_{Di}^χ . In the same vein as dispersion in prices is valued by consumers whenever demand is elastic, dispersion in domestic shares is valued as long as $\chi > 1$. In this case, the autarky price index inferred by an aggregate model is too high, making the gains from trade *upward* biased.

To fix ideas, consider an example where firms differ in their domestic shares and value added is equalized across producers. In this case, an aggregate model would conclude that efficiency is also equalized across firms - see (18). This, however, cannot be the case as the dispersion in domestic shares implies that efficiency has to vary given a common level of value added. Whether or not consumers prefer the autarky allocation with equalized efficiency depends on χ . If $\chi > 1$, such allocation features higher consumer prices and therefore higher gains from trade.

Firm-based Models. On the other side of the spectrum are firm-based models of importing. These models generate heterogeneity in firms' import shares, typically via sorting into different import markets, thereby inducing a joint distribution of import intensity and size. Gopinath and Neiman (2014), Amity et al. (2014) and Ramanarayanan (2014) for example assume that firms differ only in their efficiency and thus generate a perfect negative correlation between domestic shares and value added conditional on importing. They also imply that all importers are larger than domestic firms. By assigning the largest unit cost reductions to the most efficient firms, this tends to magnify the aggregate gains from trade. Figure 2 in Section 2, however, shows that the correlation between firm size and domestic spending is negative but far from perfect, and that many importers are small. Because models with a single source of firm heterogeneity cannot match these features of

the data, they will tend to yield biased estimates of the gains from trade.³⁰ Antràs et al. (2014) and Halpern et al. (2011) allow for heterogeneity in efficiency and fixed costs and thus generate a non-trivial distribution of value added and domestic spending. Whether or not the model-implied distribution is quantitatively consistent with the micro-data and hence informative for normative questions depends on the particular calibration.

4 Quantifying the Producer and Consumer Gains

We now take the framework laid out above to data on French firms to quantify the gains from input trade both at the firm and aggregate level. Implementing Propositions 1 and 2 empirically requires a set of parameters. We deal with their estimation in Section 4.1 and compute the producer and consumer gains in Section 4.2.

4.1 Estimation of Parameters

Our approach relies on both micro and aggregate data. We use the micro-data to estimate the production function parameters, i.e. the material elasticities $[\gamma_s]$ and the elasticities of substitution $[\varepsilon_s]$, as well as the sector-specific demand elasticities $[\sigma_s]$. We identify the input-output structure on the production side $[\zeta_j^s]$ and the aggregate demand parameters $[\alpha_s]$ from the input-output tables. This allows us also to account for the non-manufacturing sector and doing so is quantitatively important.

Data. The main source of information we use is a firm-level dataset from France. A detailed description of how the data is constructed is contained in Section 8.3 of the Appendix. Because we are interested in trade in 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. Import flows are classified at the country-product level, where products are measured at the 8-digit (NC8) level of aggregation. 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. We augment this data with two additional data sources. First, we employ data on input-output linkages in France from the STAN database of the OECD. Second, we use global trade flows from the UN Comtrade Database to measure aggregate export supplies which we use to construct an instrument to estimate the elasticity of substitution ε below.

Identification of α , ζ and σ . We compute the demand parameters α_s and the matrix of input-output linkages $[\zeta_j^s]$ using data from the French input-output tables on the distribution of firms' intermediate spending and consumers' expenditure by sector.³¹ Sectors are classified at the 2-digit

³⁰We quantify such bias in Section 5.3 below.

³¹See the Online Appendix for a detailed description of how we construct the input-output matrix.

level. 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 , our theory implies

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

We aggregate all non-manufacturing sectors into one residual sector, which we denote by S , and construct its consumption share α_S and input-output matrix ζ_j^S directly from the Input-Output Tables.

Our dataset does not have information on firm-specific prices but only revenues. We therefore use industry-specific average mark-ups to get the demand elasticities $[\sigma_s]$. 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 calculate total costs as the sum of material spending, payments to labor and the costs of capital. We compute averages at the sector level to identify σ_s .

Table 1 below contains the results. Column three reports the consumption share α_s for each sector in France. The non-manufacturing sector is important as it account for a large share of the budget of consumers. Column four reports the demand elasticities σ_s which, consistent with the literature, are estimated at around 3. For brevity, we report the input-output matrix ζ_j^s in the Online Appendix.

[Table 1 here]

Estimation of ε and γ . Of particular importance are the elasticities of substitution ε_s and the intermediate input shares γ_s , as they directly affect the producer gains. To understand our identification strategy, note that firm output can be written as³²

$$y_{is} = \tilde{\varphi}_i s D_i^{-\frac{\gamma_s}{\varepsilon_s - 1}} k_i^{\phi_{ks}} l_i^{\phi_{ls}} m_i^{\gamma_s} \times B \quad (28)$$

where m_i is total material spending by firm i and B collects all general equilibrium variables, which are constant across firms. By expressing output in terms of *spending* in materials instead of quantities, (28) shows that we can estimate ε_s by treating the domestic share as an additional input in a production function estimation exercise.^{33,34} We also see that the domestic share is akin to a productivity shifter.

Because we do not observe firm-specific prices, we rely on the demand structure assumed in Section 3.2 and express (28) in terms of firm revenue

³²In this section, we augment the production function considered in Section 3 to include capital, i.e. $y_{is} = \varphi_i k_i^{\phi_{ks}} l_i^{\phi_{ls}} x^{\gamma_s}$, where ϕ_{ks} and ϕ_{ls} denote the capital and labor output elasticities in sector s .

³³Note that it is common in the literature to rely on material spending as a measure of input use, as quantities are rarely observed. (28) shows that in this case the domestic expenditure share turns out to be the appropriate deflator for material spending. Not controlling for the domestic share therefore results in biased estimates (De Loecker and Goldberg, 2013).

³⁴Because our data does not contain a reliable measure of quality-adjusted input prices, we do not estimate ε from the firm's import demand system (6)-(8).

$$\ln(Rev_{is}) = \delta + \tilde{\phi}_{ks}\ln(k_i) + \tilde{\phi}_{ls}\ln(l_i) + \tilde{\gamma}_s\ln(m_i) + \ln(\vartheta_i), \quad (29)$$

where the productivity residual ϑ_i is given by

$$\ln(\vartheta_i) = \frac{1}{1 - \varepsilon_s} \tilde{\gamma}_s \ln(s_{Di}) + \frac{\sigma_s - 1}{\sigma_s} \ln(\tilde{\varphi}_i) \quad (30)$$

and $\tilde{\gamma}_s = \frac{\sigma_s - 1}{\sigma_s} \gamma_s$ and $\tilde{\phi}_{ks}$ and $\tilde{\phi}_{ls}$ 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 output elasticities $[\phi_{ks}, \phi_{ls}, \gamma_s]$ are obtained from (29). We consider both a factor shares approach and a proxy method. We then use such elasticities to construct productivity residuals $\ln(\vartheta_i)$ 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 estimates all parameters in (29)-(30) simultaneously. In this approach we allow for sector-specific ε_s .

Consider first the approach based on observed factor shares, which is a simple and easy-to-implement benchmark. The Cobb-Douglas production structure implies that

$$\tilde{\gamma}_s = \frac{m_i}{p_i y_i}, \quad (31)$$

so that we can measure $\tilde{\gamma}_s$ as the average share of material spending across firms. We can similarly measure $\tilde{\phi}_{ks}$ and $\tilde{\phi}_{ls}$, and hence construct the productivity residuals $\ln(\vartheta_i)$ from (29) up to an inconsequential constant. In a second step, we then use the estimated $\tilde{\gamma}_s$, the productivity residuals and the data on domestic shares to estimate equation (30).

Clearly, we cannot estimate (30) via OLS as the required orthogonality restriction fails: s_D is not orthogonal to efficiency φ under most reasonable models of import behavior. In particular, more efficient firms are likely to sort into more and different sourcing countries and this variation in the extensive margin of trade may 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 construct from the Comtrade data. More precisely, we construct the instrument

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

where WES_{ckt} denotes the total exports of product k from county c in year t to the entire world excluding France, s_{cki}^{pre} is firm i 's import share on product k from 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 efficiency. Intuitively, if we

see China’s exports of product k increasing in year t , French importers that used to 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 for the domestic share according to the following specification

$$\Delta \ln(\hat{\vartheta}_{ist}) = \delta_s + \delta_t + \frac{1}{1 - \varepsilon} \times \Delta \tilde{\gamma}_s \ln(s_{ist}^D) + o'_{ist} \xi + u_{ist}, \quad (33)$$

where δ are sector and year fixed effects, o_{ist} is a vector of firm-level controls and $\Delta \ln(\hat{\vartheta}_{ist})$ and $\Delta \tilde{\gamma}_s \ln(s_{ist}^D)$ are the changes in firm residual productivity and domestic shares respectively, which are instrumented by (32). We define products at the 6-digit level and take firms’ 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.

As an alternative to the factor shares approach, we 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 $[\phi_{ks}, \phi_{ls}, \gamma_s]$. 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. The second step is as in our previous approach: we construct productivity residuals $\ln(\vartheta_i)$ 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 $[\phi_{ks}, \phi_{ls}, \gamma_s]$ as the factor shares approach, the implied estimate for ε would be numerically identical.

Our third method consists of estimating firms’ production function with an integrated GMM approach. Instead of treating (29) and (30) as separate estimation equations, we estimate the firms’ production function in a single step with four inputs and again follow Wooldridge (2009) to estimate the four parameters via 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 3 below. Table 2 contains the estimates of ε using the factor shares approach and the proxy method based on Levinsohn and Petrin (2012) and Wooldridge (2009).³⁵ For the latter procedure, we report the results based on both the Cobb Douglas and the more general translog specifications. In the respective first column, we show the first stage relationship between changes in world 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 total world exports reduce

³⁵For brevity, we report the estimates of the other production function parameters in the Online Appendix.

their domestic spending.³⁶ Turning to the results for ε , we see that the different procedures yield relatively similar results as the estimates lie between 1.7 and 2.4.³⁷ In particular, the point estimates remain relatively unchanged when we estimate the second stage equation on importing firms only. Note, however, that the standard errors increase substantially as we lose a large amount of data by conditioning on import status.³⁸

[Table 2 here]

[Figure 3 here]

The results of the integrated GMM approach are summarized in Figure 3. Because we estimate firms' production function for each industry, this procedure gives sector-specific estimates of ε . We depict both the point estimates 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.^{39,40}

For the remainder of the paper, we take the estimate stemming from the factor shares approach, i.e. $\varepsilon = 2.38$, as the benchmark. While the two-step approaches estimate (33) in first differences, the one-step GMM approach treats firms' domestic shares as an explicit input and estimates the production function in levels. Conceptually, we prefer the identification strategy in first differences as we find the underlying exogeneity assumptions more plausible.⁴¹ It is nevertheless comforting to see that all these approaches yield consistent results. While we lock in to the factor shares estimate, 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 .

³⁶The reason why the first stage results are not numerically equivalent across the different specifications is that the estimated material elasticity is different. Recall that the independent variable is $\Delta\tilde{\gamma}_s \ln(s_{ist}^D)$.

³⁷Our estimates are close to the ones of Antràs et al. (2014) who rely on cross-country variation.

³⁸In Section 8.4 in the Appendix, we provide further robustness checks to our estimates of ε , which lead to similar conclusions. In particular, we keep the year used for the pre-sample weights s_{cki}^{pre} fixed at 2001 for all firms.

³⁹Kasahara and Rodrigue (2008) find estimates of the elasticity of substitution in the range of 3.1 to 4.4 using a related approach for Chilean data. However, they do not use an external instrument for firms' imported intermediates. Halpern et al. (2011) use Hungarian data and derive a production function equation analog to (29)-(30), as well as an import demand equation. They find an elasticity of substitution of 7.3. The main difference with our approach is that they obtain the parameters of their structural model, namely the elasticity of substitution and the quality of foreign varieties, by *simultaneously* estimating the production function and import demand equations. Because both of these equations are derived after solving for the extensive margin of trade, they only hold under restrictive assumptions. In contrast, we identify ε solely from (29)-(30) by using exogenous variation in input supplies. Hence, we can estimate ε without taking a stand on how the extensive margin of trade is determined.

⁴⁰This suggests that the relatively low value for ε found in the pooled factor shares approach is not a result of sectoral aggregation. This is in contrast to estimations performed on aggregate data, which find a heterogeneity bias (Imbs and Mejean, 2015).

⁴¹Another 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 that all firms in such sector ought to be importers.

4.2 Results

With the structural parameters at hand, we now quantify the gains from input trade in France. We proceed as in the theory. We follow Proposition 1 and use data on domestic expenditure shares to measure the producer gains from input trade, i.e. distribution of unit cost reduction across firms. We then augment this data with information on firm size and use Proposition 2 to measure the consumer price gains from input trade. 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.

Input Trade and Producer Gains. Given our estimates of ε and γ_s and the micro-data on firms’ domestic shares, Proposition 1 states that the unit cost reductions from input trade are given by $\frac{\gamma_s}{1-\varepsilon} \ln(s_D)$. We depict these producer gains in Figure 4 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 11.2% if moved to autarky, firms above the 90th percentile of the distribution would experience losses of 85% or more. According to Proposition 1, any model within the class covered in Section 3.1 will arrive at exactly the same conclusions about the distribution of the gains from trade at the micro-level, as long as it matches the micro data on domestic shares and utilizes the same values for γ_s and ε .

[Figure 4 here]

[Table 3 here]

We can also use the micro-data to learn about firm characteristics that are correlated with the producer gains. In particular, consider the following regressions:

$$\frac{\gamma_s}{1-\varepsilon} \ln(s_{Dist}) = \delta_s + \delta_t + o'_{ist} \psi + u_{ist}, \quad (34)$$

where δ_s and δ_t are industry and time fixed effects and o_{ist} is a vector of firm characteristics. To interpret ψ , recall from (8) that the observed domestic shares can reflect firm-variation in exogenous “import capabilities” (such as prices $[p_{ci}/q_{ci}]$ or the import bias β_i) and firms’ endogenous sourcing strategies Σ_i . The results are contained in Table 4 and are intuitive. Bigger firms, as measured by either value added or employment, see higher gains. Being an exporter or a member of an international group is associated with a reduction in the unit cost of 8.5% and 14.8%, respectively. When we restrict the analysis to the sample of importers, the positive relation between firm size and the producer gains becomes substantially weaker. This is consistent with the pattern documented in Section 2 above which showed a mild correlation between import intensity and value added for importers. Next, we consider the role of the firm’s sourcing strategy, which we measure by the average number of countries that the firm sources its products from. According to the theory of Section 3, firms source their inputs internationally to reduce their unit cost. Consistent with the theory, we find a strong positive relation between firms’ extensive margin of importing and the producer gains.

Note that the importance of other firm characteristics is diminished once the number of varieties is controlled for.⁴²

[Table 4 here]

Input Trade and the Consumer Price Gains. We now have all the ingredients required by Proposition 2 to quantify the effect of input trade on consumer prices. Using firm-level data on domestic expenditure shares and value added together with the parameters estimated above, we compute the sector-specific direct price reductions Λ_s given in (20). We then calculate the consumer price gains by solving the linear system in (21). Table 5 contains the results. We find that French consumer prices in the manufacturing sector would be 27.5% higher if French producers were forced to source their inputs domestically. When the price of the non-manufacturing sector is taken into account, the consumer price gains amount to 9%.⁴³ The reason why these economy-wide gains are smaller is that the non-manufacturing sector experiences only a 3% price reduction but accounts for 70% of consumers' budget - see Table 1.

In Table 5, we also report the consumer price gains predicted by a representative-firm approach that only uses data on domestic spending at the sector level - see (24). This aggregate approach implies gains of 31.4% and 9.9% in the manufacturing sector and the entire economy, respectively. Ignoring the heterogeneity in firms' import behavior within sectors therefore results in an over-estimation of the consumer price gains by 3.4 and 1 percentage points for the manufacturing sector and the entire economy, respectively. The aggregate approach is upward biased because the estimated parameters imply that, for most sectors, Λ_s is a convex aggregator of firms' domestic shares - see (25)-(26).

[Table 5 here]

Importantly, there is a second source of bias that arises when using an aggregate approach which pertains to the "correct" elasticity of substitution ε . While we treat ε as a production function parameter and estimate it from micro-data, aggregate models often estimate ε from a gravity equation using aggregate trade flows. While there is a large literature concerning this particular parameter⁴⁴, most aggregate approaches find estimates that are larger than our preferred estimate of 2.38.⁴⁵ Costinot and Rodríguez-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

⁴²In particular, firm size is substantially negatively correlated with import spending holding the number of imported varieties fixed. This is intuitive. If a small firm decides to source from the same number of sourcing countries as a large firm, then it is likely that the small firm is a proficient importer, which manifests itself in a low share of domestic spending.

⁴³Formally, the economy-wide gains P^{Aut}/P are related to the gains in the manufacturing sector P_M^{Aut}/P_M via $P^{Aut}/P = (P_M^{Aut}/P_M)^{1-\alpha_S} (P_S^{Aut}/P_S)^{\alpha_S}$, where α_S is the expenditure share in the non-manufacturing sector.

⁴⁴See e.g. Simonovska and Waugh (2013, 2014).

⁴⁵Recall that our benchmark was chosen conservatively, as all other estimates of ε in Table 2 are smaller. See also Goldberg et al. (2010) for indirect evidence on the low substitutability between domestic and imported inputs for Indian manufacturing firms.

would lead to substantially smaller gains from trade. In Section 8.5 of the Appendix, we redo the analysis of Table 5 for a range of values of ε spanning the estimates from the literature. Moving to $\varepsilon = 4$, for example, tends to reduce the consumer price gains from trade of the aggregate approach by 50%. This “elasticity bias” can therefore be substantial.

We also use the micro-data to quantify our confidence in the estimates of the gains from input trade. Table 5 reports the 90-10 confidence intervals of the bootstrap distribution of the point estimates in italics.⁴⁶ Note that the uncertainty in the point estimates stem from two sources. First, because we base our analysis on a large but finite sample, there is uncertainty in our aggregate statistics for given parameters. Second, the structural parameters ε, γ_s and σ_s are estimated with error. We see that these two forces induce quite a bit of variation in the estimates. With 80% probability, the consumer price gains in the manufacturing sector lie between 21% and 36% and the gains for the entire economy lie between 7% and 12%.⁴⁷ We also find that the aggregate approach yields more uncertain estimates (second row) and leads to an over-estimation of the gains with 80% probability (third row). A graphical depiction of this sampling uncertainty is contained in Figure 5. We see that the bootstrap distribution of the consumer price gains using the aggregate approach features a thicker right tail (upper panels) and the resulting bias has the majority of its mass on positive numbers (lower panels).

[Table 6 here]

Table 6 reports the gains by sector and provides a decomposition to isolate the importance of production linkages across sectors. We first report the sectoral consumer price gains, P_s^{Aut}/P_s , which measure the change in the price of the output bundle of sector s . We find substantial heterogeneity in the effect of input trade across sectors: while e.g. prices for textile products would be 56% higher if producers were not allowed to source their inputs from abroad, this effect is only 18% for metal products. We then decompose these price changes into the direct price reduction from firms in sector s sourcing internationally, Λ_s , and the indirect gains stemming from firms in sector s buying domestic inputs from other firms who in turn may engage in trade, $p_{D_s}^{Aut}/p_{D_s}$.⁴⁸ We find that interlinkages are important as they account for roughly 50% of the sectoral price gains. Note also that the importance of interlinkages varies substantially across industries as a result of the underlying heterogeneity in the input-output matrix: sectors that rely on relatively open sectors more intensively benefit more from input trade as their upstream suppliers experience larger unit cost reductions.

We also assess the importance of interconnections by considering the case with no cross-industry input-output linkages, i.e. where each sector uses only its own products as inputs.⁴⁹ In this case, we

⁴⁶We explain the details of the bootstrap procedure in the Online 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 redo the analysis of Section 4.1 and obtain new estimates for the structural parameters. Finally, for each iteration, we recalculate the consumer price gains and the other statistics of interest.

⁴⁷Given the large sample size, most of the uncertainty stems from the variation in the structural parameters and not from the re-sampling of firms. See Figure 8 in the Online Appendix.

⁴⁸Formally, p_{D_s} is the sector-specific price index of a unit of the bundle of domestic inputs, which is an aggregator of all the goods produced locally, see (16).

⁴⁹In this case, the matrix of input-output linkages is given by $\zeta_j^s = 0$ for $j \neq s$ and $\zeta_j^j = 1$.

find a point estimate for the consumer prices gains from trade of

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

That is, shutting down input trade would increase consumer prices by 12%. Compared to the actual gains of 9%, the economy without interlinkages over-estimates the aggregate gains by about a third. 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. As this sector is not a direct beneficiary of foreign sourcing, such linkages actually dampen the aggregate effect of input trade.

Finally, Table 6 also contains the direct price reductions that arise from a representative firm model, Λ_s^{Agg} . In line with the results of Table 5, in 12 of the 18 manufacturing sectors the gains based on aggregate data are upward biased. The reason for this pattern goes back to the condition in (26) which characterizes the sign of the bias as a function of parameters. It turns out that for most sectors the estimated σ_s and γ_s satisfy the condition for the bias to be positive. Note also that the bias can be quite substantial. Consider for example the office and computing machinery sector. While the aggregate approach would imply a direct price reduction of 37%, the exact firm-based formula tells us that this number should be only 20%.

5 Input Trade and Welfare

So far, we quantified the effect of input trade on consumer prices. By exploiting the sufficiency results in Propositions 1 and 2, we were able to measure the change in consumer prices directly from the micro-data. These consumer price gains, however, may not capture the full impact of input trade on welfare as they do not take into account the resources spent by firms to attain their sourcing strategies. To measure the welfare consequences of input trade, we therefore need to take a stand on how firms determine their extensive margin of trade and quantify such resource loss in the context of a calibrated model. To this end, we consider a model where foreign sourcing is limited by the presence of fixed costs and calibrate it to French micro-data.

5.1 A Model of Fixed Costs

We now go back to the example with fixed costs of Section 3 - see (12). In this environment, firms choose their sourcing strategy by trading off the import-induced reduction in unit costs vs the payment of fixed costs. For expositional simplicity, we consider a one-sector version of the model.⁵⁰ As discussed above, computing firms' optimal sourcing strategies can be challenging when prices, qualities and fixed costs vary by country in an arbitrary way. To ensure tractability, we assume that the fixed cost of sourcing is constant across countries, i.e. $f_c = f$ for all c . In this case, the firm selects its sourcing countries based purely on price-adjusted quality and the sourcing strategy

⁵⁰See the Online Appendix for the analysis with multiple sectors.

reduces from a set Σ to a scalar, a price-adjusted quality cutoff.⁵¹ Furthermore, we assume there is a continuum of countries so that this cutoff can be characterized by a first order condition. Finally, we impose the following functional form assumptions that ensure a parsimonious characterization of the firm’s problem.

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

1. *Imported inputs are combined according to:*

$$x_I = \left(\int_{c \in \Sigma} (q_c z_c)^{\frac{\kappa-1}{\kappa}} dc \right)^{\frac{\kappa}{\kappa-1}}. \quad (35)$$

2. *Country quality is Pareto distributed:*

$$G(q) = \Pr(q_c \leq q) = 1 - (q_{min}/q)^\theta \text{ for } q \geq q_{min},$$

and $\theta > 1$ and $q_{min} > 0$.

3. *Foreign prices p_c and qualities q_c are related by $p_c = q_c^\nu$ where $\nu \leq 1$.*

4. *The following parametric condition is satisfied: $(\kappa - 1)(1 - \nu) < \theta$.*

The isoelastic relation between prices and qualities ensures that price-adjusted qualities $q/p = q^{1-\nu}$ are increasing in country quality so that firms import high quality inputs. This, together with the Pareto distribution for country quality and the CES aggregator for the foreign bundle, imply that the import price index, given by equation (5) above, takes a convenient power form:

$$A(\Sigma) = \left(\int_{c \in \Sigma} (p_c/q_c)^{1-\kappa} dc \right)^{\frac{1}{1-\kappa}} = zn^{-\eta} \equiv A(n), \quad (36)$$

where n is the share of countries the firm sources foreign inputs from⁵² and z and η are “auxiliary” parameters which depend on the underlying parameters governing import prices ν , the distribution of quality (q_{min}, θ) and the elasticity of substitution across foreign varieties κ .⁵³ Thus, the underlying structure of the import environment matters for the firm’s problem *only* through z and η . In this way, knowledge of the deep parameters $(q_{min}, \theta, \kappa, \nu)$ is irrelevant for all aggregate outcomes as long as (z, η) are known.⁵⁴

Under the above assumptions, the firm’s profit maximization problem is given by:

⁵¹More precisely, if country c with price-adjusted quality q_c/p_c is an element of Σ so are all countries c' with $q_{c'}/p_{c'} > q_c/p_c$. See Antràs et al. (2014) for a solution method that does not require this assumption.

⁵²Note that, given the distribution of country quality, the firm’s sourcing strategy can be equivalently described by the mass of countries sourced n or a quality cutoff \bar{q} .

⁵³See Section 8.6 in the Appendix for the derivation of (36) and precise expressions for z and η . Note that the parametric condition in item 4 of Assumption 1 is required to compute the integral in (36).

⁵⁴This, however, does not mean that the degree of quality heterogeneity (θ) or the substitutability of inputs (κ) do not have a role in shaping firms’ import demand. For example, it can be easily shown that quality diversity and input substitutability decrease z .

$$\pi = \max_n \left\{ u(n)^{1-\sigma} \times B - w(nf + f_I \mathbb{I}(n > 0)) \right\}, \quad (37)$$

where f_I is a fixed cost to start importing⁵⁵, $\mathbb{I}(\cdot)$ is an indicator of import status, and the unit cost function is given by

$$u(n) \equiv \frac{1}{\varphi} w^{1-\gamma} \left[(p_D/q_D)^{1-\varepsilon} + z^{1-\varepsilon} n^{-\eta(1-\varepsilon)} \right]^{\frac{\gamma}{1-\varepsilon}}, \quad (38)$$

and $B \equiv \frac{1}{\sigma} \left(\frac{\sigma}{\sigma-1} \right)^{1-\sigma} P^{\sigma-1} S$, with P and S denoting the consumer price index and aggregate spending, which are determined in general equilibrium.⁵⁶ (37)-(38) show how the assumption of constant fixed costs across countries substantially simplifies the extensive margin problem: conditional on importing, the optimal sourcing strategy n is given by the solution to a first order condition. In the Appendix, we fully characterize the solution to this problem. There we also show that, conditional on importing, the optimal mass of sourcing countries n is increasing in φ and decreasing in f .⁵⁷

As argued above, the joint distribution of firm size and import intensity contains crucial information about the gains from input trade. To generate the rich distribution shown in Figure 2 of Section 2, we have to allow for (at least) two sources of firm heterogeneity - we choose efficiency φ_i and fixed costs f_i . As φ_i and the endogenous unit costs reduction through input trade are complements, there is a fixed-cost-specific efficiency cutoff $\bar{\varphi}(f_i)$ above which firms select into importing. This generates overlap in the size distribution of importers and non-importers. Furthermore, because the optimal sourcing strategy n_i depends on both φ_i and f_i , the model generates variation in import intensity conditional on size. In contrast, a model with a single source of heterogeneity would generate a one-to-one assignment between size and import intensity, as well as size and import status - features that are counterfactual in the French data.

To close the model, we impose equilibrium in the labor market and balanced input trade between the domestic economy and the rest of the world.⁵⁸ We also assume that foreigners demand the output of local firms with the same CES demand structure as domestic consumers and producers ((14)-(15))⁵⁹ and that the supply of foreign inputs from country c is perfectly elastic at price p_c . Letting y_i^{ROW} be the foreign demand for firm i 's production, balanced trade requires that

$$\int_i p_i y_i^{ROW} di = \int_i (1 - s_{Di}) m_i di, \quad (39)$$

where m_i denotes material spending of firm i , so that $(1 - s_{Di}) m_i$ is firm i 's spending on imported varieties, and p_i is firm i 's price.

⁵⁵Given our functional form assumptions, it is necessary to have a fixed cost to start importing f_I to guarantee the existence of non-importers.

⁵⁶In this section, we assume that β_i is constant across firms.

⁵⁷See Section 8.6 of the Appendix.

⁵⁸We abstract from trade in final goods. In this single sector economy with roundabout production, domestic inputs and final goods are equivalent. This is no longer the case with multiple sectors, see the Online Appendix.

⁵⁹This simplifies the problem of local producers, which is given by (37)-(38), where the demand of their different customers (consumers, local firms and foreigners) can be aggregated into a single iso-elastic demand function. The term B in (37) incorporates the sum of spending across the three types of customers.

Definition 1. An equilibrium is a set of prices $w, [p_i]$, labor demands for production and fixed costs $[l_i, l_i^F]$, differentiated product quantities, consumption levels and foreign demands $[y_i, c_i, y_i^{ROW}]$, domestic and international input demands by local firms $[y_{vi}], [z_{ci}]$ and sourcing strategies $[n_i]$ such that:

1. Firms maximize profits given by (37)-(38),
2. Consumers maximize utility given by (14) and (15) subject to their budget constraint

$$\int_i p_i c_i di = wL + \int_i \pi_i di, \quad (40)$$

3. Trade is balanced (39),
4. Labor and good markets clear

$$\begin{aligned} L &= \int_i (l_i + l_i^F) di \\ y_i &= c_i + y_i^{ROW} + \int_\nu y_{vi} dv. \end{aligned}$$

The following proposition characterizes the full gain in consumer welfare relative to autarky.

Proposition 3. Consider the above setup and let W and W^{Aut} denote the total welfare in the trade equilibrium and autarky respectively. Then

$$\frac{W}{W^{Aut}} = \frac{P^{Aut}}{P} \times \left(\frac{L - \int_i l_i^F di}{L} \right). \quad (41)$$

Proof. See Online Appendix. □

Proposition 3 shows that the welfare gains from input trade consists of two components. First, there is the reduction in consumer prices associated with input trade, which is captured by the term $P^{Aut}/P > 1$. This was the focus of Sections 3-4 above. Second, there is the resource loss associated with attaining the equilibrium sourcing strategies, captured by the second term in expression (41). When foreign sourcing is costly, there are fewer workers left for production.⁶⁰

5.2 Quantifying the Welfare Gains

We now have all the ingredients in place to estimate the full welfare effects of input trade along the lines of Proposition 3. A natural requirement for the structural model is to match the consumer price gains from trade which are a component of welfare, as shown by (41). It follows from Proposition 2 that this is ensured by targeting the joint distribution of value added and domestic expenditure

⁶⁰While we calculate $\int_i l_i^F di$ within a model of fixed costs, we note that (41) is consistent with other extensive margin mechanisms. For example, if importers found their trading partners through a process of network formation, (41) would still hold but the environment to calculate $\int_i l_i^F di$ would be different.

shares. Our strategy is as follows. We use the estimates of ε , γ and σ from Section 4.1 above⁶¹ and estimate η directly from the micro-data. Recall that η determines the price index of the import bundle - see (36)- and hence the demand for foreign varieties.⁶² We identify this parameter from the cross-sectional relationship between firms' extensive margin of trade and their domestic shares. Next, we parametrize the distribution of firm efficiency and fixed costs as a joint log-normal distribution

$$\begin{pmatrix} \ln(\varphi) \\ \ln(f) \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \mu_\varphi \\ \mu_f \end{pmatrix}, \begin{pmatrix} \sigma_\varphi^2 & \rho\sigma_\varphi\sigma_f \\ \rho\sigma_\varphi\sigma_f & \sigma_f^2 \end{pmatrix} \right), \quad (42)$$

where ρ controls the correlation between efficiency and fixed costs. We normalize μ_φ and calibrate the rest of the parameters in (42) to match salient features of the joint distribution of value added and domestic expenditure shares. Finally, we choose the fixed cost of being an importer f_I to match the share of importers in the French population.

Estimation of η . As seen from (36), η parametrizes the sensitivity of the import bundle price-index with respect to firms' extensive margin of trade (n). Because import prices map into import spending, we rely on the cross-sectional relation between firms' domestic expenditure shares and the extensive margin to identify η . More specifically, (36) implies that the domestic expenditure share is given by

$$s_D(n) = \left(1 + \left(\frac{p_D}{q_D} \frac{1}{z} n^\eta \right)^{\varepsilon-1} \right)^{-1}. \quad (43)$$

(43) predicts a log-linear relation between n and the term $(1 - s_D)/s_D$, with a slope given by η . At this point, we need to take a stand on what the counterpart of n is in the data. We focus on the number of countries the firm sources its products from, i.e. the number of foreign varieties.⁶³ We run the following regression:

$$\ln \left(\frac{1 - s_{Dist}}{s_{Dist}} \right) = \delta_s + \delta_t + \delta_{NK} + \eta(\varepsilon - 1) \ln(n_{ist}) + u_{ist}, \quad (44)$$

where n_{ist} denotes firm i 's average number of countries per product sourced, δ_{NK} contains a set of fixed effects for the number of products sourced and δ_s and δ_t are sector and year fixed effects. Hence, we identify η from firms sourcing the same number of products from a different number of supplier countries. We measure products at the 8-digit level. Section 8.7 in the Appendix contains the results of estimating (44) with a variety of specifications. Our preferred specification yields a

⁶¹Section 4.1 provides estimates of σ and γ by sector. In this section, we use value-added weighted averages of these sectoral estimates.

⁶²While the import-price index (36) depends also on z , it turns out that this parameter is not required for the calibration of the model. See the Online Appendix.

⁶³This notion of "varieties" is widely used in the literature - see e.g. Broda and Weinstein (2006) and Goldberg et al. (2010). Moreover, 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 may plausibly stem from love-for-variety effects, which are at the heart of the mechanism stressed by our theory. However, we note that the analysis that follows can be done under alternative interpretations of n .

value of η of 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, (ii) the share of importing firms, (iii) the standard deviation of log value added, (iv) the standard deviation of log domestic shares and (v) the correlation between log value added 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 efficiency (σ_φ) from the dispersion in domestic shares and value added, respectively. Finally, the correlation between efficiency and fixed costs (ρ) is disciplined by the correlation between value added and domestic spending.⁶⁴ Note that, by explicitly targeting the economy’s aggregate domestic share, we can again compare our results to those of an aggregate approach where the moments from the micro-data are not used.

Table 7 contains the results of the calibration and shows that the model can be calibrated to match the moments accurately. In particular, note that the correlation between firm efficiency and fixed costs turns out to be positive. This helps to match the far from perfectly negative correlation between value added and domestic shares in the data. We show in Section 8.8 of the Appendix that the calibrated model is also able to match a number of non-targeted moments relatively well.⁶⁵

[Table 7 here]

Welfare Gains from Input Trade. With the calibrated model at hand, we now compute the full welfare gains from input trade. Table 8 contains the results. The first column shows that the model-predicted consumer price gains are very close to those measured in the data.⁶⁶ This should not come as a surprise since such gains are a function of the joint distribution of value added 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 17.54%. Thus, we see that only about half of the consumer price gains translate into welfare gains once the resources spent in fixed costs are taken into account. The reason is that, as seen in column three, a move to autarky would free up about 15% of the labor force, which counteracts the increase in prices.

[Table 8 here]

⁶⁴We describe the algorithm used to calibrate the model in the Online Appendix.

⁶⁵While our calibration targeted moments for the full populations of firms, Section 8.8 of the Appendix reports various moments of the joint distribution of size and import intensity for importers.

⁶⁶The number reported in Table 8 for the French data comes from applying the one-sector consumer price gains formula (22) to the Manufacturing sector pooling all sub-sectors together. Additionally, we leave the non-manufacturing sectors out of the analysis. For these reasons, the price-index gains reported in Table 8 do not coincide with those reported for the Manufacturing sector in Table 5 above.

5.3 The Importance of Domestic Shares

In this subsection, we assess the value of the micro-data on domestic shares for estimating the gains from input trade. Note that the analysis of Section 4 to measure the producer and consumer price gains, as well as the calibration exercise of Section 5.2 to quantify the welfare gains, both relied directly on the domestic share data. When such data is not available, quantifying any type of gains from input trade - i.e. changes in producer unit costs, consumer prices or full consumer welfare - requires calibrating a structural model of importing. The role of the model is to generate a distribution of domestic expenditure shares. A natural question is: how well does a standard model of fixed costs do in predicting such shares?

To answer this question, we consider the following simple exercise. We calibrate the model of Section 5.1 without targeting the moments associated with domestic expenditure shares, i.e. the dispersion of domestic shares and their correlation with value added. Accordingly, we set the dispersion in fixed costs and their correlation with efficiency both to zero, i.e. $\sigma_f = \rho = 0$. We continue to target the dispersion in value added. Note that firm efficiency is the single source of heterogeneity in this model.⁶⁷

We report the results in Table 15 in Section 8.8 of 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 upward biased relative to those of the baseline. In particular, the NSD model over-predicts both the consumer price and the welfare gains from trade relative to the baseline. This is intuitive. By relying on efficiency as the single source of heterogeneity, the NSD model generates a perfectly negative correlation between firm efficiency and the domestic share. This means that firms with higher efficiency experience larger reductions in their unit costs, a feature that tends to make input trade more attractive. This manifests itself as a counterfactually strong negative correlation between value added and domestic shares. The resulting biases in the estimates of the gains from trade can be quantitatively meaningful, of about 14% for the consumer price gains and 24% for welfare.

6 Conclusion

Firms around the world routinely engage in input trade to reduce their costs of production, thereby benefiting domestic consumers. Quantifying the gains from input trade, however, has been limited by an inherent difficulty. As firms differ vastly in the intensity with which they participate in international markets, aggregate trade models are inapplicable. One therefore has to resort to firm-based models of importing to study the normative implications of input trade. In this paper, we identify the aspects of the data that are crucial to credibly doing so.

We first focus on the effect of input trade on consumer prices. We show that, in a class of models of importing, the joint distribution of firm size and import intensity fully determines the change in consumer prices relative a situation of input autarky. Importantly, these data, together with a

⁶⁷This exercise is in the spirit of the ones in Gopinath and Neiman (2014) and Ramanarayanan (2014).

limited set of parameters, is all that is required. To derive this result, we exploit the dichotomy between the intensive and extensive margins of importing that is present in a wide class of models. This allows us to focus on firms' cost minimization problem and apply a sufficiency statistic approach at the firm level. In particular, we can measure changes in firms' unit costs without specifying the details of the import environment or the extensive margin. We then turn to a broader notion of welfare that also takes into account the change in consumers' income due to the change in firms' sourcing strategies. To quantify such change in income, we need to commit to a particular model of the extensive margin and fully specify the import environment. In the context of a standard model with fixed costs, we show that matching the micro data on firm size and import intensity significantly affects the measurement of the welfare gains from input trade.

In an application to French data, we find substantial variation in the benefits from international sourcing across firms. While the median importer would see its costs increase by 11% if the French economy moved to input autarky, these costs would increase by more than 80% for 10% of the firms. Input trade is also important at the aggregate level as consumers would face 27% higher prices for manufacturing products under input autarky.

References

- Amiti, M., O. Itzhoki, and J. Konings: 2014, 'Importers, Exporters, and Exchange Rate Disconnect'. *American Economic Review* **104**(7), 1942–1978.
- 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. J. Redding, and P. K. Schott: 2012, 'The Empirics of Firm Heterogeneity and International Trade'. *Annual Review of Economics* **4**, 283–313.
- 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: 2015, 'Estimates of the Trade and Welfare Effects of NAFTA'. *The Review of Economic Studies* **82**(1), 1–44.
- Costinot, A. and A. Rodríguez-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'. *The Annual Review of Economics* **6**, 201–227.
- De Loecker, J., P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik: 2012, 'Prices, Markups and Trade Reform'. Working Paper.
- Eaton, J. and S. Kortum: 2002, 'Technology, Geography, and Trade'. *Econometrica* **70**(5), 1741–1779.
- 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.
- Feenstra, R. C.: 1994, 'New Product Varieties and the Measurement of International Prices'. *American Economic Review* **84**(1), 157–177.
- 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.
- Imbs, J. and I. Mejean: 2015, 'Elasticity Optimism'. *American Economic Journal: Macroeconomics* **7**(3), 43–83.
- Jia, P.: 2008, 'What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry'. *Econometrica* **76**(6), 1263–1316.
- 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: 2012, 'Measuring aggregate productivity growth using plant-level data'. *The RAND Journal of Economics* **43**(4), 705–725.
- Oberfield, E. and D. Raval: 2014, 'Micro Data and Macro Technology'. Working Paper.
- 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, 'Trade Models, Trade Elasticities, and the Gains from Trade'. Working Paper.
- Wooldridge, J.: 2009, 'On estimating firm-level production functions using proxy variables to control for unobservables'. *Economics Letters* **104**, 112–114.

7 Tables and Figures

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

Notes: σ_s denotes the demand elasticity, which is measured with industry-specific average markups. Markups are constructed as the ratio of firm revenues to total costs, which are computed as the sum of material spending, labor payments and the costs of capital. The costs of capital are measured as Rk where k denotes the firm's capital stock and R is the gross interest rate, which we take to be 0.20. α_s denotes the sectoral share in consumer expenditure, which is taken from the Input-Output Tables according to (27). γ_s denotes the sectoral share of material spending in total costs, which is measured at the firm level and then averaged at the sector level. "VA share" is the sectoral share of value added in manufacturing, computed from the firm-level data. $s_{D_s}^{Agg}$ are the sectoral aggregate domestic shares, computed as $s_{D_s}^{Agg} = \sum_{i=1}^n s_{D_{is}} \times \omega_{is}$, where ω_{is} is the firm share in sectoral value added. See Appendix for the details.

Table 1: Structural Parameters by Industry

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

Notes: Robust standard errors in parentheses with ***, **, and * respectively denoting significance at the 1%, 5% and 10% levels. The table contains the results of estimating (33) with the instrument given in (32). We employ estimates of 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 and Translog technology. For the factor share specification, we use data for the years 2002-2006. For the 2-step GMM procedure, 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. Standard errors in the 2-step GMM procedure are constructed via bootstrap to take the sampling variation in the generated regressor $\gamma_s \Delta \ln(s_D)$ into account. For non-importers, the instrument is set to zero in the full sample specifications.

Table 2: Estimating the Elasticity of Substitution ε

Mean	Quantile				
	10	25	50	70	90
24.87	0.64	2.79	11.18	33.74	85.73

Notes: The table reports quantiles of the empirical distribution of the producer gains from input trade, i.e. $(s_{Di}^{\gamma_s/(1-\varepsilon)} - 1) \times 100$ - see Proposition 1. The data for the domestic expenditure share corresponds to the cross-section of importing firms in 2004. For ε and γ_s , the estimates from the factor shares approach contained in Tables 1 and 2 are used.

Table 3: Moments of the Distribution of Producer Gains in France

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

Notes: Robust standard errors in parentheses with ***, ** and * respectively denoting significance at the 1%, 5% and 10% levels. The table contains the results of estimating (34). All regressions include year fixed effects and 3-digit industry fixed effects. The data corresponds to the full sample of firms between 2002 and 2006. The number of varieties is the number of countries the firm sources from (averaged across products). A firm is a member of an international group if at least one affiliate or the headquarter is located outside of France.

Table 4: Correlates of the Producer Gains

	Manufacturing Sector		Entire Economy	
Consumer Price Gains	27.52	[21.2,35.9]	9.04	[7.1,11.6]
Aggregate Data	30.86	[21.5,45.3]	9.92	[7.1,14]
Bias	3.34	[0.2,10]	0.88	[0,2.6]

Notes: The table reports the reduction in consumer prices for the manufacturing sector $(P_M^{Aut}/P_M - 1) \times 100$ (left panel) and the entire economy $(P^{Aut}/P - 1) \times 100$ (right panel) associated with input autarky. The measure in the first row is based on Proposition 2 where the associated Λ_s are reported in Table 6 and the structural parameters Ξ, γ_s, σ_s and α_s given in Table 1. The second row contains results based on an aggregate model with identical input-output structure and parameters. Specifically, they are based on Proposition 2 where the sectoral gains are measured by Λ_s^{agg} as per (24) instead of Λ_s . The third row reports the bias, defined as the difference between the first two rows - see (25). 90-10 confidence intervals are reported in brackets for all measures. These are calculated via a bootstrap procedure which we describe in the Online Appendix. The empirical distributions of all statistics are estimated using 200 bootstrap iterations.

Table 5: The Consumer Price Gains From Input Trade in France

<i>Estimated</i>				
Parameter	Value	Identified from		
σ	3.83	Revenue/Cost Data, Section 4.1		
ε	2.38	Prod. Function Estimation, Section 4.1		
γ	0.61	Prod. Function Estimation, Section 4.1		
η	0.38	Dom. Share and Ext. Margin, Section 5.2		
<i>Calibrated</i>				
		Moment	Data	Model
μ_f	2.43	Aggregate Dom. Share	0.72	0.72
f_I	0.12	Share of Importer	0.20	0.20
σ_φ	0.51	Dispersion in log Value Added	1.52	1.52
σ_f	2.41	Dispersion in log Dom. Shares	0.36	0.36
ρ	0.73	Correlation log Value Added - log Dom. Shares	-0.31	-0.31

Notes: The table contains the estimated and calibrated structural parameters that are used in Section 5. The first panel contains the parameters estimated directly from the micro data. The estimates for ε and γ correspond to the factor shares approach. The values for γ and σ are value added weighted averages of the corresponding sectoral estimates reported in Table 1. The estimate for η is taken from Table 13 in the Appendix. The parameters $\mu_f, f_I, \sigma_\varphi, \sigma_f$ and ρ are calibrated to match the five moments listed in column three of the second panel. The dispersion and correlation moments are computed in the full sample of firms, which includes non-importers. The calibration algorithm is described in the Online Appendix.

Table 7: Structural Parameters for Welfare Analysis

Industry	ISIC	Direct Price Reductions	Domestic Inputs	Sectoral Price Gains	Aggregate Data
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]

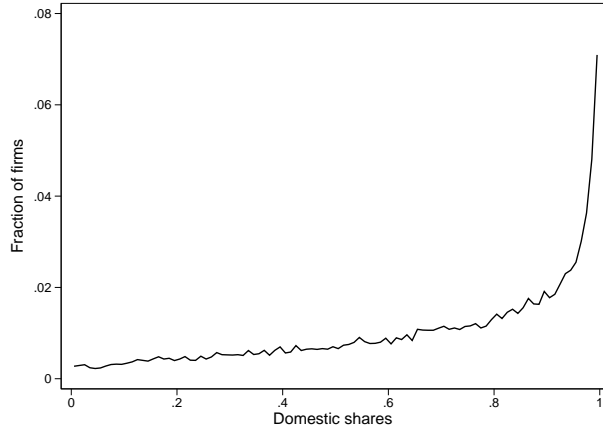
Notes: The first column reports the direct price reductions from international sourcing relative to autarky, $(exp(\Lambda_s) - 1) \times 100$, which are calculated according to (20). The second column reports the reductions in the price of domestically sourced intermediate inputs, $(P_s^{Aut}/P_s - 1) \times 100$. The third column contains the full change in sectoral prices relative to autarky, $(P_s^{Aut}/P_s - 1) \times 100$. Column four reports the direct price reductions predicted by an aggregate approach, $(exp(\Lambda_s^{agg}) - 1) \times 100$, as per (24). 90-10 confidence intervals are reported in brackets for all measures. These are calculated via a bootstrap procedure which we describe in the Online Appendix. The empirical distributions of all statistics are estimated using 200 bootstrap iterations.

Table 6: The Gains from Input Trade: Sectoral Variation

	Consumer Price Gains	Welfare Gains	% of Labor in Fixed Cost Production
Model	38.07	17.53	14.87
French Data	41.53	-	-

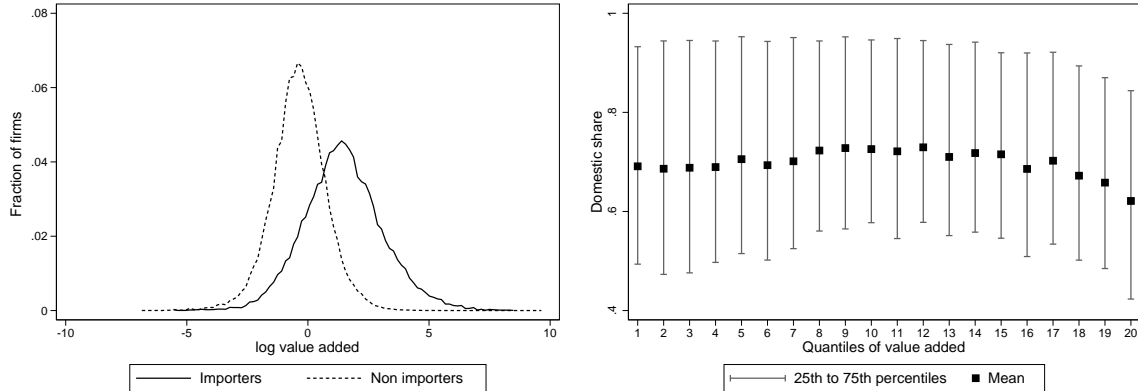
Notes: The first row of the table contains the results of Section 5.2. The consumer price gains are given by $(P^{Aut}/P - 1) \times 100$. The welfare gains are given by $(W/W^{Aut} - 1) \times 100$ and are computed according to (41). The third column reports the percentage of the labor force allocated to fixed costs production, i.e. $\int l_i^F di/L \times 100$. The second row reports the consumer price gains for the French manufacturing sector, which are computed applying (22) to a sample of firms that pools all manufacturing sectors together. The parameters used are reported in Table 7. See main text for details.

Table 8: Welfare Gains from Input Trade



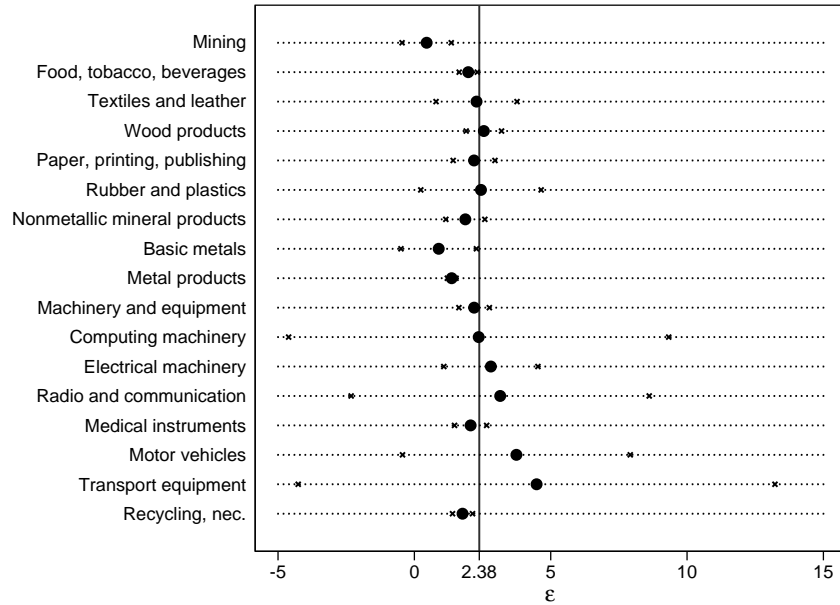
Notes: The figure shows the cross-sectional distribution of domestic expenditure shares, i.e. the share of material spending allocated to domestic inputs, for the population of importing manufacturing firms in France in 2004.

Figure 1: The Dispersion in Domestic Shares



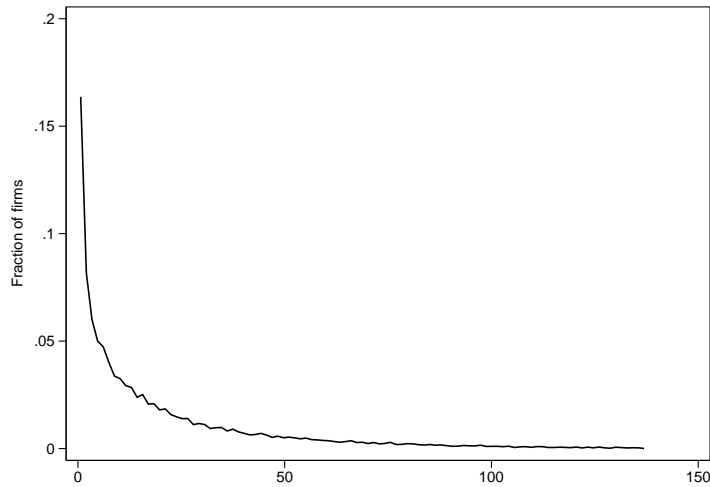
Notes: The left panel displays the distribution of log value added by import status. The right panel shows the mean and the 25th and 75th percentiles of domestic shares for twenty quantiles of value added for importers. The data corresponds to the population of manufacturing firms in France in 2004.

Figure 2: Domestic Shares and Firm Size



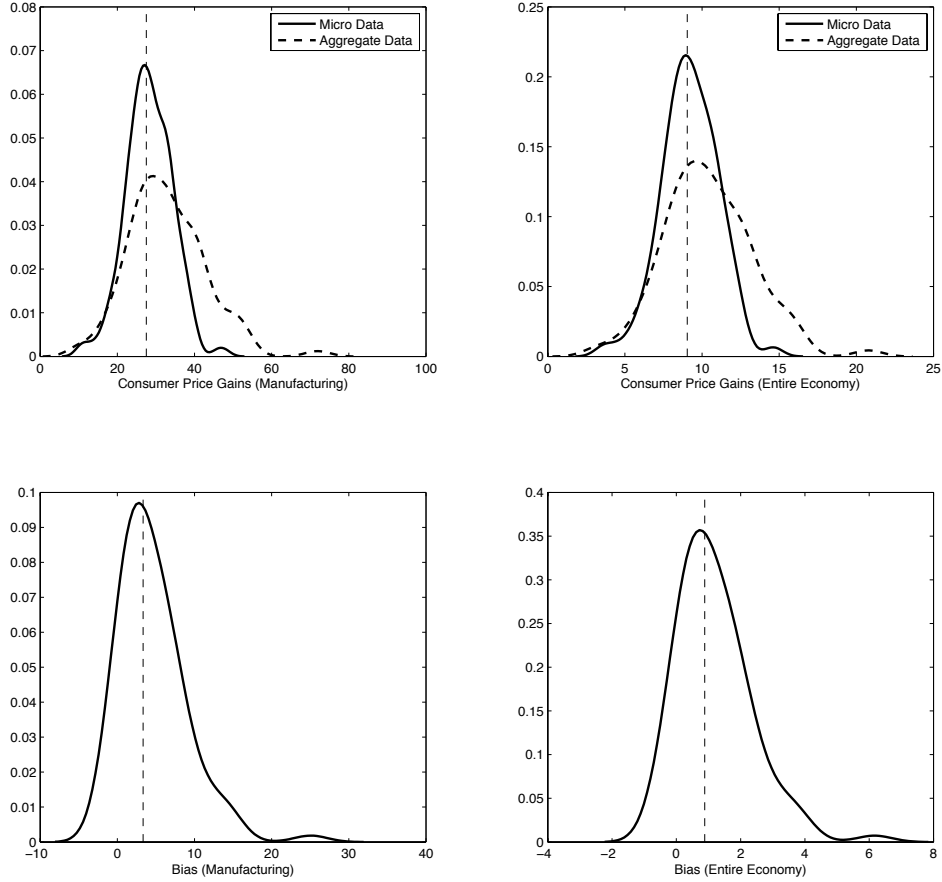
Notes: The figure displays the sector-specific estimates for ε from the integrated GMM approach, which estimates (29)-(30) in a single step following Wooldridge (2009). We display the point estimates and two standard deviations confidence intervals. We also show the benchmark estimate for the factor shares approach $\varepsilon = 2.38$ as a vertical line. We omit the Chemicals sector from the figure because its point estimate is negative and imprecisely estimated. The full results are contained in Table 20 in the Online Appendix.

Figure 3: Estimates of ε_s from Integrated GMM Approach



Notes: The figure reports the empirical distribution of the producer gains from input trade, i.e. $(\varepsilon_s^{\gamma_s / (1-\varepsilon)} - 1) \times 100$ - see Proposition 1. The data for the domestic expenditure shares corresponds to the cross-section of French importing firms in 2004. The values for ε and γ_s are taken from the factor shares approach contained in Table 2.

Figure 4: The Producer Gains from Input Trade in France



Notes: The top panels of the figure depict the bootstrap distribution of the consumer price gains from input trade for the manufacturing sector $(P_M^{Aut}/P_M - 1) \times 100$ (left panel) and the entire economy $(P^{Aut}/P - 1) \times 100$ (right panel). These are computed according to Proposition 2. We display the gains based on the micro data, i.e. using Λ_s , and aggregate data, i.e. using Λ_s^{Agg} . The bottom panels depict the bootstrap distribution of the bias from using aggregate data, which is computed according to (25). The bootstrap procedure is described in the Online Appendix. We use 200 iterations.

Figure 5: Sampling Variation in the Consumer Price Gains and the Bias

8 Appendix

8.1 Generalizations of Proposition 1

In this section, we consider three generalizations of equation (9), which states that the firm's unit costs is given by

$$u_i = \frac{1}{\bar{\varphi}_i} \times (s_{Di})^{\frac{\gamma}{\varepsilon-1}} \times \left(\frac{p_D}{q_D}\right)^\gamma w^{1-\gamma}. \quad (45)$$

(45) was derived under the restrictions: (i) the production function has a constant elasticity of materials γ , (ii) domestic and foreign inputs are combined in a CES fashion with elasticity of substitution ε and (iii) foreign inputs are differentiated at the country, but not at the product level. We now relax these assumptions and derive expressions akin to (45).

Extension 1: CES Upper Tier. Suppose that the production function between materials x and primary factors l is CES instead of Cobb-Douglas, i.e.

$$y = \varphi \left((1-\gamma) l^{\frac{\zeta-1}{\zeta}} + \gamma x^{\frac{\zeta-1}{\zeta}} \right)^{\frac{\zeta}{\zeta-1}}.$$

The rest of the environment is exactly as in Section 3. Let Q denote again the price index of materials x and w denote the price of primary factors l . In this case, the firm's unit cost is given by

$$u = \frac{1}{\varphi} \left(\gamma^\zeta Q^{1-\zeta} + (1-\gamma)^\zeta w^{1-\zeta} \right)^{\frac{1}{1-\zeta}}.$$

Noting that the optimal expenditure share on materials is given by

$$s_M = \frac{\gamma^\zeta Q^{1-\zeta}}{\gamma^\zeta Q^{1-\zeta} + (1-\gamma)^\zeta w^{1-\zeta}}, \quad (46)$$

we can write the firm's unit cost as

$$u = \frac{1}{\varphi} s_M^{\frac{1}{\zeta-1}} \left(\frac{1}{\gamma}\right)^{\frac{\zeta}{\zeta-1}} s_D^{\frac{1}{\varepsilon-1}} \left(\frac{1}{\beta}\right)^{\frac{\varepsilon}{\varepsilon-1}} \left(\frac{p_D}{q_D}\right) \propto s_M^{\frac{1}{\zeta-1}} s_D^{\frac{1}{\varepsilon-1}}, \quad (47)$$

where we have substituted for Q using (8). (47), which is a generalization of (9), shows that measuring the effect of input trade on the unit cost requires knowledge of the counterfactual material share in the autarky equilibrium, s_M^{Aut} .⁶⁸ Because this object is not observed in the data, we can use (6) and

⁶⁸The Cobb-Douglas assumption in (1) in the main text bypasses this issue because it implies that the material share is constant and given by γ . In the non-Cobb-Douglas case, the material share endogenously reacts to changes in the import environment. A move to autarky, for example, makes materials relatively more expensive and should induce firms to substitute towards primary inputs.

(46) to compute it:

$$s_M^{Aut} = \frac{\left(\frac{\gamma}{1-\gamma}\right)^\zeta \beta^{-\frac{\varepsilon}{\varepsilon-1}(1-\zeta)} \left(\frac{p_D/q_D}{w}\right)^{1-\zeta}}{1 + \left(\frac{\gamma}{1-\gamma}\right)^\zeta \beta^{-\frac{\varepsilon}{\varepsilon-1}(1-\zeta)} \left(\frac{p_D/q_D}{w}\right)^{1-\zeta}}. \quad (48)$$

The producer gains from input trade are therefore given by

$$\ln\left(\frac{u^{Aut}}{u}\right)\Big|_{p_D, w} = \ln\left(\frac{1 + \left(\frac{\gamma}{1-\gamma}\right)^\zeta \beta^{-\frac{\varepsilon}{\varepsilon-1}(1-\zeta)} \left(\frac{p_D/q_D}{w}\right)^{1-\zeta} s_D^{\frac{1-\zeta}{\varepsilon-1}}}{1 + \left(\frac{\gamma}{1-\gamma}\right)^\zeta \beta^{-\frac{\varepsilon}{\varepsilon-1}(1-\zeta)} \left(\frac{p_D/q_D}{w}\right)^{1-\zeta}}\right)^{\frac{1}{\zeta-1}}. \quad (49)$$

(49) is the generalization of Proposition 1 for the case where the aggregator between materials and primary factors is CES. We see that, in this case, quantifying the change in the unit cost relative to autarky requires knowledge of additional parameters $[\beta, \zeta, p_D/q_D]$ to predict the material share in autarky. Under the *additional* assumption that there is no variation in β and p_D/q_D across firms, we can bypass the estimation of some of these additional parameters. In this case, all firms would feature the same material share in autarky, which is given by the material share of a domestic firm in the observed trade equilibrium, s_M^D . In this case, (49) reduces to

$$\ln\left(\frac{u^{Aut}}{u}\right)\Big|_{p_D, w} = \ln\left(1 - s_M^D + s_D^{\frac{1-\zeta}{\varepsilon-1}} \times s_M^D\right)^{\frac{1}{\zeta-1}}, \quad (50)$$

so that only micro-data on domestic expenditure shares s_D and the two elasticities of substitution ε and ζ are required.⁶⁹

Extension 2: General Production Function for Materials. In Section (3), we assumed that material services were a CES aggregator of a domestic variety z_D and a foreign input bundle x_I . Suppose now that the aggregator for materials is given by a general function

$$x = g(q_D z_D, x_I). \quad (51)$$

We continue to assume that materials x and primary factors l are combined with a Cobb-Douglas production function given in (1). Again let $A(\Sigma)$ be the price index of the import bundle and $Q(\Sigma)$ be the price index of materials. Consider any shock to the trading environment that affects $A(\Sigma)$. Then

$$d\ln(u)|_{p_D, w} = \gamma \times d\ln(Q)|_{p_D} = \gamma \frac{z_I A}{u} \frac{dA}{A} = \gamma(1 - s_D) d\ln(A). \quad (52)$$

⁶⁹Note that, when $\zeta \rightarrow 1$, (50) reduces to the expression in Proposition 1:

$$\lim_{\zeta \rightarrow 1} \ln\left(\frac{u^{Aut}}{u}\right)\Big|_{p_D, w} = \frac{\gamma}{1-\varepsilon} \ln(s_D)$$

The optimality conditions from the cost-minimization problem imply that

$$d\ln(A) = -\frac{\left(-\frac{1}{\varepsilon_L}\right)}{1 - \frac{1}{\varepsilon_L}} \frac{1}{1 - s_D} d\ln(s_D),$$

where

$$-\frac{1}{\varepsilon_L} \equiv \frac{\partial \ln\left(\frac{\partial g(q_D z_D, x_I)/\partial x_D}{\partial g(q_D z_D, x_I)/\partial x_I}\right)}{\partial \ln\left(\frac{q_D z_D}{x_I}\right)}$$

is the local elasticity of substitution. Substituting this into (52) yields

$$d\ln(u)|_{p_D, w} = \gamma \frac{\frac{1}{\varepsilon_L}}{1 - \frac{1}{\varepsilon_L}} d\ln(s_D) = -\frac{\gamma}{1 - \varepsilon_L} d\ln(s_D). \quad (53)$$

In case the elasticity of substitution is constant, i.e. $\varepsilon_L = \varepsilon$, (53) can be integrated to yield (9).

Extension 3: Multiple Foreign Products. In the main analysis, we assumed that firms source a single product from each sourcing country. In the data, firms often import multiple products from a given country. We now explore how (45) would change in a multi-product environment. Consider first the case where the product aggregator is nested in the country aggregator, i.e. the production structure is given by (1)-(3), where

$$q_{ci} z_c \equiv \psi_{ci} \left([q_{kci} z_{kc}]_{k \in K_{ci}} \right), \quad (54)$$

k is a product index, K_{ci} denotes the set of products that firm i sources from country c , ψ_{ci} is a constant-returns-to-scale production function and (54) applies also to the domestic variety. As long as the number of products sourced domestically does not change when firms are forced into input-autarky, the analysis in the main text remains entirely unchanged and the producer gains are still given by Proposition 1.

Consider next the case where the country aggregator is nested in the product aggregator. Suppose for example that the production structure for intermediates x is given by

$$x = \left(\sum_{k=1}^K (\eta_k x_k)^{\frac{\iota-1}{\iota}} \right)^{\frac{\iota}{\iota-1}} \quad (55)$$

$$x_k = \left(\beta_{ki} (q_{kD} z_{kD})^{\frac{\varepsilon_k-1}{\varepsilon_k}} + (1 - \beta_{ki}) x_{kI}^{\frac{\varepsilon_k-1}{\varepsilon_k}} \right)^{\frac{\varepsilon_k}{\varepsilon_k-1}} \quad (56)$$

$$x_{kI} = h_{ki} \left([q_{kci} z_{kc}]_{c \in \Sigma_{ki}} \right). \quad (57)$$

Note that the sourcing strategy is now a list of countries for each product. Letting Q_i and Q_{ki} denote the price indices for materials x and product-specific material services x_k respectively, it can be easily

shown that

$$Q_i = \left(\sum_{k=1}^K (Q_{ki}/\eta_k)^{1-\iota} \right)^{\frac{1}{1-\iota}}$$

$$Q_{ki} = (s_{kDi})^{\frac{1}{\varepsilon_k-1}} \beta_{ki}^{-\frac{\varepsilon_k}{\varepsilon_k-1}} p_{kD}/q_{kD},$$

where s_{kDi} is firm i 's domestic expenditure share for product k . The producer gains are therefore given by

$$G_i = \frac{\gamma}{\iota-1} \times \ln \left(\sum_{k=1}^K \chi_{ki} (s_{kDi})^{\frac{\iota-1}{1-\varepsilon_k}} \right), \quad (58)$$

where

$$\chi_{ki} \equiv \frac{\left(\beta_{ki}^{-\frac{\varepsilon_k}{\varepsilon_k-1}} p_{kD}/q_{kD} \right)^{1-\iota}}{\sum_{k=1}^K \left(\beta_{ki}^{-\frac{\varepsilon_k}{\varepsilon_k-1}} p_{kD}/q_{kD} \right)^{1-\iota}}.$$

We see that the producer gains are akin to a weighted average of the product-specific producer gains $(s_{kDi})^{\frac{\iota-1}{1-\varepsilon_k}}$. In our empirical application, we cannot implement (58) because we do not observe domestic shares at the product level s_{kDi} in the French data. Note that implementing (58) also requires measuring the weights χ_{ki} . In the case where (55) takes the Cobb-Douglas form, i.e. $\iota = 1$ as in Halpern et al. (2011), (58) simplifies to

$$G_i = \sum_{k=1}^K \eta_k \frac{\gamma}{1-\varepsilon_k} \ln \left(s_{Di}^k \right).$$

Thus, in the Cobb-Douglas case, the producer gains are a weighted average of the product-specific producer gains.

8.2 Proof of Proposition 2

The consumer price index associated with (14)-(15) is given by

$$P = \prod_{s=1}^S (P_s/\alpha_s)^{\alpha_s} \quad \text{where } P_s = \left(\int_0^{N_s} p_{is}^{1-\sigma_s} di \right)^{\frac{1}{1-\sigma_s}}, \quad (59)$$

and P_s is the price index for sector s . Using (59), the consumer price gains from input trade can be expressed as

$$G \equiv \ln \left(\frac{P^{Aut}}{P} \right) = \sum_{s=1}^S \alpha_s \ln \left(\frac{P_s^{Aut}}{P_s} \right).$$

We now express P_s^{Aut}/P_s in terms of observables. Note that monopolistic competition implies a constant markup pricing rule, $p_{is} = \mu_s u_{is}$. Using the expression for the firm's unit cost in terms of

its domestic expenditure share in (9), we find that

$$P_s = \mu_s \left(\frac{p_{Ds}}{q_{Ds}} \right)^{\gamma_s} \left(\int_0^{N_s} \left(\tilde{\varphi}_i^{-1} (s_{Di})^{\gamma_s / (\varepsilon_s - 1)} \right)^{1 - \sigma_s} di \right)^{\frac{1}{1 - \sigma_s}}, \quad (60)$$

which is (17) in the main text. Given the aggregator in (16), the price index of the domestic bundle is given by

$$p_{Ds} = \zeta_s^* \prod_{j=1}^S P_j^{\zeta_j^s} \text{ where } \zeta_s^* \equiv \prod_{j=1}^S (\zeta_j^s)^{-\zeta_j^s} \quad (61)$$

Note that (60) implies

$$\frac{P_s^{Aut}}{P_s} = \left(\frac{p_{Ds}^{Aut}}{p_{Ds}} \right)^{\gamma_s} \left(\frac{\int_0^{N_s} \tilde{\varphi}_i^{\sigma_s - 1} di}{\int_0^{N_s} \left(\tilde{\varphi}_i s_{Di}^{\frac{\gamma_s}{1 - \varepsilon_s}} \right)^{\sigma_s - 1} di} \right)^{\frac{1}{1 - \sigma_s}} = \left(\frac{p_{Ds}^{Aut}}{p_{Ds}} \right)^{\gamma_s} \left(\int_0^{N_s} \omega_i s_{Di}^{\frac{\gamma_s(1 - \sigma_s)}{1 - \varepsilon_s}} di \right)^{\frac{1}{1 - \sigma_s}}, \quad (62)$$

where ω_i is firm i 's share in total value added in sector s and the second equality follows from $va_i = \kappa_s \tilde{\varphi}_i^{\sigma_s - 1} s_{Di}^{\frac{\gamma_s(\sigma_s - 1)}{1 - \varepsilon_s}}$. With (62) at hand, we can express the consumer price gains as

$$G = \sum_{s=1}^S \gamma_s \alpha_s \pi_s + \sum_{s=1}^S \alpha_s \Lambda_s \quad (63)$$

where $\pi_s \equiv \ln \left(\frac{p_{Ds}^{Aut}}{p_{Ds}} \right)$ and Λ_s is given by (20) in the main text. As Λ_s are observable from the micro-data, obtaining G reduces to solving for $[\pi_s]_{s=1}^S$. Note that (61) and (62) jointly imply

$$\pi_s = \sum_{j=1}^S \zeta_j^s \gamma_j \pi_j + \sum_{j=1}^S \zeta_j^s \Lambda_j. \quad (64)$$

(64) gives an $S \times S$ system of equations that characterizes the equilibrium $[\pi_s]_{s=1}^S$. Letting $\pi \equiv [\pi_1, \pi_2, \dots, \pi_s]$ be a column vector, we can express the system in (64) in matrix form as $\pi = \Xi \Gamma \pi + \Xi \Lambda$. Its solution is given by $\pi = (\mathcal{I} - \Xi \Gamma)^{-1} \Xi \Lambda$. Using (63), the consumer price gains G are therefore given by

$$G = \alpha' \Gamma \pi + \alpha' \Lambda = \alpha' \Gamma (\mathcal{I} - \Xi \Gamma)^{-1} \Xi \Lambda + \alpha' \Lambda. \quad (65)$$

This proves Proposition 2. For counterfactuals other than autarky, (62) should be replaced by

$$\frac{P'_s}{P_s} = \left(\frac{p'_{Ds}}{p_{Ds}} \right)^{\gamma_s} \left(\frac{\int_0^{N_s} \left(\tilde{\varphi}_i (s'_{Di})^{\gamma_s/(1-\varepsilon_s)} \right)^{\sigma_s-1} di}{\int_0^{N_s} \left(\tilde{\varphi}_i (s_{Di})^{\gamma_s/(1-\varepsilon_s)} \right)^{\sigma_s-1} di} \right)^{\frac{1}{1-\sigma_s}} = \left(\frac{p'_{Ds}}{p_{Ds}} \right)^{\gamma_s} \left(\int_0^{N_s} \omega_i \left(\frac{s_{Di}}{s'_{Di}} \right)^{\frac{\gamma_s(1-\sigma_s)}{1-\varepsilon_s}} di \right)^{\frac{1}{1-\sigma_s}}, \quad (66)$$

where s'_{Di} denotes the counterfactual domestic share and P'_s, p'_{Ds} denote the counterfactual price indices. It follows that the consumer price gains associated with the policy, $G \equiv \ln(P'/P)$, are given by (65) where Λ_s is given by (23).

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.⁷⁰ The data is collected at the 8-digit (NC8) level. A firm located within the French metropolitan territory must report 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 exceeds 1,000 Euros. For within EU imports, import flows have to be reported as long as the firm's annual trade value exceeds 100,000 Euros.⁷¹ However, some firms that are below the threshold (ca. 10,000 firm-year observations out of ca. 130,000) voluntarily report.⁷²

In spite of this limitation, the attractive feature of the French data is the presence of unique firm identifiers (the SIREN code) that is available in all French administrative files. Hence, various 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.⁷³ 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 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) that describes the structure of ownership of all firms in the private sector whose financial investments in other firms (participation) are higher than 1.2 million Euros or have sales above 60 million Euros or have more than 500

⁷⁰For 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.

⁷¹This threshold was in effect between 2001 and 2006, which is period we focus on. 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.

⁷²The existence of this administrative threshold induces a censoring of small EU importers. In results available upon request, we use the time-variation in the reporting thresholds (see footnote 71) to show that this concern is unlikely to severely affect our results. The reason is related to weak relation between domestic expenditure shares and firm size shown in Figure 2.

⁷³The 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 purposes these two data sets contain essentially the same information. Their union covers nearly the entire universe of 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

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 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. As expected, importers are larger, more capital intensive and have higher revenue productivity - see also Bernard et al. (2012). Furthermore, import and export status are highly correlated.

8.4 Estimates of the Elasticity of Substitution ε : Robustness

In Section 4.1, we constructed the instrument for the estimation of ε using firms' first year of import activity as the initial period for the pre-sample weights, s_{cki}^{pre} - see (32). We now redo the analysis of estimating (33) keeping the year used for the pre-sample weights fixed at 2001 for all firms. Note that this reduces the size of the sample as all firms that start to import after 2001 are dropped. Table 10 contains the results for the factor shares and the 2-step GMM procedures with Cobb-Douglas and translog technology. The results of the main text are also reported for comparison - see the "All weights" rows. As before, we run the regressions on both the entire population and the sample of importers. Depending on the specification, the estimates of ε range from 1.4 to 2.4.

			$\tilde{\gamma}_s \times \Delta \ln(s_D)$	ε	N
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	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	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)	

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 Translog. 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 $\gamma_s \Delta \ln(s_D)$ into account. See the main text for details regarding the different subsamples for the respective specifications.

Table 10: Estimating the Elasticity of Substitution ε

	Micro Data	Aggregate Data				
		ε				
	2.38	2.38	3	4	5	6
Entire Economy	9.04	9.9	6.72	4.43	3.31	2.64
Manufacturing Sector	27.52	30.8	20.32	13.12	9.69	7.68

Notes: The table reports the reduction in consumer prices for the entire economy $(P^{Aut}/P-1)\times 100$ (first row) and the manufacturing sector $(P_M^{Aut}/P_M-1)\times 100$ (second row) for different values of the elasticity of substitution ε . In the first two columns, we report the baseline results under $\varepsilon = 2.38$ for comparison. Column one is based on Proposition 2 where Λ_s are computed with micro data as reported in Table 6. The remaining columns contain results based on an aggregate model, i.e. they are based on Proposition 2 where the sectoral gains are measured by Λ_s^{Agg} as per (24) instead of Λ_s . The values for Ξ, γ_s, σ_s and α_s employed for all calculations are given in Table 1.

Table 11: The Consumer Price Gains for Different Values of ε

8.5 The Elasticity Bias: Using ε from Aggregate Trade Flows

Tables 11 and 12 report the consumer price gains from input trade for the entire economy and at the sector level for different values of the elasticity of substitution ε . Columns one and two replicate the results for our baseline estimate $\varepsilon = 2.38$. While column one reports the results based on the micro-data, column two reports the gains based on aggregate data, Λ_s^{Agg} . These results correspond to the ones reported in Tables 5 and 6 above. In the remaining columns, we report the gains for a range of values of ε from studies that rely on aggregate data to estimate such elasticity. For example, Costinot and Rodríguez-Clare (2014) take $\varepsilon = 4$ as their baseline. The tables show that the gains are very sensitive to the value of ε . Table 11 shows that the economy-wide gains predicted by an aggregate approach under $\varepsilon = 4$ are about 50% lower than the gains predicted by the approach that relies on micro-data.

8.6 Characterizing the Model with Fixed Costs

Consider the setup of Section 5.1 where Assumption 1 holds. Note that, under item 3 of this assumption, countries can be indexed by their quality q . We first show that the price index of the import bundle takes the power form in (36). Under items 1 and 3 of Assumption 1, this price index is given by

$$A(\Sigma) = \left(\int_{q \in \Sigma} (p(q)/q)^{1-\kappa} dG(q) \right)^{\frac{1}{1-\kappa}} = \left(\int_{q \in \Sigma} q^{(1-\nu)(\kappa-1)} dG(q) \right)^{\frac{1}{1-\kappa}}. \quad (67)$$

As quality is Pareto distributed, (67) becomes

$$A(\Sigma)^{1-\kappa} = \theta q_{min}^{\theta} \int_{q \in \Sigma} q^{(1-\nu)(\kappa-1)} q^{-\theta-1} dq.$$

Because fixed costs are constant across countries and quality-adjusted prices $q^{\nu-1}$ are decreasing in quality, the sourcing set Σ can be parametrized by a quality cutoff \bar{q} . In particular, the firm selects

Industry	ISIC	Micro Data	Aggregate Data				
		ε					
		2.38	2.38	3	4	5	6
Mining	10-14	2.96	2.5	1.72	1.14	0.86	0.68
Food, tobacco, beverages	15-16	11.06	12.62	8.53	5.61	4.18	3.33
Textiles and leather	17-19	31.14	31.87	20.99	13.55	10	7.92
Wood and wood products	20	8.23	9.58	6.51	4.29	3.2	2.55
Paper, printing, publishing	21-22	12.15	10.96	7.43	4.89	3.65	2.91
Chemicals	24	27.23	28.14	18.62	12.06	8.91	7.07
Rubber and plastics products	25	20.12	21.53	14.37	9.37	6.95	5.52
Non-metallic mineral products	26	13.42	13.29	8.98	5.9	4.39	3.5
Basic metals	27	21.8	28.83	19.07	12.34	9.12	7.23
Metal products (ex machinery and equipment)	28	8.17	7.7	5.24	3.47	2.59	2.07
Machinery and equipment	29	17.64	18.23	12.23	7.99	5.94	4.72
Office and computing machinery	30	20.42	37	24.22	15.56	11.45	9.06
Electrical machinery	31	19.77	21.64	14.45	9.41	6.98	5.55
Radio and communication	32	21.55	22.15	14.78	9.62	7.13	5.67
Medical and optical instruments	33	17.9	15.9	10.7	7.01	5.21	4.15
Motor vehicles, trailers	34	6.24	11.23	7.61	5.01	3.73	2.98
Transport equipment	35	15.32	11.83	8.01	5.27	3.93	3.13
Manufacturing, recycling	36-37	12.87	14.06	9.48	6.23	4.63	3.69
Non-manufacturing		0	0	0	0	0	0

Notes: The table reports the reduction in consumer prices at the sectoral level $(P_s^{Aut}/P_s - 1) \times 100$ for different values of the elasticity of substitution ε . In the first two columns, we report the baseline results under $\varepsilon = 2.38$ for comparison. Column one is based on Proposition 2 where Λ_s are computed with micro data as reported in Table 6. The remaining columns contain results based on an aggregate model, i.e. they are based on Proposition 2 where the sectoral gains are measured by Λ_s^{Agg} as per (24) instead of Λ_s . The values for Ξ, γ_s, σ_s and α_s employed for all calculations are given in Table 1.

Table 12: The Sectoral Consumer Price Gains for Different Values of ε

countries with high enough quality, i.e. $q \in \Sigma$ as long as $q \geq \bar{q}$. It follows that

$$A(\bar{q})^{1-\kappa} = q_{min}^{\theta} \frac{\theta}{\theta - (1-\nu)(\kappa-1)} \bar{q}^{(1-\nu)(\kappa-1)-\theta}. \quad (68)$$

We can rewrite this expression in terms of the mass of countries sourced, n , which is given by

$$n = P(q \in \Sigma) = P(q \geq \bar{q}) = q_{min}^{\theta} \bar{q}^{-\theta}. \quad (69)$$

Substituting (69) into (68) yields

$$A(n) = q_{min}^{-(1-\nu)} \left(\frac{\theta}{\theta - (1-\nu)(\kappa-1)} \right)^{\frac{1}{1-\kappa}} n^{-\left(\frac{1}{\kappa-1} - \frac{1-\nu}{\theta}\right)}, \quad (70)$$

which is (36) in the main text where

$$z \equiv q_{min}^{\nu-1} \left(\frac{\theta}{\theta - (1-\nu)(\kappa-1)} \right)^{\frac{1}{1-\kappa}} \quad (71)$$

$$\eta \equiv \frac{1}{\kappa-1} - \frac{1-\nu}{\theta}. \quad (72)$$

Note that $\eta > 0$ follows from item 4 of Assumption 1. This completes the characterization of (36).

The following proposition characterizes the solution to the extensive margin problem.

Proposition 4. *Consider the setup above and suppose that*

$$\eta(\varepsilon-1) < 1 \text{ and } \eta\gamma(\sigma-1) < 1. \quad (73)$$

Then, the firm's profit maximization problem (37) has a unique solution for any value of φ and f . The optimal mass of countries sourced is given by a function $n(\varphi, f)$ and an efficiency cutoff $\bar{\varphi}(f)$ such that $n(\varphi, f) = 0$ for $\varphi \leq \bar{\varphi}(f)$ with $\bar{\varphi}(\cdot)$ increasing. Furthermore, $n(\varphi, f)$ is increasing in efficiency φ and decreasing in the fixed costs of sourcing f .

Proof. Conditional on importing, the optimal mass of countries is characterized by the following first order condition:

$$\left((q_D/p_D)^{\varepsilon-1} + \left(\frac{1}{z}\right)^{\varepsilon-1} n^{\eta(\varepsilon-1)} \right)^{\frac{\gamma(\sigma-1)}{\varepsilon-1}-1} z^{1-\varepsilon} n^{\eta(\varepsilon-1)-1} = \frac{1}{\eta\gamma(\sigma-1)} \frac{1}{\Gamma} \frac{f}{\varphi^{\sigma-1}}, \quad (74)$$

where $\Gamma \equiv Bw^{(1-\sigma)(1-\gamma)-1}$. The second order condition is given by

$$\Gamma\gamma(\sigma-1)\eta\varphi^{\sigma-1}z^{1-\varepsilon} \left((q_D/p_D)^{\varepsilon-1} + z^{1-\varepsilon}n^{\eta(\varepsilon-1)} \right)^{\frac{\gamma(\sigma-1)}{\varepsilon-1}-1} n^{\eta(\varepsilon-1)-2} \times \quad (75)$$

$$\{(\eta(\varepsilon-1)-1) + (\gamma(\sigma-1) - \varepsilon + 1)\eta l(n)\} < 0$$

where

$$l(n) \equiv \frac{z^{1-\varepsilon} n^{\eta(\varepsilon-1)}}{(q_D/p_D)^{\varepsilon-1} + z^{1-\varepsilon} n^{\eta(\varepsilon-1)}} \in [0, 1].$$

It follows that (75) is satisfied if and only if

$$\eta(\varepsilon - 1) - 1 + (\gamma(\sigma - 1) - \varepsilon + 1)\eta l(n) < 0. \quad (76)$$

Because we allow for arbitrary values of φ and f , we need to verify that (76) holds for any value of n . Sufficient conditions for this are given by

$$\eta(\varepsilon - 1) < 1 \quad (77)$$

and

$$\eta\gamma(\sigma - 1) < 1. \quad (78)$$

If (77) is not satisfied, there exists a range of values of n close enough to zero such that (76) is violated.⁷⁴ (77) is therefore necessary. If $\gamma(\sigma - 1) - \varepsilon + 1 \leq 0$, then (76) is satisfied for all n . If $\gamma(\sigma - 1) - \varepsilon + 1 > 0$, then (76) holds for all n if and only if

$$\eta(\varepsilon - 1) - 1 + (\gamma(\sigma - 1) - \varepsilon + 1)\eta l(1) < 0. \quad (79)$$

As $l(1) < 1$, a sufficient condition for (79) is given by (78). This proves that, under (77)-(78), the optimal mass of countries conditional on importing is uniquely characterized by (74) for any values of φ and f .^{75,76} The firm becomes an importer whenever $\pi_I \geq \pi_D$, where π_I are optimal profits conditional on importing and π_D are profits as a non-importer. It can be shown that this condition is satisfied whenever

$$\left[\left(1 + \left(\frac{p_D}{q_D} z^{-1} n^\eta \right)^{\varepsilon-1} \right)^{\frac{\gamma(\sigma-1)}{\varepsilon-1}} - 1 \right] (q_D/p_D)^{\gamma(\sigma-1)} \Gamma \varphi^{\sigma-1} - fn - f_I > 0, \quad (80)$$

where n is the solution to (74). It follows the firm's profit maximization problem in (37) has a unique solution for any value of φ and f .

Note that, under (77)-(78), the left hand side of (74) is decreasing in n . Therefore, the optimal mass of countries sourced is weakly increasing in φ and weakly decreasing in f . Holding n fixed, an increase in φ tends to increase the left hand side of (80). Additionally, π_I is increasing in φ . It follows that $\pi_I - \pi_D$ is increasing in φ for a given f . This proves that $n = 0$ if and only if $\varphi \leq \bar{\varphi}(f)$ where $\bar{\varphi}(\cdot)$ is implicitly defined as the value of φ that makes the left hand side of (80) equal to zero. The fact that $\bar{\varphi}(f)$ is increasing in f follows from the fact that $\pi_I - \pi_D$ is decreasing in f for a given

⁷⁴This follows from the fact that $l(n)$ is continuous and strictly increasing.

⁷⁵Note that our calibrated and estimated parameters satisfy (77)-(78) - see Table 7.

⁷⁶When the solution to (74) exceeds unity, the solution is given by $n = 1$. Clearly, $n = 0$ cannot be a solution as the firm always prefers to be a non-importer and avoid payment of f_I .

φ .⁷⁷ This proves Proposition 4. □

8.7 Estimation of η

Table 13 contains the results of estimating (44). Columns one and two show that it is important to control for the number of products sourced as import-intensive firms source both more varieties per-product and more products on international markets - without the product fixed effects, the estimated η increases substantially.⁷⁸ Columns three and four show that the estimate of η is virtually unaffected by additional firm-level controls that can affect firms' import behavior conditional on the number of varieties sourced. In column five, we focus on a subsample of firm-product pairs that source their respective products from at least two supplier countries. In this case, the estimated η decreases as the single-variety importers have very high domestic shares in the data. For our quantitative analysis, we take column five as the benchmark.⁷⁹ The implied value of η is 0.382 and it is precisely estimated.

Dep. Variable: $\ln\left(\frac{1-s_D}{s_D}\right)$						
	All Importers				> 1 variety	> 2 varieties
ln (Number of Varieties)	1.308*** (0.009)	0.707*** (0.010)	0.733*** (0.010)	0.739*** (0.010)	0.526*** (0.011)	0.463*** (0.019)
ln (Capital / Employment)				-0.070*** (0.006)		
Exporter Dummy			-0.395*** (0.013)	-0.388*** (0.013)	-0.254*** (0.017)	-0.198*** (0.029)
International Group			0.150*** (0.016)	0.174*** (0.016)	0.216*** (0.016)	0.223*** (0.019)
Control for Num of products	No	Yes	Yes	Yes	Yes	Yes
Implied Eta	0.950*** (0.260)	0.513*** (0.142)	0.532*** (0.147)	0.536*** (0.148)	0.382*** (0.106)	0.336*** (0.096)
Observations	120,344	120,344	120,344	120,344	73,651	35,751

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 average number of countries a firm sources its foreign products from. To back out the value for η , we use our benchmark value for $\varepsilon = 2.378$ from Section 4.1.

Table 13: Estimating η

8.8 Non-Targeted Moments and Model Comparison

Table 14 and Figures 6 and 7 report how the model fits additional non-targeted moments. Table

⁷⁷To see why this is the case, note that the left hand side of (80) is decreasing in f given φ and n . Additionally, π_I is decreasing in f .

⁷⁸Recall that η is a combination of different structural parameters of the economy. While η is sufficient to characterize the welfare gains from trade, one might be interested in decomposing the returns to international sourcing into the the elasticity of substitution across varieties κ , the dispersion in input quality θ , and the elasticity of input prices with respect to quality ν . To do so, we need two additional pieces of information: import prices (to identify ν) and data on firms' expenditure shares across trading partners (to identify θ).

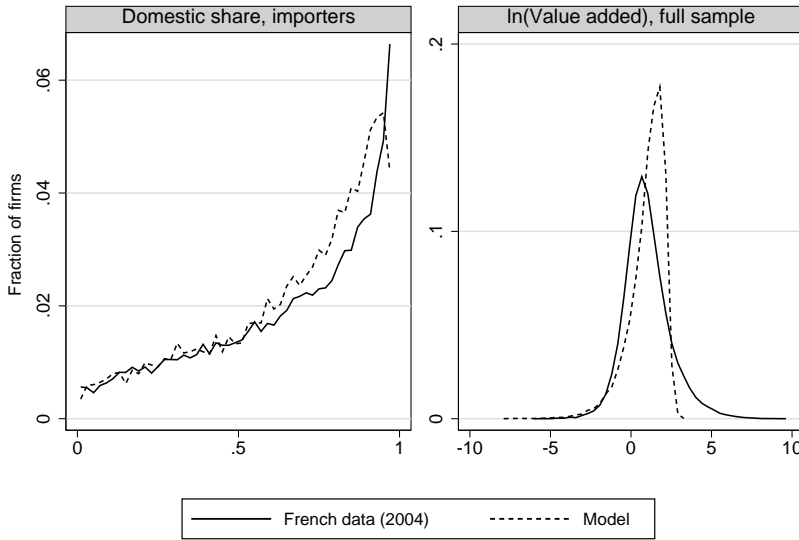
⁷⁹We are concerned that the single-variety observations may not help identify the extensive margin channel emphasized by our theory but rather pick-up other variation across firms. Additionally, a non-parametric regression shows that the log linear relation between n and $(1 - s_D)/s_D$ in (44) fits the data better in the sample with at least two varieties than in the full sample (results available upon request).

<i>Non-Targeted Moments</i>	French Data	Baseline
Avg Domestic Share (Importers)	0.70	0.67
Avg Domestic Share (Population)	0.94	0.93
Agg Domestic Share (Importers)	0.63	0.52
Dispersion log Value Added (Importers)	1.62	1.02
Dispersion log Dom Shares (Importers)	0.69	0.65
Correlation log Value Added - log Dom Shares (Importers)	-0.01	-0.06
Share of Value Added by Importers	0.79	0.59

Notes: This table reports some non-targeted moments in the micro-data (column 1) and in our baseline calibration (column 2). The calibrated parameters of the benchmark model are contained in Table 7.

Table 14: Non-Targeted Moments

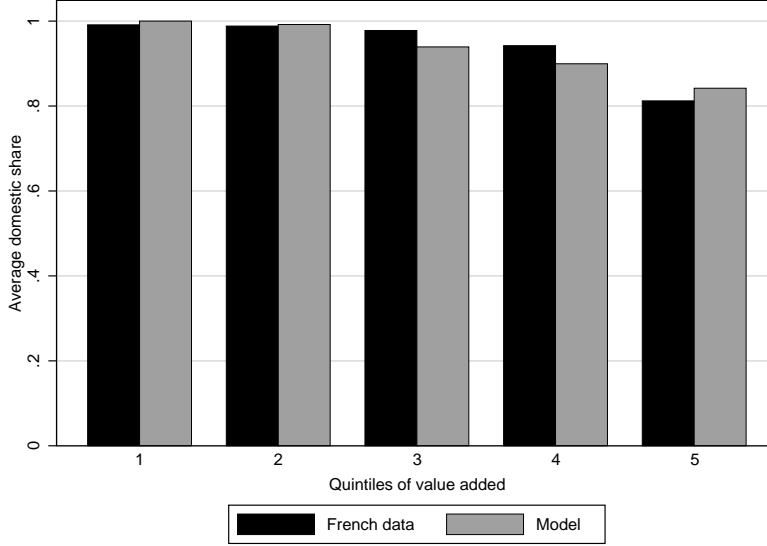
14 shows that the model performs relatively well. That fact that it under predicts the dispersion of value added for importers is related to the fact that the log-normal distribution of efficiency has too thin tails. This also explains why the model under predicts the share of value added by importers - there are too few very large firms in our model, the majority of which are importers.



Notes: The left panel (right panel) shows the distribution of domestic shares (log value added) in the data (solid line) and in the model (dashed line). The firms are grouped in 50 bins of equal length. The distributions of log value added have been normalized to have means of unity.

Figure 6: Marginal Distributions: Model and Data

Figure 6 reports the marginal distributions of domestic shares and log value added 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 may not be worth it to pay the fixed costs to then import tiny amounts. When we compare the distribution of value added between the model and the data, we again see that the model generates too little dispersion, especially on the right tail.



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 7: Correlation Structure: Model and Data

Figure 7 reports the average domestic share by value added quintile for the full population. The model fits this moment relatively well.

Finally, Table 15 reports the calibrated parameters and targeted moments for the alternative calibration of Section 5.3 where no information on domestic shares is used (last two columns).

	Baseline		No s_D Data	
	Model	Parameter	Model	Parameter
Aggregate Domestic Share	0.72	$\mu_f = 2.43$	0.72	$\mu_f = 2.88$
Share of Importers	0.20	$f_I = 0.12$	0.20	$f_I = 0.16$
Dispersion in log Value Added	1.52	$\sigma_\varphi = 0.51$	1.52	$\sigma_\varphi = 0.51$
Dispersion in Domestic Shares	0.36	$\sigma_f = 2.41$	0.14	$\sigma_f = 0$
Correlation log Value Added - Dom Shares	-0.31	$\rho = 0.73$	-0.72	$\rho = 0$
Consumer Price Gains	38.07 %		43.26%	Bias:13.65 %
Welfare Gains	17.53%		21.69%	Bias:23.73 %

Notes: The table contains an alternative calibration which does not use information on domestic expenditure shares. This model is only calibrated to the moments in the first three rows. We report our baseline calibration of Section 5.2 for comparison.

Table 15: Calibration Without Domestic Shares

Note that the model is calibrated to match the moments contained in the first three rows. The results in Table 15 show that the model is unsuccessful in predicting a distribution of import shares that is consistent with the data. Not only is there too strong a negative correlation between import intensity and firm value added, but the model also under-predicts the cross-sectional dispersion in domestic shares.