

The Gains from Input Trade with Heterogeneous Importers[†]

By JOAQUIN BLAUM, CLAIRE LELARGE, AND MICHAEL PETERS*

Firms differ substantially in their participation in foreign input markets. We develop a methodology to measure the aggregate effects of input trade that takes such heterogeneity into account. We provide a theoretical result that holds in a variety of settings: the firm-level data on value added and domestic expenditure shares in material spending is sufficient to compute the change in consumer prices due to a shock to the import environment. We characterize the bias of approaches that rely on aggregate statistics. In an application to French data, input trade reduces the prices of manufacturing products by 27 percent. (JEL D24, E31, F12, F14, L11, L25, L60)

International trade benefits consumers by lowering the prices of the goods they consume. An important distinction is that between trade in final goods and trade in intermediate inputs. While the former benefits consumers directly, the latter operates only indirectly. By allowing firms to access novel, cheaper, or higher quality inputs from abroad, input trade reduces firms' production costs and thus the prices of locally produced goods. Because intermediate inputs account for about two-thirds of the volume of world trade, understanding the normative consequences of input trade is important.

A recent body of work has incorporated input trade into quantitative trade models—see, e.g., Eaton, Kortum, and Kramarz (2011); Caliendo and Parro (2015); and Costinot and Rodríguez-Clare (2014). As in Arkolakis, Costinot, and Rodríguez-Clare (2012), these frameworks have the convenient implication that the change in consumer prices can be measured with aggregate data, and hence firm heterogeneity is irrelevant. This property, however, only holds when firms' import intensities are equalized—a feature that is at odds with the data. This is shown in Figure 1, which depicts the cross-sectional distribution of French manufacturing

*Blaum: Department of Economics, Brown University, 64 Waterman Street, Providence, RI 02912 (email: jblaum@gmail.com); Lelarge: University of Paris-Saclay, 54 Boulevard Desgranges, 92330 Sceaux, France, and CEPR (email: claire.lelarge@gmail.com); Peters: Department of Economics, Yale University, 28 Hillhouse Avenue, New Haven, CT 06511, and NBER (email: m.peters@yale.edu). We thank Costas Arkolakis, Arnaud Costinot, Dave Donaldson, Jonathan Eaton, Pablo Fajgelbaum, Lionel Fontagné, Penny Goldberg, David Weinstein, and Daniel Xu. We are grateful to seminar participants at Berkeley, Brown, Columbia, Dartmouth, ERWIT 2017, LSE, Penn State, Princeton, Stanford, UCLA, and Yale. A previous version of this paper circulated under the title "The Gains from Input Trade in Firm-Based Models of Importing." Blaum thanks the International Economics Section at Princeton for their support and hospitality.

[†]Go to <https://doi.org/10.1257/mac.20160426> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

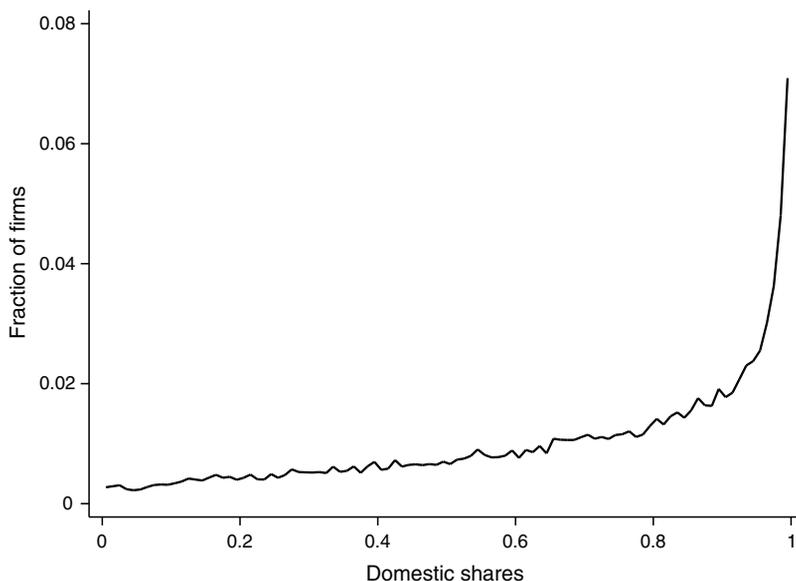


FIGURE 1. THE DISPERSION OF DOMESTIC SHARES

Note: 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.

firms' domestic expenditure shares, i.e., the share of material spending allocated to domestic inputs. These differ markedly. While the majority of importers spend more than 90 percent of their material spending on domestic inputs, some firms are heavy importers with import shares exceeding 50 percent. In this paper, we show that accounting for this heterogeneity in import behavior is crucial to quantify the aggregate effects of input trade.

We provide a methodology to measure the effect of input trade on consumer prices in environments with heterogeneous firms. In particular, we show that changes in consumer prices can be computed from firm-level data on domestic expenditure shares and value added, and we provide a closed-form expression to do so. Importantly, this formula holds in a wide class of models of importing because it does not require specific assumptions on firms' import environment.¹ By relying on firms' observable domestic shares, we circumvent the need to structurally estimate a particular model. Moreover, we do not require information on the prices and qualities of the foreign inputs, nor how firms find their suppliers, e.g., whether importing is limited by fixed costs or a process of network formation. Therefore, many positive aspects of heterogeneous import behavior across firms, such as the number of supplier countries or the distribution of spending across trading partners, are irrelevant for the link between input trade and consumer prices.

¹ Besides the aggregate quantitative models mentioned above, this class nests several firm-based frameworks used in the literature, including Halpern, Koren, and Szeidl (2015); Gopinath and Neiman (2014); Antràs, Fort, and Tintelnot (2017); and Goldberg et al. (2010).

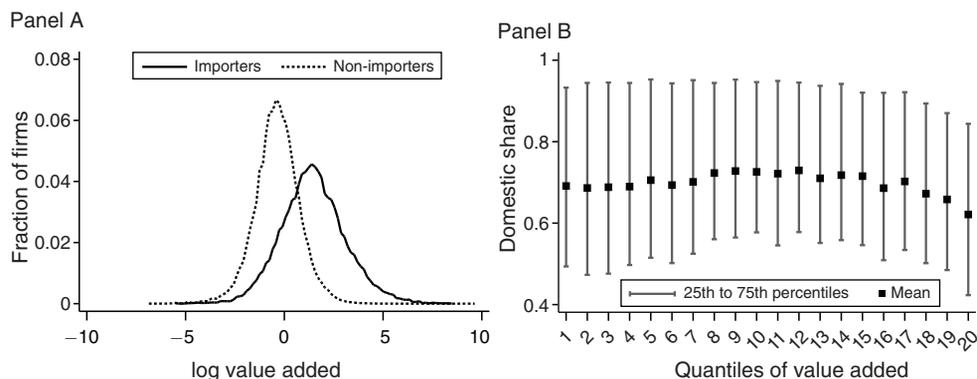


FIGURE 2. DOMESTIC SHARES AND FIRM SIZE

Notes: Panel A displays the distribution of log value added by import status. Panel B shows the mean and the twenty-fifth and seventy-fifth percentiles of domestic shares for 20 quantiles of value added for importers. The data corresponds to the population of manufacturing firms in France in 2004.

The intuition behind this result is simple. Consider the case of a reversal to input autarky, where firms can only use local inputs. Domestic consumers are affected by input trade solely through firms' unit costs. By inverting the demand system for intermediates, we can link each firm's unit cost to its spending pattern on domestic inputs. When such a demand system is CES, the unit cost reduction from importing can be recovered from the observable domestic expenditure share. In particular, a low domestic share indicates that the firm benefits substantially from input trade. In this sense, Figure 1 shows that the gains from input trade are heterogeneous at the micro-level. To correctly aggregate these firm-level gains, one needs to know each firm's relative importance in the economy. In a multi-sector general equilibrium trade model with inter-sectoral linkages, we show that the aggregate effect of input trade on the consumer price index is akin to a value-added weighted average of the firm-level gains. Hence, a key aspect of the data is how firm size and domestic shares correlate: if bigger firms feature lower domestic shares, the aggregate effects of input trade will turn out to be large.

The extent to which this is the case in France is depicted in Figure 2. In panel A, we display the distribution of value added by import status. In panel B, we focus on the population of importers and show the distribution of domestic shares for different value-added quantiles. We see that importing and firm size are far from perfectly aligned. While importers are significantly larger than non-importers, there is ample overlap in their distribution of value added. Furthermore, conditional on importing, the relationship between import intensity and size is essentially flat and there is substantial dispersion in import shares conditional on size. We show in this paper that these patterns of heterogeneity are important: models that do not match the data displayed in Figure 2 yield biased estimates of the effects of input trade on consumer prices.

This logic can be extended to study shocks other than a reversal to input autarky. More precisely, we show that the effect of any shock to the import environment (e.g., a trade liberalization episode or an increase in foreign input prices) on the

domestic consumer price index is fully determined from the joint distribution of value added and the *changes* in firms' domestic shares. If such changes are observed, one can directly calculate the aggregate consequences of the shock. A limitation of our approach is that it is not well suited to study counterfactual shocks or policies where such changes are unobserved (see Arkolakis, Costinot, and Rodríguez-Clare 2012).

A key aspect of our methodology is that we can measure firms' unit cost changes independently of the macroeconomic environment. We show that in a wide class of models of importing the micro and macro parts of the model can be effectively separated and we exploit this property to easily handle rich macroeconomic settings. In particular, we can consider multi-sector general equilibrium environments with realistic input-output linkages and different assumptions about competition in output markets. We consider both a CES monopolistic competition model and a setting with variable markups and we show that the micro data on firm size and domestic expenditure shares is sufficient to compute changes in consumer prices in these different settings. Moreover, we provide closed-form expressions to do so.

To assess the importance of the micro data, we provide an explicit expression for the difference in the gains from trade implied by aggregate models and our approach based on micro data. By relying on aggregate statistics, instead of the micro data in Figures 1 and 2, aggregate models yield biased results. While the magnitude of this bias depends on the underlying micro data, its sign only depends on a small set of parameters. In particular, aggregate models imply gains from trade that are too high whenever the elasticity of substitution between domestic and foreign inputs is small, or the elasticity of consumers' demand is large.

We apply our methodology to the population of manufacturing firms in France. We first measure the change in consumer prices relative to autarky. We estimate the distribution of trade-induced changes in unit costs across firms implied by the distribution of domestic expenditure shares displayed in Figure 1 above. While the median unit cost reduction is 11 percent, it exceeds 80 percent for 10 percent of the firms. We then aggregate these firm-level gains to compute the consumer price gains by relying on the joint distribution of domestic shares and value added displayed in Figure 2 above. We find that consumer prices of manufacturing products would be 27 percent higher if French firms were not allowed to source intermediate inputs from abroad. An analysis based on aggregate data would overestimate this change in consumer prices by about 10 percent.

There are three reasons why our estimate of the consumer price gains exceeds the median firm-level gains. First, the dispersion in firm-level gains, displayed in Figure 1, is valued by consumers given their elastic demand. Second, the weak but positive relation between import intensity and firm size, shown in Figure 2, is beneficial because the endogenous productivity gains from importing and firm efficiency are complements. Third, there are important linkages between firms whereby non-importers buy intermediates from importing firms.

An important parameter in our analysis is the elasticity of substitution between domestically sourced and imported inputs. Because firm-based models of importing do not generate a standard gravity equation for aggregate trade flows, we devise a strategy to identify this elasticity 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. Using the variation across firms is important as we obtain a value for the elasticity close to two.

We then turn to counterfactuals other than autarky. In particular, we study shocks that make foreign inputs more expensive (e.g., a currency devaluation). To do so, we need to fully specify the import environment to predict firms' domestic shares after the shock. We consider a standard framework where importing is subject to fixed costs and evaluate quantitatively whether the micro data on size and domestic shares is important for the estimates of the effects. More precisely, we compare different parameterizations of the model which vary in the extent to which they match the data displayed in Figures 1 and 2. First, we find that versions of the model that do not match the data in Figures 1 and 2 tend to overpredict the increase in consumer prices by 13 to 18 percent. For example, models where efficiency is the single source of heterogeneity imply a one-to-one, and hence counterfactual, relation between firm size and domestic shares and predict effects that are too large. Second, different models that match the data in Figures 1 and 2 predict very similar effects of the shock. Hence, conditional on the observable micro data, the details of the import environment, e.g., whether firms differ in fixed costs or in their efficiency as importers, are not crucial to predict changes in consumer prices. These results suggest that the sufficiency of the data in Figures 1 and 2, which holds exactly for the case autarky, quantitatively extends to other counterfactuals.

Another reason why approaches based on aggregate data may yield biased estimates of the gains from trade pertains to the elasticity of substitution between domestic and imported inputs. A common approach in the literature is to discipline this parameter with the aggregate trade elasticity. Holding the elasticity of substitution constant, the implied trade elasticity varies across models. With our baseline parameters, in particular with an elasticity of substitution close to two, a model with fixed costs calibrated to the data in Figures 1 and 2 implies a trade elasticity of 4.5. This is in the ballpark of the estimates in the literature. In contrast, an aggregate model implies a value close to 1. To match a trade elasticity of 4.5, the aggregate model requires an elasticity of substitution of 5.5 which reduces the gains from trade by a factor of 4. Thus, relying on aggregate data can lead to substantial biases.

Related Literature.—Our paper contributes to the literature that measures how consumer prices are affected by international trade—see Feenstra (1994); Broda and Weinstein (2006); and Arkolakis, Costinot, and Rodríguez-Clare (2012). This literature studies trade in final goods and uses observable expenditure shares to measure the change in the consumer price index. We apply a similar methodology to the context of firms importing intermediate inputs from abroad. In this environment, two important differences arise. First, we measure the distribution of firms'

unit costs rather than final good prices directly. To do so, we exploit firm-level customs data which allows us to compute expenditure shares at the micro level. Second, to map the firms' units costs into the consumer price index, we specify a macroeconomic environment including the structure of product market competition and input-output linkages. We find that, when firms' import intensities are heterogeneous, the results in Arkolakis, Costinot, and Rodríguez-Clare (2012) do not apply: aggregate statistics are no longer sufficient to compute the change in consumer prices and the entire distribution of domestic shares and firm size is required. We provide a formula to map this distribution to the change in the consumer price index.

Our paper is also related to a literature that studies input trade in quantitative models with firm heterogeneity, e.g., Gopinath and Neiman (2014); Halpern, Koren, and Szeidl (2015); Antràs, Fort, and Tintelnot (2017); and Ramanarayanan (2017). While these contributions analyze input trade by fully specifying and estimating structural models of importing, we follow a different approach and measure firms' unit cost changes directly from the micro data. Our approach has two benefits. First, we can be agnostic about several components of the theory. Hence, our estimates do not rely on particular assumptions about firms' import environment, such as the qualities of foreign inputs or whether firms' extensive margin is limited by fixed costs. Second, our methodology is particularly useful to study the macroeconomics implications of input trade, because we can take general equilibrium effects into account and allow for input-output linkages and variable markups. Building these features into a structural estimation would entail substantial computational complexities. Using our methodology, we can incorporate these elements into the analysis easily. The limitation of our approach is that our formula can be directly applied only in situations where expenditure shares are observed, for instance to infer the consumer price gains of historical episodes of trade liberalization or to measure the gains from trade relative to autarky. In addition, our approach is not suited to measure changes in welfare as the resources spent by firms in attaining their sourcing strategies cannot be recovered from observable data.²

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 Topalova and Khandelwal (2011).³ Our results are complementary to this literature. We show that the domestic expenditure share can be used to measure the effect of the policy on firm productivity holding technology constant. These static productivity gains do not capture the effect that a trade liberalization may have on firms' technologies via R&D or quality upgrading—see Fieler, Eslava, and Xu (2018) or Bøler, Moxnes, and Ulltveit-Moe (2015). By focusing on the domestic expenditure share as the outcome of interest (instead of standard measures of firm productivity), one can disentangle the static from the dynamic effects of the policy. If micro data on value added is also available, our results can be used to gauge the full effect of the policy on consumer

² We show, however, that the change in consumer prices provides an upper bound for the change in welfare.

³ Kasahara 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 (2014) for a more general empirical framework to study firm performance in international markets.

prices in general equilibrium. Amiti et al. (2017) use a related methodology to study the effect of China's WTO entry on the US consumer price index.

The remainder of the paper is structured as follows. Section I lays out the theory that we consider to measure the effect of input trade on firms' unit costs and consumer prices. Section II deals with the biases associated with models that do not fully exploit the micro data. The application to France is contained in Section III. Section IV concludes.

I. Theory

In this section, we lay out the theoretical framework of importing that we use to measure the effects of input trade. In Section IA, we study the import problem faced by a single firm and relate the domestic expenditure share to the effect of input trade on the unit cost. In Section IB, we embed this problem into a general equilibrium macroeconomic model and show that the information contained in firms' domestic spending shares and size is sufficient to calculate the impact of shocks to the import environment on consumer prices. In particular, a wide class of models predicts the exact same change in consumer prices given the micro data.

A. Micro: Firms and 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:

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

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

$$(3) \quad x_I = h_i \left([q_{ci} z_c]_{c \in \mathcal{L}_i} \right),$$

where $\gamma, \beta_i \in (0, 1)$ and $\varepsilon > 1$. Hence, the firm combines intermediate inputs x with primary factors l , which for simplicity we refer to as labor, in a Cobb-Douglas fashion with efficiency φ_i .⁴ 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 . We refer to β_i as the firm's home bias. 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.⁵

⁴We 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 III, we consider labor and capital.

⁵We discuss below how to generalize the results of this section when the Cobb-Douglas and CES functional forms in (1)–(2) are not satisfied. In particular, we consider the cases where (1) takes a CES form so that intermediate spending shares are not equalized, and a multi-product version of (2), where domestic and foreign inputs are closer substitutes within a product nest. We abstract from the product dimension in the main text because we do not observe firms' domestic spending at the product level.

Foreign inputs are aggregated according to a constant returns to scale, potentially firm-specific production function $h_i(\cdot)$.⁶ An important endogenous object in the production structure is the set of foreign countries the firm sources from, which we denote by \mathcal{L}_i and henceforth refer to as the sourcing strategy. We do not impose any restrictions on how \mathcal{L}_i is determined until Section IIID. As far as the market structure is concerned, we assume that the firm takes prices of domestic and foreign inputs $(p_D, [p_{ci}])$ as parametric, i.e., it can buy any quantity at given prices. Note that p_{ci} includes all variable trade costs and is allowed to be firm-specific. Finally, we assume that labor can be hired frictionlessly at a given wage w .

This setup describes a class of models of importing that have been used in the literature. First, it nests aggregate quantitative trade models (Eaton, Kortum, and Kramarz 2011; Costinot and Rodríguez-Clare 2014; Caliendo and Parro 2015). In these models, firms' import intensities are equalized. In the above setup, this corresponds to the case where firms' sourcing strategies are equalized, all firms face the same prices and qualities, and there is no heterogeneity in the home bias (i.e., $\mathcal{L}_i = \mathcal{L}$, $[p_{ci}, q_{ci}] = [p_c, q_c]$, $\beta_i = \beta$). Second, it nests a variety of recent examples of firm-based models of importing, e.g., Halpern, Koren, and Szeidl (2015); Gopinath and Neiman (2014); Antràs, Fort, and Tintelnot (2017); Kasahara and Rodrigue (2008); Lu, Mariscal, and Mejía (2017); Amiti, Itskhoki, and Konings (2014); and Goldberg et al. (2010).⁷ A unifying feature of these models is that firms engage in input trade because it lowers their unit cost of production via love of variety and quality channels. Additionally, these contributions generate heterogeneity in firms' import intensities through variation in the sourcing strategies \mathcal{L}_i . Characterizing firms' optimal sourcing strategy \mathcal{L}_i in economies with fixed costs can be nontrivial and requires stringent assumptions. One of the main results of this paper is that, to measure the effect of input trade on consumer prices, the solution to this problem is not required.

The assumptions made above, most importantly parametric input prices and constant returns to scale, guarantee that the unit cost is constant *given the sourcing strategy* \mathcal{L} . This separability between the intensive and extensive margin allows us to characterize the unit cost at the firm level without solving for the optimal sourcing set nor specifying the demand the firm faces. Formally, the unit cost is given by

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

⁶ Note 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 (2012) and Blaum, Lelarge, and Peters (2017).

⁷ To nest the contributions that allow for multiple products (see, e.g., Halpern, Koren, and Szeidl 2015 or Goldberg et al. 2010), our production function needs to be extended. In particular, the intermediate input bundle x would be given by $x = \prod_k x_k^{1/\theta}$. Here x_k denotes the intermediate input bundle for product k , which is given by equation (2). In online Appendix O1, we extend the theoretical results of this section to this case. Regarding Antràs, Fort, and Tintelnot (2017), note that they consider a model of importing in the spirit of Eaton and Kortum (2002) instead of a variety-type model. Their Fréchet assumption implies that these models are isomorphic.

subject to (2)–(3). For simplicity, we refer to the unit cost as $u_i(\mathcal{L}_i)$. Standard calculations imply that there is an *import price index* given by

$$A(\mathcal{L}_i, [q_{ci}], [p_{ci}], h_i) \equiv m_I/x_I,$$

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

$$(4) \quad Q_i(\mathcal{L}_i) = \left(\beta_i^\varepsilon (p_D/q_D)^{1-\varepsilon} + (1 - \beta_i)^\varepsilon A_i(\mathcal{L}_i)^{1-\varepsilon} \right)^{\frac{1}{1-\varepsilon}}.$$

Because $A_i(\mathcal{L}_i)$ is decreasing in the size of \mathcal{L}_i , i.e., $A_i(\mathcal{L}_i) \leq A_i(\mathcal{L}'_i)$ whenever $\mathcal{L}_i \supseteq \mathcal{L}'_i$, firms with more trading opportunities abroad benefit from lower input prices. Additionally, this price index depends on a number of unobserved parameters related to the trading environment, e.g., the prices and qualities of the foreign inputs $[q_{ci}, p_{ci}]$ and firms' production function for imported inputs h_i . Instead of imposing sufficient structure to be able to estimate $A_i(\mathcal{L}_i)$ and $Q_i(\mathcal{L}_i)$, we use the fact that the *unobserved* price index $Q_i(\mathcal{L}_i)$ is related to the *observed* expenditure share on domestic inputs $s_{Di} = p_D z_D / (m_I + p_D z_D)$ via

$$(5) \quad s_{Di} = (Q_i(\mathcal{L}_i))^{\varepsilon-1} \beta_i^\varepsilon \left(\frac{q_D}{p_D} \right)^{\varepsilon-1}.$$

It then follows that the firm's unit cost is given by

$$(6) \quad u_i(\mathcal{L}_i) = \frac{1}{\tilde{\varphi}_i} w^{1-\gamma} (Q_i(\mathcal{L}_i))^\gamma = \frac{1}{\tilde{\varphi}_i} \times (s_{Di})^{\frac{\gamma}{\varepsilon-1}} \times \left(\frac{p_D}{q_D} \right)^\gamma w^{1-\gamma},$$

where $\tilde{\varphi}_i \equiv \varphi_i \beta_i^{\frac{\varepsilon\gamma}{\varepsilon-1}} (1-\gamma)^{1-\gamma} \gamma^\gamma$. Equation (6) is a sufficiency result: conditional on the firm's domestic expenditure share s_{Di} , no aspects of the import environment, including the sourcing strategy \mathcal{L}_i , the prices p_{ci} , the qualities q_{ci} , or the technology h_i , affect the unit cost. The domestic expenditure share conveniently encapsulates all the information from the import environment that is relevant for the unit cost.

This equation allows us to express trade-induced changes in firms' unit costs in terms of observables. To see this, consider an arbitrary shock to the import environment, i.e., a change in foreign prices, qualities, trade costs or the sourcing strategy. The change in the firm's unit cost resulting from the shock, holding prices (p_D, w) constant, is given by

$$(7) \quad \ln \left(\frac{u'_i}{u} \right) \Big|_{p_D, w} = \frac{\gamma}{1-\varepsilon} \times \ln \left(\frac{s_{Di}}{s'_{Di}} \right),$$

where u'_i and s'_{Di} denote the unit cost and the domestic expenditure share after the shock. Intuitively, an adverse trade shock, such as an increase in foreign prices or a reduction in the set of trading partners, hurts the firm by increasing 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. Hence, the effect of input trade on firm productivity can be directly measured from the data, without having to fully specify and estimate a structural model of importing.⁸

Equation (7) is akin to a firm-level analogue of Arkolakis, Costinot, and Rodríguez-Clare (2012). In the same vein as consumers gain purchasing power by sourcing cheaper or complementary products from abroad, firms can lower the effective price of material services by tapping into foreign input markets. While this analogy works at the firm level, it breaks at the aggregate level. We show below, in the context of a macro model, that there is no aggregate statistic that is sufficient to measure changes in consumer prices.

This result implies that domestic expenditure shares are the crucial empirical object to learn about the relationship between input trade and production costs. Other micro moments such as characteristics of the set of sourcing partners, the distribution of expenditure across sourcing countries, or whether or not international sourcing is hierarchical, while potentially interesting per se, are not important for the relationship between input trade and unit costs. This property can be useful for applied empirical work, e.g., to study the effect of trade liberalizations on firm productivity (see Pavcnik 2002, Amiti and Konings 2007, or De Loecker et al. 2016). According to (7), the causal effect of the policy on firms' unit costs can be measured from the change in firms' domestic shares which are due to the policy.⁹

While equation (7) is a partial equilibrium result, we note that it identifies the dispersion in unit cost changes across firms in general equilibrium and hence the distributional effects of input trade. One special case where this is especially apparent is input autarky. In that case, $s'_{Di} = 1$ and (7) reduces to $\gamma/(1 - \varepsilon) \times \ln(s_{Di})$. Hence, Figure 1 fully summarizes how each importer's unit cost (relative to a domestic producer) would change if it was forced to source its inputs only domestically.

⁸In online Appendix O1, we show how the result in (7) can be extended to a more general production function than (1)–(3). In particular, we consider the cases where (i) domestic and foreign inputs are not combined in a CES fashion; (ii) the output elasticity of material inputs is not constant; and (iii) firms source multiple products from different countries. We also discuss what additional data, relative to the result in (7), is required to perform counterfactual analysis in these cases. For the multi-product version of our model, (7) generalizes to $\ln(u'_i/u_i)|_{p_D, w} = \gamma/(1 - \varepsilon) \sum_{k=1}^K \eta_k \ln(s^k_{Di}/s^k_{Di})$, where η_k are the Cobb-Douglas weights in the intermediate input production (see footnote 7). In our application, we consider the setup in (1)–(3) because we do not observe firms' domestic shares at the product level.

⁹We note that opening up to trade may induce firms to engage in productivity enhancing activities that directly increase efficiency φ such as R&D—see, e.g., Fieler, Eslava, and Xu (2018). Such increases in complementary investments are not encapsulated in (7), which only measures the static gains from trade holding efficiency fixed.

B. Macro: Consumer Prices in General Equilibrium

We now embed the above model of firm behavior in a macroeconomic environment to link input trade to consumer prices. The micro result in (7) is crucial as it allows us to measure the firm-level unit cost reductions directly from the micro data, albeit in partial equilibrium. To aggregate these firm-level gains taking general equilibrium effects into account, 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 the 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. To isolate the effect of input trade, we abstract from trade in final goods.

As a baseline case, we consider a multi-sector CES monopolistic competition environment. We generalize our results to a setting with variable markups in Appendix C. 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 labor inelastically and whose preferences are given by

$$(8) \quad U = \prod_{s=1}^S C_s^{\alpha_s} \quad \text{and} \quad C_s = \left(\int_0^{N_s} c_{is}^{\frac{\sigma_s-1}{\sigma_s}} di \right)^{\frac{\sigma_s}{\sigma_s-1}},$$

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) 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. To allow for the fact that consumers spend part of their budget on goods outside of the manufacturing sector, we assume sector S to be comprised of firms that do not trade inputs and refer to it as the non-manufacturing sector.

We assume the following structure of roundabout production, which is also used in Caliendo and Parro (2015). Firms use a sector-specific domestic input that is produced using the output of all other firms in the economy according to

$$(9) \quad z_{Ds} = \prod_{j=1}^S Y_{js}^{\zeta_j^s} \quad \text{and} \quad Y_{js} = \left(\int_0^{N_j} y_{\nu js}^{\frac{\sigma_j-1}{\sigma_j}} d\nu \right)^{\frac{\sigma_j}{\sigma_j-1}},$$

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 above, we can express the effect of input trade on the consumer price index associated with (8) in terms of observables. Given the expression for firms' unit costs (7), the CES demand and monopolistic competition structure, the consumer price index for sector s is given by

$$(10) \quad P_s = \mu_s \left(\int_0^{N_s} u_i^{1-\sigma_s} di \right)^{\frac{1}{1-\sigma_s}}$$

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

where $\mu_s \equiv \sigma_s/(\sigma_s - 1)$ is the markup in sector s and we treat labor as the numeraire. Equation (10) shows that, holding domestic input prices fixed, the effect of input trade on the sectoral price index is an efficiency-weighted average of the firm-level gains. While firm efficiency $\tilde{\varphi}_i$ is not observed, it can be recovered up to scale from data on value added and domestic spending as¹⁰

$$(11) \quad va_i \propto \left(\tilde{\varphi}_i (s_{Di})^{\frac{\gamma_s}{1-\varepsilon_s}} \right)^{\sigma_s-1}.$$

Consider again any shock to the import environment, i.e., a change in foreign prices, qualities, trade costs, or the sourcing strategies. Combining (10) and (11), the change in the sectoral consumer price index due to the shock is given by

$$(12) \quad \ln \left(\frac{P'_s}{P_s} \right) = \gamma_s \ln \left(\frac{p'_{Ds}}{p_{Ds}} \right) + \frac{1}{1-\sigma_s} \ln \left(\int_0^{N_s} \omega_i \left(\frac{s_{Di}}{s'_{Di}} \right)^{\frac{\gamma_s}{1-\varepsilon_s}(1-\sigma_s)} di \right),$$

where ω_i denotes firm i 's share in sectoral value added. Equation (12) shows that shocks to firms' ability to source inputs from abroad affect consumer prices through two channels. First, there is a direct effect stemming from firms in sector s changing their intensity to source inputs internationally—this is the last term in (12). Importantly, this term can be directly computed from the micro data. Second, there is an indirect effect as the price of domestic inputs changes because of input-output linkages, p'_{Ds}/p_{Ds} . Because of the structure of roundabout production in (9), this indirect effect can be computed from the system of equations in (12). This is the content of our main proposition.

PROPOSITION 1: *Consider a shock to firms' import environment and let P and P' be the consumer price indices before and after the shock. Define the direct cost reduction of input trade in sector s as*

$$(13) \quad \Lambda_s = \frac{1}{1-\sigma_s} \ln \left(\int_0^{N_s} \omega_i \left(\frac{s_{Di}}{s'_{Di}} \right)^{\frac{\gamma_s}{1-\varepsilon_s}(1-\sigma_s)} di \right).$$

¹⁰This assumes that the data on value added does not record firms' expenses to attain their sourcing strategies. If it did, one should express (11) in terms of sales or employment.

The change in consumer prices is then given by

$$(14) \quad \ln\left(\frac{P'}{P}\right) = \boldsymbol{\alpha}'(\boldsymbol{\Gamma}(\mathbf{I} - \boldsymbol{\Xi} \times \boldsymbol{\Gamma})^{-1}\boldsymbol{\Xi} + \mathbf{I}) \times \boldsymbol{\Lambda},$$

where $\boldsymbol{\Lambda} = [\Lambda_1, \Lambda_2, \dots, \Lambda_S]$, Λ_s is given in (13), $\boldsymbol{\Xi} = [\zeta_j^s]$ is the $S \times S$ matrix of production interlinkages, $\boldsymbol{\alpha}$ is the $S \times 1$ vector of demand coefficients, \mathbf{I} is an identity matrix, and $\boldsymbol{\Gamma} = \text{diag}(\gamma)$, where γ is the $S \times 1$ vector of input intensities.

In the special case of a reversal to input autarky, the increase in consumer prices is given by (14), where Λ_s is given by

$$(15) \quad \Lambda_s^{\text{Aut}} = \frac{1}{1 - \sigma_s} \ln\left(\int_0^{N_s} \omega_i s_{Di}^{\frac{\gamma_s}{1 - \varepsilon_s}(1 - \sigma_s)} di\right) \geq 0.$$

PROOF:

See Appendix A. ■

Proposition 1 shows that the micro data on value added and changes in domestic shares is sufficient to fully characterize the consumer price consequences of trade-induced shocks in the class of models considered in this section. In particular, the change in consumer prices can be directly computed by using (13) and (14) given parameters for consumer demand and production.¹¹ Moreover, these equations highlight that the micro data is required only to compute the direct cost reductions of input trade, i.e., Λ_s . The other terms in (14) reflect the general equilibrium effect of input-output linkages across firms, by which changes in importers' unit costs diffuse through the economy. To see this, note that in the case of a single sector economy (14) simplifies to

$$\ln\left(\frac{P'}{P}\right) = \frac{\Lambda}{1 - \gamma},$$

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

A key aspect of Proposition 1 is that “changes in consumer prices can be measured” without specifying many details of the micro part of the model. In particular, we do not have to parametrize (and estimate) the import environment $[p_{ci}, q_{ci}, h_i, \beta_i]$ or characterize firms' sourcing strategy \mathcal{L}_i . Hence, we do not need to take a stand on whether firms' extensive margin is shaped by the presence of fixed costs or a process of search or network formation. These aspects are irrelevant for consumer prices conditional on the data on size and domestic spending. Moreover, any estimated

¹¹ As in Feenstra (1994) and Broda and Weinstein (2006), our results focus on changes in consumer prices and therefore may not capture the full welfare effects of input trade if firms need to spend resources to find their trading partners. This feature is not specific to theories of importing but also arises in models of exporting. For example, the welfare formula of Arkolakis, Costinot, and Rodríguez-Clare (2012) precisely relies on the condition that profits are a constant share of aggregate income. Whether or not this condition is satisfied depends on details of the environment which we did not have to specify to derive Proposition 1. We note, however, that the change in consumer prices provides an upper bound for the change in welfare in all models in our class. In Section IIID, we provide examples of fully-specified models of importing where this bound is tight or where consumer prices and welfare are substantially different.

model in our class will arrive at the exact same number as long as it is successfully calibrated to the observable micro data.

Proposition 1 can be applied to the analysis of observed policies, i.e., in situations where the researcher has access to both s_{Di} and s'_{Di} . To the extent that the change in domestic shares can be attributed to the policy, the effect of the policy on consumer prices can be readily calculated from equation (14).¹² A special case of interest is a reversal to input autarky. Because firms' counterfactual domestic shares are given by unity, the change in domestic spending shares between the current trade equilibrium and autarky is simply given by their level in the observed equilibrium. Input autarky is therefore a policy which is trivially observable in any firm-level dataset that contains information on firms' domestic spending patterns. The data contained in Figures 1 and 2 is therefore sufficient to calculate the gains from input trade. We measure these gains for the French economy in Section III.

An advantage of our methodology relative to approaches that estimate an entire model of import behavior is related to computational complexity. Our approach allows for multiple sectors with a rich input-output structure, strategic pricing, and takes general equilibrium interactions into account. Antràs, Fort, and Tintelnot (2017), for example, assume that wages are not affected by input trade but determined in an outside sector. Halpern, Koren, and Szeidl (2015) use a single sector partial equilibrium framework. Estimating these models in full general equilibrium with sectoral interlinkages and variable markups would entail substantial computational difficulty. The reason is that the solution to firms' optimal sourcing problem, which is already challenging in models with fixed costs, interacts with finding the equilibrium market clearing prices.

A limitation of our methodology is that it cannot be directly applied when firms' domestic shares after the shock are not observed. In this case, the entire import environment $[p_{ci}, q_{ci}, h_i, \beta_i]$ and the extensive margin mechanism need to be spelled out in the context of a particular model.

Variable Markups.—While Proposition 1 was derived for the familiar CES monopolistic competition environment, it can be extended to more general settings where competition among firms, and hence the distribution of markups, responds endogenously to changes in the trading environment. This might be important in the context of input trade. If markups are increasing in productivity and importing increases productivity differentially across firms, changes in trade policy will change firms' relative unit costs and hence the markups they post. In particular, if large, productive firms have higher import shares, the possibility of input trade increases the dispersion in unit costs and hence the dispersion of markups and the extent of misallocation. This channel, by which input trade may be anticompetitive, is different from the mechanisms studied in the literature where imports of final goods promote domestic competition—see Edmond, Midrigan, and Xu (2015).

¹² In practice, one needs to use changes in firms' domestic shares that are only due to the policy. In the context of a trade liberalization episode, one can often use the change in policy to construct instruments. Note that a similar identification challenge arises in structural exercises. Gopinath and Neiman (2014), for example, assume that the entire decline in aggregate import spending is due to an increase in foreign import prices caused by the devaluation.

Our methodology is well suited to take these considerations into account. In Appendix C, we show that the data on domestic expenditure shares and firm size continues to be sufficient to calculate the change in consumer prices in any model where markups are only a function of relative prices. One specific example where this is the case is the Atkeson and Burstein (2008) model with either Cournot or Bertrand competition. We derive the analogue of Proposition 1 for that model in Appendix C.

II. The Importance of Firm Heterogeneity

The analysis so far established that size and the domestic expenditure share are the only two relevant dimensions of firm heterogeneity as far as the effect of trade shocks on consumer prices is concerned. Existing approaches in the literature either abstract from firm heterogeneity altogether and rely on aggregate statistics or do not target the joint distribution of domestic expenditure shares and size. In this section, we use Proposition 1 to assess the extent to which this is consequential.

The Bias of Aggregate Models.—Consider first aggregate approaches where firms' domestic expenditure shares are by construction equalized—see Eaton, Kortum, and Kramarz (2011); Caliendo and Parro (2015); and Costinot and Rodríguez-Clare (2014). The gains from input trade relative to autarky in these models can be computed via Proposition 1 where the direct cost reductions Λ_s^{Aut} are given by

$$(16) \quad \Lambda_{Agg,s}^{Aut} = \frac{\gamma_s}{1 - \varepsilon_s} \ln(s_{D_s}^{Agg}),$$

where $s_{D_s}^{Agg}$ is the aggregate domestic expenditure share in sector s . Hence, as in Arkolakis, Costinot, and Rodríguez-Clare (2012), these frameworks have the benefit of only requiring aggregate data. Figure 1, however, shows that their implication of equalized domestic shares is rejected in the micro data and Proposition 1 shows that this has aggregate consequences. Using (15) and (16), we define the bias from measuring the price reduction in sector s through the lens of an aggregate model as

$$(17) \quad Bias_s \equiv \Lambda_{Agg,s}^{Aut} - \Lambda_s^{Aut} = \frac{\gamma_s}{\varepsilon_s - 1} \times \ln \left[\frac{\left(\int_0^{N_s} \omega_i s_{D_i}^{\chi_s} di \right)^{1/\chi_s}}{\int_0^{N_s} \omega_i s_{D_i} di} \right],$$

where $\chi_s = \gamma_s(\sigma_s - 1)/(\varepsilon_s - 1)$. As long as $\chi_s \neq 1$, the heterogeneity in import shares induces a bias in the estimates of the gains from trade of aggregate models. 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, (17) together with Jensen's inequality directly implies that

$$(18) \quad Bias_s > 0 \quad \text{if and only if} \quad \chi_s = \frac{\gamma_s(\sigma_s - 1)}{\varepsilon_s - 1} > 1.$$

As long as $\chi_s > 1$, which is the case when consumer demand is elastic (σ_s is large) or the elasticity of unit costs with respect to the domestic share is large ($\gamma/(\varepsilon - 1)$ is high), an analysis based on aggregate data would imply consumer gains that are too large. The economic intuition of this result 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. 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, (11) 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.¹³

Note also that $\Lambda_{Agg,s}^{Aut}$ provides a bound for the change in consumer prices resulting from trade-induced shocks. More specifically, (17) and (18) directly imply that if $\chi > 1$ ($\chi < 1$) an aggregate model provides an upper (lower) bound for the effect of input trade on consumer prices for *any* model that is calibrated to the aggregate domestic share. Thus, aggregate approaches in the spirit of Arkolakis, Costinot, and Rodríguez-Clare (2012) can be used to derive a bound in cases where the micro data is not available. In the quantitative analysis in Section IIID, we show that this intuition carries over to counterfactuals beyond autarky.

The Bias of 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 non-degenerate joint distribution of import intensity and firm size. Gopinath and Neiman (2014) and Ramanarayanan (2017) 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. These models also imply that all importers are larger than domestic firms. Figure 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 yield biased estimates of the gains from trade by construction. Moreover, by assigning the largest unit cost reductions to the most efficient firms, these models tend to magnify the aggregate gains from trade.

Antràs, Fort, and Tintelnot (2017) and Halpern, Koren, and Szeidl (2015) allow for heterogeneity in efficiency and fixed costs and thus generate a nontrivial distribution of value added and domestic spending shares. The question is whether the model-implied joint distribution of domestic shares and firm size looks like the one

¹³ The following example may be instructive. Consider an economy where firms differ in their domestic shares but value added is equalized. Looking at the data through the lens of an aggregate model, one would conclude that innate efficiency is also equalized across firms. Equation (11) however implies that firm efficiency has to vary given a common level of value added. Whether or not consumers prefer the misspecified autarky allocation with equalized efficiency depends on χ . If $\chi > 1$, the absence of productivity dispersion will imply higher consumer prices and therefore higher gains from trade in an aggregative framework.

in the data. In Section III D, we show quantitatively that failing to match such data can lead to substantial differences in the estimates of the gains from trade both for the case of autarky and other counterfactuals.

III. Empirical Application

We now take the framework laid out above to data on French firms to measure the effect of input trade on consumer prices. In order to emphasize the link between input trade and domestically produced goods, we focus our analysis on manufacturing firms. We first focus on a reversal to input autarky and compute the resulting change in consumer prices directly from the observed micro data. We then study shocks that make foreign inputs more expensive without leading the economy into autarky.

A. Data

The main source of information we use is a firm-level dataset from France.¹⁴ A detailed description of how the data is constructed is contained in Appendix B. We observe import flows for every manufacturing firm in France from the official custom files. Manufacturing firms account for 30 percent of the population of French importing firms and 53 percent 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 match this dataset to fiscal files which contain detailed information on firm characteristics. Most importantly, we retrieve the total input expenditures from these files and then compute domestic spending as the difference between total input expenditures and total imports. 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. Appendix Table B1 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 .

B. 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 γ_s and the elasticities of substitution ε_s , as well as the sector-specific demand elasticities σ_s . We identify the input-output structure ζ_j^s and the aggregate demand parameters α_s from the input-output tables. This allows us to account for the non-manufacturing sector and doing so is quantitatively important.

¹⁴ This dataset is also used in Eaton, Kortum, and Kramarz (2011) and Blaum, Lelarge, and Peters (2017). Similar data is available for many countries, among others Hungary (Halpern, Koren, and Szeidl 2015), Belgium (De Loecker 2011), Slovenia (De Loecker and Warzynski 2012), Indonesia (Amiti and Konings 2007), and Chile (Kasahara and Rodrigue 2008).

TABLE 1—STRUCTURAL PARAMETERS

Industry	ISIC	α_s (percent)	σ_s	γ_s	Value added share (percent)	$s_{D_s}^{Agg}$
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 plastic 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	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
Recycling, nec.	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 (19). 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. The value-added 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 $s_{D_s}^{Agg} = \sum_{i=1}^n s_{Dis} \times \omega_{is}$, where ω_{is} is the firm share in sectoral value added. See the Appendix for the details.

Identification of α and ζ .—We compute the demand parameters α_s and the matrix of input-output linkages ζ_j^s using data from the French input-output tables on the distribution of firms' intermediate spending and consumers' expenditure by sector.¹⁵ Sectors are classified at the 2-digit 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

$$(19) \quad \zeta_j^s = \frac{Z_j^s}{\sum_{j=1}^S Z_j^s} \quad \text{and} \quad \alpha_s = \frac{E_s}{\sum_{j=1}^S E_j}.$$

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. The results for the consumption shares α_s for each sector are contained in column 1 of Table 1. The non-manufacturing sector is important as it accounts for a large share of the budget of consumers. For brevity, we report the input-output matrix ζ_j^S in online Appendix O2.

¹⁵ See online Appendix O2 for a detailed description of how we construct the input-output matrix.

Estimation of ε , σ , and γ .—To identify the elasticity of substitution ε , the intermediate input share γ , and the demand elasticity σ , we turn to the French micro data. We follow Oberfield and Raval (2014) to measure the demand elasticities σ_s from firms' profit margins, i.e., the ratio of revenue to total costs,

$$(20) \quad \frac{p_i y_i}{Cost_i} = \frac{\sigma}{\sigma - 1},$$

where $p_i y_i$ is firm revenue and $Cost_i$ denotes production costs, encompassing the wage bill and total input expenditure. We compute averages at the sector level to obtain σ_s . Column 2 of Table 1 contains the estimates which, consistent with the literature, are between 3 and 4.

To identify ε , note that firm output can be written as

$$(21) \quad y_i = \tilde{\varphi}_i s_{Di}^{-\frac{\gamma_s}{\varepsilon_s - 1}} l_i^{1 - \gamma_s} m_i^{\gamma_s} \times B,$$

where m_i is total material spending by firm i and B collects all general equilibrium variables, which are constant across firms within an industry. By expressing output in terms of *spending* in materials instead of quantities, (21) shows that we can identify ε_s from variation in domestic expenditure shares holding material spending fixed. Intuitively, the domestic share is akin to a productivity shifter. We implement this idea with two complementary approaches. Our first method relies on simple accounting identities using firms' factor shares. It has the benefit that it is straightforward to implement and maps directly into our theory. Our second approach follows the recent literature on production function estimation, in particular De Loecker (2011), and is discussed in more detail in Appendix B.

The approach based on observed factor shares is a simple and easy-to-implement benchmark (see, e.g., Syverson 2011 and Oberfield and Raval 2014). According to our theory, we can identify γ directly from firms' material spending shares¹⁶, i.e.,

$$(22) \quad \frac{m_i}{p_i y_i} = \gamma \frac{\sigma - 1}{\sigma}.$$

To estimate ε , we express (21) in terms of firm revenue:

$$(23) \quad \ln(p_i y_i) = \Phi + \beta \ln(l_i) + \gamma \ln(m_i) + \ln(\vartheta_i),$$

¹⁶In our analysis, we assume that material shares are constant. With nonconstant material shares, firms' unit costs would be determined from the micro data on domestic shares and material shares. We discuss this case in more detail in online Appendix O1. Empirically, the dispersion in domestic shares exceeds the one of material shares. Focusing on the sample of importing firms, the average interquartile range of material spending shares (domestic expenditure shares) within 2-digit industries is 0.25 (0.42). The average difference between the ninetieth and tenth percentile is 0.46 (0.71).

TABLE 2—ESTIMATING THE ELASTICITY OF SUBSTITUTION

	First-stage	Reduced form estimates:		ε	N
		$\Delta \ln$ productivity	$\Delta \ln$ value added		
<i>Full sample</i>					
All weights	-0.019 (0.003)	0.014 (0.004)	0.050 (0.005)	2.378 (0.523)	526,687
Pre-sample weights	-0.017 (0.004)	0.024 (0.004)	0.068 (0.006)	1.711 (0.166)	443,954
<i>Importers</i>					
All weights	-0.010 (0.004)	0.005 (0.004)	0.030 (0.006)	2.322 (1.014)	65,799
Pre-sample weights	-0.010 (0.005)	0.009 (0.004)	0.033 (0.006)	1.892 (0.541)	54,604

Notes: Robust standard errors are in parentheses. The table contains the results of estimating (26) with the instrument given in (25). We employ estimates of γ_s and σ_s based on (22) and (20), which are contained in Table 1. We use data for the years 2002–2006. The pre-sample period is 2001. In column 1, we report the first-stage relationship between our instrument and the changes in firms' domestic expenditure shares. The F -statistic for the main specification is 10.5. Columns 2 and 3 show that the instrument is correlated with two measures of firm performance: productivity (defined in (24)) and value added. Column 4 reports the implied value of ε as per (26). In the top panel, we include all firms; in the bottom panel, we only focus on the set of importing firms. In rows 1 and 3, we calculate firms' pre-sample import shares in the first year they appear as importers. In rows 2 and 4, we take the first year in our data as the pre-sample period for all firms and hence only include firms who were already active in that year. We retrieve $\hat{\varepsilon}$ from (26) using the delta method. The obtained estimator is convergent and asymptotically Gaussian. Because $\hat{\gamma}_s$ and $\widehat{1/(1-\varepsilon)}$ are estimated in separate regressions, we estimate the standard error associated with $\hat{\varepsilon}$ using a bootstrap procedure with 200 replications.

where Φ contains general equilibrium variables, $\tilde{\rho} = (1 - \gamma)(\sigma - 1)/\sigma$, $\tilde{\gamma} = \gamma(\sigma - 1)/\sigma$, and

$$(24) \quad \ln(\vartheta_i) = \frac{1}{1-\varepsilon} \tilde{\gamma} \ln(s_{Di}) + \frac{\sigma-1}{\sigma} \ln(\varphi_i).$$

Given estimates for γ and σ , we can use (23) to recover the productivity residual $\ln(\vartheta_i)$ up to a constant. We can then use (24) to estimate ε from the variation in changes in firms' domestic expenditure shares.

Clearly, we cannot estimate (24) via OLS as the required orthogonality restriction fails: most models of import behavior predict that changes in firms' domestic share are correlated with changes in firm efficiency φ_i . Hence, we employ an instrumental variable strategy. In particular, we follow Hummels et al. (2014) and instrument s_D with shocks to world export supplies, which we construct from the Comtrade data. More precisely, we construct the instrument

$$(25) \quad z_{it} = \Delta \ln \left(\sum_{ck} WES_{ckt} \times s_{cki}^{pre} \right),$$

where WES_{ckt} denotes aggregate exports of product k from country c in year t to the entire world excluding France, s_{cki}^{pre} is firm i 's import share of product k from country 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 induce exogenous variation

in firms' domestic shares as long as changes in firm efficiency $\tilde{\varphi}$ are uncorrelated with changes in aggregate exports of the countries in the firm's initial sourcing set. Intuitively, if we see China's global 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. Hence, we estimate ε from the second-stage regression

$$(26) \quad \Delta \ln(\hat{\vartheta}_{ist}) = \delta_s + \delta_t + \frac{1}{1-\varepsilon} \times \Delta \tilde{\gamma}_s \widehat{\ln(s_{ist}^D)} + u_{ist},$$

where δ are sector and year fixed effects, $\Delta \ln(\hat{\vartheta}_{ist})$ is the change in firm residual productivity, and $\Delta \tilde{\gamma}_s \widehat{\ln(s_{ist}^D)}$ is the change in domestic shares, which is instrumented by (25).

We implement this procedure in the following way. First, we augment the production function to include capital, i.e., we consider $y_{is} = \varphi_i l^{\phi_{ls}} x^{\gamma_s} k^{1-\phi_{ls}-\gamma_s}$.¹⁷ Second, we measure all parameters ϕ_{ls} , γ_s , and σ_s at the 2-digit sectoral level. The estimated material elasticities γ_s are reported in column 3 of Table 1. To construct our instrument (25), 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} . Finally, to increase the power of the estimation, we estimate (26) by pooling all firms from all sectors and estimate a single ε .

Table 2 contains the results. In the first column, we show the first-stage relationship between changes in world export supply z_{it} and firms' changes in domestic spending shares. 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 their domestic spending shares. Columns 2 and 3 show the reduced form results of regressing performance measures such as changes in residual productivity (defined in (24)) or log value added on the export supply shock z_{it} . As expected, there is a positive correlation. Column 4 then contains the results for ε . In our preferred specification, which does not condition on import status and where we calculate firms' initial expenditure shares s_{cki}^{pre} from the year before they start importing, we estimate this elasticity to be 2.38. If we keep the year used for the pre-sample weights s_{cki}^{pre} fixed for all firms, the implied elasticity is lower. When we restrict our sample to continuing importers, these point estimates are essentially unchanged. Note, however, that the standard errors increase substantially as we lose a large amount of data by conditioning on import status.¹⁸

¹⁷ Naturally, we include capital in our measure of total costs in (20) and calculate $Costs_i = wl_i + m_i + Rk_i$, where we assume a rental rate of 20 percent. We redid our analysis for a rental rate of 10 percent with quantitatively very similar results. These are available upon request. Similarly, we augment (23) to include capital and infer the labor elasticity ϕ_l from the optimality condition $m_i/wl_i = \gamma/\phi_l$.

¹⁸ Our estimates are close to the ones of Antràs, Fort, and Tintelnot (2017) who rely on cross-country variation. Kasahara 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, Koren, and Szeidl (2015) use Hungarian data and derive a production function equation analog to (23), as well as an import demand equation. They also find an elasticity of substitution of 4. The main difference with our approach is that they obtain the parameters of their structural model by *simultaneously* estimating the production function and import demand equations. In contrast, we identify ε mainly by using exogenous variation in input supplies.

As a complementary approach to identify γ and ε , we also consider methods to structurally estimate production functions—see, e.g., Levinsohn and Petrin (2003); De Loecker (2011); and Akerberg, Caves, and Frazer (2015), who build on the seminal work by Olley and Pakes (1996). A description of this approach and its application is contained in Appendix B. We estimate ε_s by treating the domestic share as an additional input in a production function estimation exercise. In contrast to the factor shares approach, we estimate all parameters simultaneously and allow ε to vary by sector. For the majority of industries, the point estimates are precisely estimated and in the same ballpark as the pooled estimate from the factor shares approach. For a few other industries, we lack precision and cannot reject that ε is below one. The existence of non-importers, however, implies that ε has to exceed unity. We therefore take the estimate stemming from the factor shares approach, i.e., $\varepsilon = 2.38$, as the benchmark for the remainder of the paper. While we lock in to this particular point estimate, we report confidence intervals for all quantitative results which take the sampling variation in this benchmark estimate into account. Because the estimated elasticities from the production function approach are within the sampling variation of the factor share estimate, our quantitative results will also be informative for these estimates.

C. Consumer Prices in Autarky

With the structural parameters at hand, we now quantify the gains from input trade in France using Proposition 1. We focus on one particular observed shock—a hypothetical reversal to input autarky, which can be directly analyzed using the cross-sectional data on firm-size and domestic shares displayed in Figures 1 and 2. The crucial ingredient to quantify the aggregate gains from input trade is the distribution of unit cost changes in the population of firms, which are simply given by $(\gamma_s/(1 - \varepsilon)) \ln(s_D)$ (see (7)). We depict this distribution in Figure 3 and summarize it in Table 3. We see that there is substantial heterogeneity across firms. While the median firm would see its unit cost increase by 11.2 percent if moved to autarky, firms above the ninetieth percentile of the distribution would experience losses of 85 percent or more.¹⁹

We now turn to Proposition 1 to aggregate these firm-level cost changes to the change in aggregate consumer prices. Panel A in Table 4 contains the results. We find that the price level in the manufacturing sector would be 27.5 percent higher if French producers were forced to source their inputs domestically. When the non-manufacturing sector is taken into account, the consumer price gains amount to 9 percent. The reason why these economy-wide gains are substantially smaller is that the non-manufacturing sector (which in our setting is assumed to be closed to

¹⁹ This heterogeneity is partly systematic in that bigger firms and exporters see higher gains. When we condition on import status, the positive relation between firm size and the firm-level gains essentially disappears. This is consistent with the pattern documented in Figure 2. See online Appendix O3 for details.

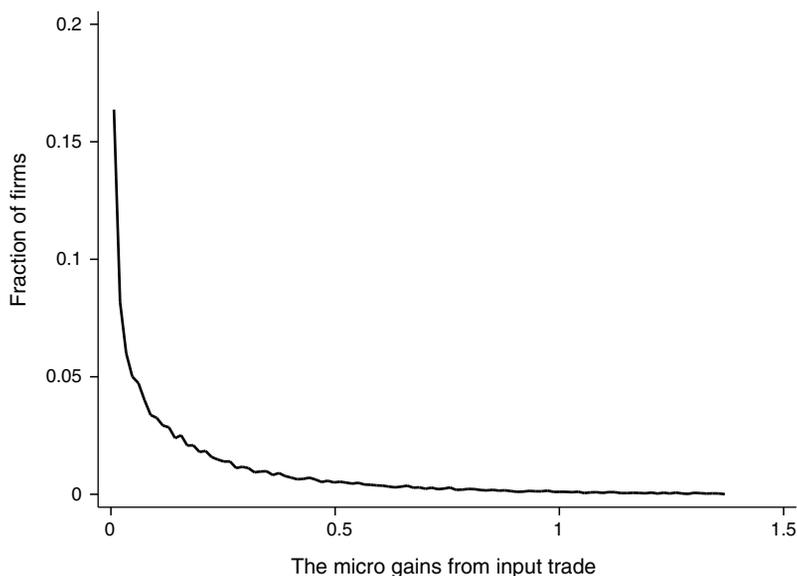


FIGURE 3. THE FIRM-LEVEL GAINS FROM INPUT TRADE IN FRANCE

Notes: The figure reports the empirical distribution of the firm-level gains from input trade relative to autarky, i.e., $(s_{Di}^{\gamma_s/(1-\varepsilon)} - 1) \times 100$ —see (7). 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 Tables 1 and 2.

TABLE 3—MOMENTS OF THE DISTRIBUTION OF THE FIRM-LEVEL GAINS IN FRANCE

Mean	Quantiles				
	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 firm-level gains from input trade relative to autarky shown in Figure 3, i.e., $(s_{Di}^{\gamma_s/(1-\varepsilon)} - 1) \times 100$ —see (7). 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 Tables 1 and 2.

trade) experiences only a 3 percent reduction in prices but accounts for 70 percent of consumers' budget (see Table 1).^{20,21}

To quantify the importance of firm heterogeneity, we also report the consumer price gains predicted by an aggregate approach that only uses data on domestic

²⁰ Allowing for variable markups, as in Edmond, Midrigan, and Xu (2015) and Atkeson and Burstein (2008), does not affect this estimate substantially. As we show in Appendix C, we find that consumer prices would be 28.7 to 28.9 percent higher in autarky, when markups are allowed to respond to input trade.

²¹ Note also that Proposition 1 infers unobserved physical productivity φ from data on firm value added. This procedure is valid if value added is only generated domestically or if firms do not vary in their export intensity. As we discuss in online Appendix O4, when export participation is a function of firm productivity, the formula in Proposition 1 continues to apply once the weights ω are calculated with domestic sales only. In this case, the economy-wide gains from input trade would amount to 8.2 percent and the gains in the manufacturing sector would be 24.4 percent. Finally, we note that these results are robust to using different weighting schemes. We redid the analysis weighting firms by employment and sales, instead of value added, and found very similar results. These are available upon request.

TABLE 4—THE GAINS FROM INPUT TRADE IN FRANCE

	Manufacturing sector		Entire economy					
	Sectoral price gains	Price of domestic inputs	Direct price reductions	Aggregate data				
<i>Panel A. Aggregate results</i>								
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]				
<i>Panel B. Sectoral results</i>								
Mining	7.8	[5.2, 10.3]	14.9	[11.1, 19.2]	3.0	[1.8, 4.2]	2.5	[1.6, 3.6]
Food, tobacco, beverages	17.8	[12.4, 23.4]	8.4	[6.2, 10.6]	11.1	[7.5, 14.6]	12.6	[7.8, 18.2]
Textiles and leather	55.6	[42.4, 74]	31.4	[24.3, 40.3]	31.1	[24.2, 39.9]	31.9	[22.4, 46.9]
Wood and wood products	14.4	[11.1, 18.2]	9.6	[7.4, 12.1]	8.2	[6.4, 10.5]	9.6	[6.7, 13.7]
Paper, printing, publishing	20.1	[14.7, 26.5]	14.5	[10.9, 18.7]	12.2	[9, 16]	11.0	[7.7, 15.4]
Chemicals	45.1	[32.7, 60.7]	21.6	[16.1, 28.2]	27.2	[20.1, 36.4]	28.1	[18.7, 41.8]
Rubber and plastic products	38.4	[27.5, 50.9]	27.3	[20.2, 36]	20.1	[14.3, 26.5]	21.5	[13.9, 31]
Non-metallic mineral products	20.8	[15.3, 27.4]	12.7	[9.7, 16.3]	13.4	[9.6, 17.9]	13.3	[9, 19]
Basic metals	38.9	[28.2, 50.2]	21.5	[16.4, 27.3]	21.8	[16.3, 27.7]	28.8	[19, 41.6]
Metal products	18.3	[13.8, 23.5]	20.5	[15.5, 26.2]	8.2	[6.2, 10.5]	7.7	[5.5, 10.8]
Machinery and equipment	31.7	[23, 41.6]	20.0	[15, 25.7]	17.6	[12.8, 23.2]	18.2	[12.2, 26.2]
Office and computing machinery	44.6	[31.9, 57]	25.2	[18.3, 32.1]	20.4	[15.4, 25.5]	37.0	[22.4, 60.3]
Electrical machinery	36.1	[26.4, 46.6]	23.9	[17.7, 30.6]	19.8	[14.6, 25.6]	21.6	[14.8, 30.7]
Radio and communication	38.5	[23.5, 54.8]	23.3	[16.6, 30.5]	21.5	[13.1, 31.1]	22.1	[12.5, 36.1]
Medical and optical instruments	29.2	[21.1, 38.3]	20.4	[15.1, 26.2]	17.9	[12.8, 23.4]	15.9	[10.7, 22.5]
Motor vehicles, trailers	23.3	[17.4, 39]	21.7	[17, 29.3]	6.2	[3.2, 16.4]	11.2	[6.1, 24.3]
Transport equipment	22.9	[16, 33.2]	19.9	[14.5, 27.2]	15.3	[10.5, 22]	11.8	[7.9, 18.2]
Recycling, nec.	26.0	[19.2, 33.4]	19.0	[14.5, 24]	12.9	[9.7, 16.3]	14.1	[9.5, 20.4]
Non-manufacturing	3.0	[2.3, 3.8]	7.5	[5.7, 9.4]	0.0	[0, 0]	0.0	[0, 0]

Notes: Panel A 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 1 where the associated Λ_s^{Aut} are reported in panel B and the structural parameters Ξ , γ_s , σ_s , and α_s are 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 1 where the sectoral gains are measured by $\Lambda_{Agg,s}^{Aut}$ as per (16) instead of Λ_s^{Aut} . The third row reports the bias, defined as the difference between the first two rows—see (17). Panel B contains the sectoral results. The first column contains the full change in sectoral prices relative to autarky, $(P_s^{Aut}/P_s - 1) \times 100$. The second column reports the reductions in the price of domestically sourced intermediate inputs, $(P_{Ds}^{Aut}/P_{Ds} - 1) \times 100$. The third column reports the direct price reductions from international sourcing relative to autarky, $(\exp(\Lambda_s^{Aut}) - 1) \times 100$, which are calculated according to (15). Column 4 reports the direct price reductions predicted by an aggregate approach, $(\exp(\Lambda_{Agg,s}^{Aut}) - 1) \times 100$, as per (16). 90–10 confidence intervals are reported in brackets for all measures. These are calculated via a bootstrap procedure, which we describe in online Appendix O5. The empirical distributions of all statistics are estimated using 200 bootstrap iterations.

spending shares at the sector level. This aggregate approach implies gains of 31.4 percent and 9.9 percent in the manufacturing sector and the entire economy, respectively. Ignoring the heterogeneity in firms' import behavior within sectors therefore results in an overestimation 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 (17)–(18).

To assess our confidence in these estimates, we calculate 90–10 confidence intervals of the bootstrap distribution.²² These are reported in squared brackets in Table 4. Note that the uncertainty in the point estimates stems from two sources.

²² We explain the details of the bootstrap procedure in online Appendix O5. 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 IIIB and obtain new estimates for the structural parameters. Finally, for each iteration, we recalculate the consumer price gains and the other statistics of interest.

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. These two forces induce quite a bit of variation in the estimates, with the majority of the variation stemming from the estimation of ε . With 80 percent probability, the consumer price gains in the manufacturing sector lie between 21 percent and 36 percent and the gains for the entire economy lie between 7 percent and 12 percent. We also find that the aggregate approach yields more uncertain estimates (second row) and leads to an overestimation of the gains with 80 percent probability (third row).

In panel B of Table 4, we provide additional details of the gains from input trade. More specifically, we report the gains by sector and provide a decomposition to isolate the importance of production linkages across sectors. In column 1, we 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. For example, while prices for textile products would be 56 percent higher if producers were not allowed to source their inputs from abroad, this effect is only 18 percent for metal products. Columns 2 and 3 decompose these price changes into the indirect gains stemming from firms in sector s buying domestic inputs from other firms who in turn may engage in input trade, $p_{D_s}^{Aut}/p_{D_s}$, and the direct price reduction from firms in sector s sourcing internationally, Λ_s . We find that interlinkages are important as they account for roughly 50 percent 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. To assess the importance of such heterogeneous interconnections, consider the case without cross-industry input-output linkages, where each sector uses only its own products as inputs. In this case, we find a point estimate for the consumer prices gains from trade of 12 percent. Compared to the actual gains of 9 percent, the economy without interlinkages overestimates 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 in the model, such linkages actually dampen the aggregate effect of input trade.

Finally, panel B of Table 4 also contains the direct price reductions that arise from an aggregate model, Λ_s^{Agg} , at the sectoral level. 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 (18) which characterizes the sign of the bias as a function of parameters encapsulated in χ_s . It turns out that for most sectors our estimated parameters imply that $\chi_s > 1$ so that the aggregate models are upward biased. 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 percent, the exact firm-based formula tells us that this number should be only 20 percent. Taken together, our results suggest that the heterogeneity in firms' domestic expenditure shares is important to credibly quantify the gains from input trade.

D. Beyond Autarky: A Shock to Import Prices

We now extend the analysis to shocks that make all foreign varieties more expensive without leading the economy into autarky. This is, for example, the case under a currency devaluation (see, e.g., Gopinath and Neiman 2014). While the effect of the shock on consumer prices is still fully determined from the changes in firms' domestic shares, such changes are no longer observed and one needs a model to predict them. Hence, we now specify additional components of the theory. Doing so also allows us to quantify the effect of input trade on welfare, taking into account the resources (if any) spent by firms to attain their sourcing strategies. The main goal of this section is to assess whether the *observable* micro data is important to quantify the effects of counterfactuals which—by construction—are *unobserved*. To do so, we consider different models of importing which vary in the extent to which they match the micro data and compare their counterfactual implications.

A Model of Import Behavior.—To construct a model of firms' domestic shares, we start from the general framework laid out in Section I and impose restrictions.²³ For brevity, we provide all derivations in online Appendix O8. We consider settings where firms' extensive margin is limited through the presence of fixed costs so that firms choose their sourcing strategy by trading off the import-induced reduction in unit costs versus the payment of fixed costs. While this seems a natural starting point, one could extend the analysis to other models of the extensive margin.

We assume that the fixed cost of sourcing from a given country is constant across countries but potentially varies across firms. In this case, firms select their sourcing countries based purely on their price-adjusted qualities and the sourcing strategy reduces from a set \mathcal{L} to a scalar, a price-adjusted quality cutoff.²⁴ We also impose the following functional form assumptions: (i) the import bundle x_I takes a CES form with elasticity of substitution κ and (ii) country quality q_c is Pareto distributed, i.e., $\Pr(q_c \leq q) = 1 - (q_{\min}/q)^\theta$, where $\theta > \min[1, \kappa - 1]$. Setting the price of imported inputs, p_c , to unity, these assumptions imply that the import price index only depends on the mass of countries sourced from and takes a convenient power form:

$$(27) \quad A(\mathcal{L}) = \left(\int_{c \in \mathcal{L}} q_c^{-(1-\kappa)} dc \right)^{\frac{1}{1-\kappa}} \\ = \frac{1}{q_{\min}} \left(\frac{\theta}{\theta - (\kappa - 1)} \right)^{\frac{1}{1-\kappa}} n^{-\left(\frac{1}{\kappa-1} + \frac{1}{\theta}\right)} \equiv zn^{-\eta} = A(n).$$

Here n is the share of countries the firm sources foreign inputs from and z and η are auxiliary parameters which depend on the parameters governing the distribution

²³ For expositional simplicity, we consider a one-sector version of the model. See the online Appendix for the analysis with multiple sectors.

²⁴ More precisely, as long as fixed costs are constant across countries, if country c with price-adjusted quality q_c/p_c is an element of \mathcal{L} so are all countries c' with $q_{c'}/p_{c'} > q_c/p_c$. Computing firms' optimal sourcing strategies can be challenging when prices, qualities and fixed costs vary by country in an arbitrary way—see Antràs, Fort, and Tintelnot (2017). Allowing for country-specific fixed costs will only matter for normative questions as long as it translates into a different predicted distribution of domestic shares.

of quality (q_{\min}, θ) and the elasticity of substitution across foreign varieties κ . In particular, z parametrizes the average quality-adjusted price of foreign inputs. We will consider changes in z as our counterfactual shock. While the reversal to autarky considered above corresponds to $z \rightarrow \infty$, we are now able to study finite increases in z .

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

$$(28) \quad \pi = \max_n \left\{ u(n)^{1-\sigma} \times B - w(nf + f_I \mathbf{1}(n > 0)) \right\},$$

where f denotes the fixed cost per country, f_I is an additional fixed cost to start importing, and $\mathbf{1}(\cdot)$ is an indicator of import status. Finally, B is defined as $B \equiv (1/\sigma)(\sigma/(\sigma-1))^{1-\sigma} P^{\sigma-1} S$, with P and S denoting the consumer price index and aggregate spending, which are determined in general equilibrium. The unit cost function $u(\cdot)$ is given by the analogues of equations (5) and (6), which we replicate here for convenience:

$$(29) \quad u_i(n) \equiv \frac{1}{\varphi_i} w^{1-\gamma} \left(\frac{P_D}{Q_D} \right)^\gamma s_{Di}(n) \frac{\gamma}{\varepsilon-1},$$

$$(30) \quad s_{Di}(n) = \left(1 + \left(\frac{1-\beta_i}{\beta_i} \right)^\varepsilon \left(\left(\frac{P_D}{Q_D} \right) \frac{1}{z} n^\eta \right)^{\varepsilon-1} \right)^{-1}.$$

While (29) shows that the effect of input trade on firms' unit costs is fully summarized by the domestic share, (30) now contains *a theory of domestic shares*: these can be small either because the firm sources from many countries (n is large) or because of a technological bias toward foreign inputs (β is low).

Equations (28)–(30) fully describe firms' optimal import behavior.²⁵ To close the model in general equilibrium, we impose equilibrium in the labor market and balanced trade between the domestic economy and the rest of the world. In particular, we assume that foreigners demand the output of local firms with the same CES demand structure as domestic consumers and producers and that the supply of foreign inputs from country c is perfectly elastic. Letting y_i^{ROW} be the foreign demand for firm i 's production, balanced trade requires that

$$(31) \quad \int_i p_i y_i^{ROW} di = \int_i (1 - s_{Di}) m_i di,$$

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. An equilibrium is attained when firms maximize profits, consumers maximize utility, trade is balanced and labor and good markets clear.

In this context, it can be shown that the equilibrium change in consumer welfare relative to autarky is given by

$$(32) \quad \frac{W}{W^{Aut}} = \frac{P^{Aut}}{P} \times \left(\frac{L - \int_i l_{\mathcal{I}_i} di}{L} \right),$$

²⁵ In online Appendix O8, 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 .

where W denotes consumer welfare and $\int_i l_{\varphi_i} di$ denotes the aggregate resource loss due to fixed costs. Hence, the welfare gains from input trade consist of two components. First, there is the reduction in consumer prices associated with input trade. This is the main focus of our paper. Second, there is the resource loss due to fixed costs, which results in (weakly) fewer workers left for production. Because this second term in (32) is weakly smaller than unity, the change in the consumer price index provides an upper bound for the change in welfare in the class of models of Section I. While we calculate $\int_i l_{\varphi_i} di$ within a model of fixed costs, we note that (32) is consistent with any extensive margin mechanism. For example, if importers found their trading partners through a process of network formation, (32) would still hold but the environment to calculate $\int_i l_{\varphi_i} di$ would be different.

Dimensions of Firm-Level Heterogeneity.—We now calibrate this model to the French micro data. In order to generate the rich distribution of domestic shares and value added shown in Figures 1 and 2, we have to allow for (at least) two sources of firm heterogeneity. As is standard, we allow firms to differ in efficiency φ_i . For the second source of heterogeneity, we consider two options. Our first option is a model with *heterogeneous fixed costs* where firms differ in their f_i . For simplicity, we assume that the fixed cost to start importing f_I is constant across firms. Second, we consider a model with *heterogeneous home bias* where firms differ in their β_i . In this model, we assume there are no fixed costs of importing per country ($f = 0$), but still assume a positive fixed cost to start importing ($f_I > 0$) to match the existence of non-importing firms.

As φ_i and the endogenous unit costs reduction through input trade are complements, there is a firm-specific efficiency cutoff, either $\bar{\varphi}(f_i)$ or $\bar{\varphi}(\beta_i)$, above which firms select into importing. This sorting generates overlap in the size distribution of importers and non-importers as seen in Figure 2. Furthermore, both models generate variation in import intensity conditional on size. While the heterogeneous fixed-cost model generates dispersion in import shares fully via variation in the extensive margin n_i , the home bias model is the polar opposite in that firms gain differentially from international trade because of variation in β_i .

In order to calibrate these parameterizations of the model to the data, we adopt the following strategy. First, we use the estimates of ε , γ , and σ from Section IIIB above.²⁶ Next, for the model with heterogeneous fixed costs, we need to estimate η , which determines the price index of the import bundle and hence the demand for foreign varieties (see (27)). We estimate η directly from the micro data and identify it from the cross-sectional relationship between firms' extensive margin of trade and their domestic shares.²⁷ Without loss of generality, we can normalize the quality of the domestic variety (q_D) to unity.

²⁶ Section IIIB provides estimates of σ and γ by sector. In this section, we use value-added weighted averages of these sectoral estimates, which yield $\sigma = 3.83$ and $\gamma = 0.61$.

²⁷ In particular, (30) predicts a log-linear relation between n and $(1 - s_D)/s_D$, with a slope given by η . See online Appendix O9 for details and results. Our preferred specification yields a value of η of 0.382 that is precisely estimated. Note also that we do not require η for the heterogeneous bias model as all firms decide to source from all countries (conditional on importing), i.e., $n = 1$, see (30).

Finally, we parametrize the distributions of firm heterogeneity. For efficiency, we take a log-normal distribution with variance σ_φ^2 . We normalize mean efficiency to unity. For the heterogeneous fixed-cost model, we parametrize the *conditional* distribution of fixed costs also as log-normal and denote the mean of log fixed costs by μ_f , their variance by σ_f^2 , and their correlation with efficiency by $\rho_{f\varphi}$. Similarly, we assume that the degree of home bias, $\tilde{\beta} \equiv \beta/(1 - \beta) \in [0, \infty]$, is conditionally log-normally distributed with mean $\mu_{\tilde{\beta}}$, variance $\sigma_{\tilde{\beta}}^2$, and correlation with efficiency $\rho_{\tilde{\beta}\varphi}$.²⁸

Calibration.—Our calibration strategy is as follows. The distributions of firm heterogeneity are parametrized by four parameters. For the model with heterogeneous fixed costs (resp. home bias), such parameters control the dispersion in efficiency, the dispersion in fixed costs (resp. home bias), the average fixed cost (resp. home bias) and the correlation of fixed costs (resp. home bias) with efficiency. For each model, we calibrate these parameters by targeting salient features of the joint distribution of value added and domestic shares displayed in Figures 1 and 2. In particular, we match the aggregate domestic share, the dispersion in value added, the dispersion in domestic shares and their correlation with value added. Finally, we also need to calibrate the fixed cost to start importing f_I and to do so we target the share of non-importing firms.

To assess the value of the micro data, we also calibrate the above models *without* targeting the moments associated with the heterogeneity in domestic shares, i.e., their dispersion and correlation with firm size. As we drop these two moments, we also drop two parameters in each model. In the heterogeneous fixed costs model, we set $\sigma_f = \rho_{f\varphi} = 0$, which corresponds to assuming constant fixed costs across firms as in Gopinath and Neiman (2014) and Ramanarayanan (2017). We call this parametrization the *homogeneous fixed cost* model. In the heterogeneous home-bias model, we set $\sigma_{\tilde{\beta}} = \rho_{\tilde{\beta}\varphi} = 0$, which corresponds to a *homogeneous home bias model*. Finally, we consider a model with no fixed costs of any kind, $f_i = f_I = 0$, and a constant home bias. This version of the model implies that firms' import intensities are equalized and aggregate statistics are sufficient. We therefore refer to it as the *aggregate model*.

Table 5 summarizes the five parameterizations of the model we consider and contains the calibration results. In panel A, we report the calibrated parameters. Panel B contains the model-generated moments as well as the targeted ones in bold letters. We first note that all versions of the model match the targeted moments *exactly*. As expected, the aggregate model (column 5) generates full participation in import markets and equalized domestic shares. The homogeneous fixed cost and home-bias models (columns 3 and 4) improve on these margins by allowing for fixed costs. However, because they feature efficiency as the single source of firm heterogeneity, these models predict too strong a correlation between firm-size and domestic shares relative to the data, as well as no overlap in the size distribution of importers and

²⁸ Formally, we assume that $\ln(\varphi) \sim \mathcal{N}(-(1/2)\sigma_\varphi^2, \sigma_\varphi^2)$ and parametrize the conditional distribution of fixed costs and home bias as $\ln(f)|_{\ln(\varphi)} \sim \mathcal{N}(a_0 + a_\varphi \ln(\varphi), \sigma_{f|\varphi}^2)$ and $\ln(\tilde{\beta})|_{\ln(\varphi)} \sim \mathcal{N}(b_0 + b_\varphi \ln(\varphi), \sigma_{\tilde{\beta}|\varphi}^2)$.

TABLE 5—CALIBRATING MODELS OF IMPORTING: THE VALUE OF THE MICRO DATA

		Firm-based models				Aggregate model
		Heterogeneous fixed costs	Heterogeneous home bias	Homogeneous fixed costs	Homogeneous home bias	
<i>Panel A. Structural parameters</i>						
Dispersion in efficiency	σ_φ	0.528	0.528	0.513	0.496	0.556
Fixed cost of importing	f_I	0.035	0.058	0.047	0.562	—
Average home bias	μ_{β}	1 [†]	2.597	1 [†]	1.193	1.284
Dispersion in home bias	σ_{β}	—	1.028	—	0	—
Correlation of home bias and efficiency	$\rho_{\beta\varphi}$	—	0.124	—	0	—
Average fixed cost	μ_f	5.061	—	5.475	—	—
Dispersion in fixed costs	σ_f	2.374	—	0	—	—
Correlation of fixed cost and efficiency	$\rho_{f\varphi}$	0.739	—	0	—	—
	Data					
<i>Panel B. Moments</i>						
Aggregate domestic share	0.720	0.720	0.720	0.720	0.720	0.720
Dispersion in ln VA	1.520	1.520	1.520	1.520	1.520	1.520
Share of importers	0.199	0.199	0.199	0.200	0.199	1.000
Dispersion in ln s_D	0.360	0.360	0.360	0.137	0.179	0.000
Correlation of lnVA and ln s_D	-0.310	-0.310	-0.310	-0.720	-0.768	0.000
<i>Panel C. Reversal to autarky</i>						
Percentage change in consumer prices	$\frac{P^{Aut} - P}{P}$	37.87	38.01	43.09	43.89	44.73
	Bias			13.78	15.90	18.10
Percentage change in welfare	$\frac{W - W^{Aut}}{W}$	17.43	36.42	21.59	27.81	44.73
<i>Panel D. Increase in import prices</i>						
Change in aggregate import share by . . .						
... 5%	$\frac{P' - P}{P}$, pct.	1.85	1.87	2.08	2.15	2.19
	Pct. diff.		0.79	12.47	16.05	18.21
... 10%	$\frac{P' - P}{P}$, pct.	3.71	3.73	4.17	4.30	4.39
	Pct. diff.		0.67	12.58	16.08	18.31
... 20%	$\frac{P' - P}{P}$, pct.	7.42	7.47	8.37	8.63	8.80
	Pct. diff.		0.67	12.86	16.31	18.55

Notes: In panel A, we report the calibrated structural parameters for the respective models. In the models of columns 1 and 3, we normalize the level of the home bias to unity without loss of generality, and we denote this normalization by “†”. Panel B contains the moments. We report both the moments observed in the data and generated by the models. The moments that the respective models are calibrated to are displayed in bold figures. The number of calibrated moments equals the number of parameters for the respective model. In panel C, we report the results from a reversal to autarky. We report the change in consumer prices (row 1) and the change in welfare (row 3). We also report the bias of the change in consumer prices relative to the model with heterogeneous fixed costs (row 2). We calculate this number as the percentage difference between the respective change in consumer prices. In panel D, we report the effects of a shock that increases the prices of all foreign varieties to reduce the aggregate import share by 5, 10, and 20 percent. We report the implied change in consumer prices ($(P' - P)/P$ in rows 1, 3, and 5) and the difference relative to the model with heterogeneous fixed costs (Pct. diff. in rows 2, 4, and 6). See online Appendix O10 for details of the computational procedure.

domestic firms. By allowing for an additional dimension of heterogeneity, the heterogeneous fixed costs or home-bias models (columns 1 and 2) improve the fit along these dimensions. First, they increase the dispersion in domestic shares by introducing variation in import demand conditional on efficiency. Second, they reduce the correlation between size and domestic shares. Intuitively, to be consistent

with the low correlation of size and import intensity both parameterizations require that some efficient firms have a *lower* incentive to import compared to a model with a single source heterogeneity. This is achieved by having a positive correlation between firm efficiency and fixed costs (resp. home bias).

Results.—With the calibrated models at hand, we can now study the effect of any shock to the trading environment on both consumer prices and welfare. We focus on two counterfactuals: (i) a reversal to input autarky ($z \rightarrow \infty$) and (ii) an increase in the relative price of all foreign inputs. More precisely, the latter exercise corresponds to increasing z to attain a decrease in the aggregate import share of 5 percent, 10 percent, or 20 percent. Table 5 contains the results, from which we draw three main conclusions.

First, we find that the two models that match the micro data on size and domestic shares predict the same counterfactual change in consumer prices. To see this, consider the two models in columns 1 and 2. While both models perfectly match the four moments of the joint distribution of value added and domestic shares, their underlying micro-structure is very different. They nevertheless give very similar predictions for the change in consumer prices across the different counterfactuals. That this result is exact for a reversal to input autarky (panel C) is the content of Proposition 1: both models predict an increase in consumer prices of 38 percent.²⁹ Panel D shows that this is also the case for the non-autarky counterfactuals: the difference in the implied changes in consumer prices between the two models is less than 1 percent.

Second, the models that do *not* match the data on domestic shares and value added (columns 3–5) yield quantitatively meaningful biases. In panel C, we report the change in consumer prices relative to autarky: the three models predict changes that are 14 to 18 percent too high. That such biases are not confined to the autarky counterfactual is seen in panel D. The estimated effects of the increases in the price of foreign inputs are also upward biased by similar magnitudes. To understand why these biases are positive it is helpful to go back to our theoretical results. That the aggregate model in column 5 predicts the largest change in consumer prices in the autarky counterfactual follows from our characterization of the bias in (18): because $\gamma(\sigma - 1)/(\varepsilon - 1) > 1$, the aggregate model provides an upper bound for *any* model of importing. It is also intuitive that the models with homogeneous fixed costs and home bias are upward biased vis-à-vis the models that match the micro data on firm-size and domestic shares. By relying on efficiency as the single source of firm heterogeneity, the models in columns 3 and 4 generate a perfectly negative correlation between efficiency and the domestic share. This means that more efficient firms experience a larger reduction in their unit cost, a feature that tends to make input trade more attractive.

Finally, we also calculate the change in welfare taking the resource loss of fixed costs into account. We report the results for the reversal to autarky in the last row of

²⁹ This number does not coincide with that reported for the manufacturing sector in Table 4 above, i.e., 27.5 percent. The reason is that we calibrated a one-sector version of the model of Section I to moments obtained from pooling all industries. Additionally, we targeted only five moments of the joint distribution of size and domestic shares.

panel C. In contrast to the results for consumer prices, the implications for welfare can vary substantially across models, even conditional on fully matching the micro data. Specifically, columns 1 and 2 in panel C show that the heterogeneous fixed cost and home-bias models predict very different changes in welfare relative to input autarky. While the former predicts an increase of 17 percent in welfare, the latter predicts an increase of 36 percent.³⁰ Thus, the share of the consumer price gains that is lost by firms' attaining their sourcing strategies crucially depends on the underlying source of variation in firms' domestic shares.

Elasticity Bias.—In the above experiments, we keep the elasticity of substitution ε constant across models. In particular, we treat ε as a production function parameter and estimate it directly from the micro-data. In contrast, approaches that rely on aggregate data often discipline this parameter with the aggregate trade elasticity—see, e.g., Costinot and Rodríguez-Clare (2014). We do not target this moment in our calibration exercises of Table 5 and, as a result, the implied trade elasticity varies across models. In particular, the model with heterogeneous fixed costs generates an aggregate trade elasticity of 4.5, while the aggregate model features an elasticity of 1.38.³¹ Hence, while the elasticity of substitution ε and the aggregate trade elasticity coincide in aggregate Armington-style models (note that indeed $1.38 = \varepsilon - 1$), this is not the case in models with richer firm heterogeneity. For example, in the model with heterogeneous fixed costs, the additional extensive margins of adjustment directly affect the trade elasticity and dissociate it from the structural parameter ε . In our context, we find that the relatively low estimate for ε stemming from micro data is perfectly consistent with an aggregate trade elasticity between 4 and 5, which is close to the consensus estimates in the literature (see, e.g., Costinot and Rodríguez-Clare 2014; Edmond, Midrigan, and Xu 2015; or Simonovska and Waugh 2014a, b).

We explore an alternative approach where the aggregate trade elasticity is kept constant across models. In particular, for the aggregate model to generate an aggregate trade elasticity of 4.5 as in the heterogeneous fixed-cost model, the implied elasticity of substitution has to be equal to 5.5. With this higher value of ε , the aggregate model predicts an increase in consumer prices of 11 percent under input autarky, instead of the 45 percent predicted by the baseline aggregate model of Table 5. Relying on aggregate data therefore results in a substantial reduction in the predicted gains from trade via the estimated value of ε .³² This negative “elasticity

³⁰Note that this difference is not due to the fact that the home-bias model does not feature any fixed cost per sourcing country. The homogeneous home-bias model of column 4 does not feature any fixed costs per country either, but implies that the fixed costs to start importing account for about 40 percent of the consumer price gains. The reason why heterogeneity in the efficiency of using imported inputs generates a tighter bound between welfare and consumer prices is the discrepancy between the marginal importer, whose cost reductions determine the calibrated value of fixed costs, and the set of *inframarginal* firms, who might benefit from input trade substantially.

³¹Formally, we calculate the trade elasticity as $d \ln((1 - s_D)/s_D)/d \ln \tau$, where $d \ln \tau$ denotes the increase in iceberg trade costs (or relative foreign prices). Empirically, this elasticity is usually estimated from the cross-country variation in trade costs and import shares conditional on origin and destination fixed effects. We therefore hold the domestic price level p_D constant when calculating the elasticity in the model. The implied trade elasticities in the remaining models of Table 5 are between 1.5 and 3.5.

³²We conduct this exercise for the aggregate model only, as this model implies common domestic shares across firms and hence is inconsistent with the approach of Section IIIB, which estimates ε from variation in such shares.

bias” of aggregate models is in sharp contrast to the positive bias found in Table 5 holding ε constant and further highlights the importance of using micro-data to draw aggregate conclusions.

IV. Conclusion

Firms around the world routinely engage in input trade to reduce their costs of production, thereby benefiting domestic consumers through lower prices. Moreover, firms differ vastly in the intensity with which they participate in international markets. In this paper, we develop a methodology to measure how consumer prices are affected by input trade in environments which explicitly take the heterogeneity in import behavior into account.

Our main theoretical result is a sufficiency result that shows that the change in consumer prices due to changes in the import environment (e.g., a change in trade costs or a change in foreign prices) is fully determined from the joint distribution of firm size and changes in domestic expenditure shares. Importantly, a wide class of models used in the literature features the same predictions for consumer prices as long as they are calibrated to the same micro data. Approaches that abstract from firm heterogeneity altogether and rely on aggregate statistics give biased results. A focal point of our analysis is the case of a reversal to input autarky. As firms’ counterfactual domestic shares in autarky are equal to unity, the gains from input trade are fully determined from firms’ value added and domestic shares. In our application to France, we find that consumers would face 27 percent higher prices for manufacturing products under input autarky.

We then show quantitatively that this result extends to non-autarky counterfactuals such as an increase in the price of foreign inputs. In the context of a model with fixed costs, we find that parameterizations of the model that are calibrated to the micro data on firm size and domestic shares imply similar changes in consumer prices. Conversely, models that do not match this data give biased predictions. We conclude that the information contained in the joint distribution of firm-level domestic shares and size is crucial to quantify the effect of input trade on consumer prices in settings with heterogeneous importers.

APPENDIX A: PROOF OF PROPOSITION 1

The consumer price index associated with (8) is given by

$$(A1) \quad P = \prod_{s=1}^S \left(\frac{P_s}{\alpha_s} \right)^{\alpha_s} \quad \text{where} \quad P_s = \left(\int_0^{N_s} p_{is}^{1-\sigma_s} di \right)^{\frac{1}{1-\sigma_s}},$$

Of course, we re-calibrate all other parameters to still match all moments reported in Table 5. Similarly, we find smaller effects from the different experiments analyzed in panel D of Table 5. For example, an increase in foreign input prices that reduces the aggregate import share by 10 percent increases consumer prices by 2.6 percent instead of 4.4 percent. See online Appendix O6 for details.

and P_s is the price index for sector s . Using (A1), 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, $p_{is} = (\sigma_s/(\sigma_s - 1))u_{is}$. Using the expression for the firm's unit cost in terms of its domestic expenditure share in (6), we find that

$$(A2) \quad P_s = \frac{\sigma_s}{\sigma_s - 1} \left(\frac{1}{\gamma_s}\right)^{\gamma_s} \left(\frac{1}{1 - \gamma_s}\right)^{1 - \gamma_s} \left(\frac{p_{Ds}}{q_{Ds}}\right)^{\gamma_s} \left(\int_0^{N_s} (\tilde{\varphi}_i^{-1} (s_{Di})^{\frac{\gamma_s}{\varepsilon_s - 1}})^{1 - \sigma_s} di\right)^{\frac{1}{1 - \sigma_s}},$$

which is (10) in the main text. Given the aggregator in (9), the price index of the domestic bundle is given by

$$(A3) \quad p_{Ds} = \zeta_s^* \prod_{j=1}^S P_j^{\zeta_j^s} \quad \text{where} \quad \zeta_s^* \equiv \prod_{j=1}^S (\zeta_j^s)^{-\zeta_j^s}.$$

Note that (A2) implies

$$(A4) \quad \frac{P_s^{Aut}}{P_s} = \left(\frac{p_{Ds}^{Aut}}{p_{Ds}}\right)^{\gamma_s} \left(\frac{\int_0^{N_s} \tilde{\varphi}_i^{\sigma_s - 1} di}{\int_0^{N_s} (\tilde{\varphi}_i s_{Di})^{\frac{\gamma_s}{1 - \varepsilon_s}} di}\right)^{\frac{1}{1 - \sigma_s}}$$

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

where ω_i is firm i 's share in total value added in sector s , and the second equality follows from $va_i = \kappa_s \tilde{\varphi}_i^{\sigma_s - 1} s_{Di}^{\gamma_s(\sigma_s - 1)/(1 - \varepsilon_s)}$. With (A4) at hand, we can express the consumer price gains as

$$(A5) \quad G = \sum_{s=1}^S \gamma_s \alpha_s \pi_s + \sum_{s=1}^S \alpha_s \Lambda_s,$$

where $\pi_s \equiv \ln(p_{Ds}^{Aut}/p_{Ds})$ and Λ_s is given by (13) 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 (A3) and (A4) jointly imply

$$(A6) \quad \pi_s = \sum_{j=1}^S \zeta_j^s \gamma_j \pi_j + \sum_{j=1}^S \zeta_j^s \Lambda_j.$$

Equation (A6) describes a system of $S \times S$ equations that characterizes the equilibrium $[\pi_s]_{s=1}^S$. Letting $\boldsymbol{\pi} \equiv [\pi_1, \pi_2, \dots, \pi_S]$ be a column vector, we can express

the system in (A6) in matrix form as $\pi = \Xi\Gamma\pi + \Xi\Lambda$. Its solution is given by $\pi = (\mathbf{I} - \Xi\Gamma)^{-1}\Xi\Lambda$. Using (A5), the consumer price gains G are therefore given by

$$(A7) \quad G = \alpha'\Gamma\pi + \alpha'\Lambda = \alpha'\Gamma(\mathbf{I} - \Xi \times \Gamma)^{-1}\Xi\Lambda + \alpha'\Lambda.$$

For counterfactuals other than autarky, (A4) should be replaced by

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

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 (A7) where Λ_s is given by (13). This proves Proposition 1.

APPENDIX B: ADDITIONAL RESULTS FOR EMPIRICAL ANALYSIS

A. Data Description

Our main dataset stems from the information system of the French custom administration (DGDDI) and contains the majority of import and export flows by French manufacturing firms. 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.³³ However, some firms that are below the threshold (ca. 10,000 firm-year observations out of ca. 130,000) voluntarily report.³⁴

In spite of this limitation, the attractive feature of the French data is the presence of unique firm identifiers (the SIREN codes) that are 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

³³This threshold was in effect between 2001 and 2006, which is the 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.

³⁴The 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 33) to show that this concern is unlikely to severely affect our results. The reason is related to the weak relation between domestic expenditure shares and firm size shown in Figure 2.

TABLE B1—CHARACTERISTICS OF IMPORTERS, EXPORTERS, AND DOMESTIC 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: All amounts are expressed in thousand euros. Sales, wages, expenditures on imports or exports are 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"). A firm is member of an international group if at least one affiliate or the head-quarter is located outside of France.

fiscal files.³⁵ 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 LIFIDIANE (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 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 B1. 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.

³⁵The firm-level accounting information is retrieved from two different files: the Bénéfices Réels Normaux (BRN) and the Régime Simplifié d'Imposition (RSI). 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 datasets contain essentially the same information. Their union covers nearly the entire universe of French firms.

TABLE B2—PRODUCTION FUNCTION COEFFICIENT ESTIMATES BY 2-DIGIT SECTOR: FACTOR SHARES

Industry	ISIC	ϕ_k	ϕ_l	γ
Mining	10–14	0.374 (0.039)	0.293 (0.017)	0.333 (0.043)
Food, tobacco, beverages	15–16	0.098 (0.004)	0.177 (0.003)	0.725 (0.006)
Textiles and leather	17–19	0.081 (0.003)	0.293 (0.009)	0.626 (0.012)
Wood and wood products	20	0.113 (0.004)	0.285 (0.006)	0.602 (0.006)
Paper, printing, publishing	21–22	0.134 (0.007)	0.362 (0.011)	0.504 (0.011)
Chemicals	24	0.124 (0.008)	0.204 (0.01)	0.671 (0.014)
Rubber and plastic products	25	0.124 (0.005)	0.289 (0.007)	0.587 (0.011)
Non-metallic mineral products	26	0.178 (0.01)	0.294 (0.012)	0.529 (0.015)
Basic metals	27	0.124 (0.01)	0.202 (0.015)	0.674 (0.021)
Metal products	28	0.108 (0.002)	0.412 (0.008)	0.479 (0.009)
Machinery and equipment	29	0.071 (0.003)	0.313 (0.015)	0.616 (0.018)
Office and computing machinery	30	0.037 (0.012)	0.150 (0.032)	0.813 (0.04)
Electrical machinery	31	0.096 (0.008)	0.306 (0.011)	0.598 (0.014)
Radio and communication	32	0.055 (0.006)	0.322 (0.048)	0.624 (0.052)
Medical and optical instruments	33	0.071 (0.004)	0.435 (0.026)	0.494 (0.029)
Motor vehicles, trailers	34	0.106 (0.009)	0.135 (0.016)	0.759 (0.014)
Transport equipment	35	0.152 (0.019)	0.499 (0.03)	0.349 (0.044)
Recycling, nec.	36–37	0.084 (0.003)	0.283 (0.009)	0.633 (0.012)

Notes: The table contains the production function parameters based on observed factor shares. See the main text for details. Bootstrapped standard errors are in parentheses.

B. Estimation of Parameters

Table B2 contains the full results of the factor share approach, where we compute standard deviations with a bootstrap procedure (with 200 replications). Remember that this method imposes the assumption of constant returns, so that $\phi_{ks} + \phi_{ls} + \gamma_s = 1$.

We now describe a complementary approach to identify γ and ε based on methods to structurally estimate production functions—see, e.g., Levinsohn and Petrin (2003); De Loecker (2011); and Akerberg, Caves, and Frazer (2015), who build on

the seminal work by Olley and Pakes (1996). Our implementation and identification strategy closely follows De Loecker (2011).

Estimation.—In this section, we show how the setup laid out in Section III can be used to retrieve structural estimates of γ_s and ϵ_s , as an alternative approach to the baseline factor share approach of Section IIIB. First, we augment firms' production function to include physical capital and also explicitly introduce time subscripts to clarify the timing assumptions underlying our structural estimation. We therefore write our production function as

$$(B1) \quad y_{its} = \varphi_{it} \cdot k_{it}^{\phi_{ks}} \cdot l_{it}^{\phi_{ls}} \cdot x_{it}^{\gamma_s},$$

where ϕ_{ks} and ϕ_{ls} denote the capital and labor elasticities in sector s . We make the standard assumption that the productivity term φ_{it} can be decomposed into two components according to $\varphi_{it} = \varphi_{it}^{obs} \exp(\epsilon_{it})$, where φ_{it}^{obs} is observed by the firm and ϵ_{it} captures both measurement error and idiosyncratic shocks to production. We assume that ϵ_{it} is i.i.d. across firms and time and is independent of all other shocks. We also assume that log productivity follows a flexible AR(1) process

$$(B2) \quad \ln(\varphi_{it}^{obs}) = n(\varphi_{it-1}^{obs}) + \chi_{it},$$

where χ_{it} is an i.i.d. shock.

In equation (B1), we can again replace the (unobserved) physical quantities of intermediate inputs x_{it} with the (observed) level of spending in materials and express the appropriate deflator via the firm's domestic share (see (5)). This implies that firm output is given by

$$(B3) \quad y_{its} = \varphi_{it}^{obs} \cdot s_{Dit}^{-\frac{\gamma_s}{\epsilon_s-1}} \cdot k_{it}^{\phi_{ks}} \cdot l_{it}^{\phi_{ls}} \cdot m_{it}^{\gamma_s} \times \beta_i^{\frac{\epsilon_s \gamma_s}{\epsilon_s-1}} \times B_{st} \times \exp(\epsilon_{it}),$$

where m_{it} is total material *spending* by firm i and B_{st} collects general equilibrium variables which are common to all firms in industry s . Importantly, notice that by not directly observing physical quantities of intermediate inputs (but only expenditures), the firm's production function includes the term $\beta_i^{\epsilon_s \gamma_s / (\epsilon_s - 1)}$. Allowing for heterogeneity in terms of this home bias β_i will break the usual identifying assumptions for the production function parameters: even if intermediate inputs are strictly increasing in φ_{it} (see Akerberg, Caves, and Frazer 2015; and Levinsohn and Petrin 2003), variation in β_i will still affect both y_{its} and the demand for intermediary inputs. Alternatively, the "compound" productivity variable $\tilde{\varphi}_{it} = \varphi_{it} \times \beta_i^{\epsilon_s \gamma_s / (\epsilon_s - 1)}$ does not satisfy the structure in (B2) but contains a fixed component, which is generically correlated with firms' material spending.³⁶ We therefore assume in this section that there is no such heterogeneity in home bias, i.e., $\beta_i = \beta$.

To turn (B3) into an estimation equation, note that in our data firms' physical output is not directly observable—the fiscal files only report firm revenue. We therefore

³⁶ Introducing fixed effects would fix this problem, but at the usual cost of worsening measurement error problems (Griliches and Hausman 1986).

follow De Loecker (2011) and rely on specific assumptions about the demand structure. In particular, the CES demand implies that prices can be written by

$$\frac{p_i}{P_s} = \left(\frac{y_i}{Y_s}\right)^{-\frac{1}{\sigma_s}} \exp(\xi_{ist}),$$

where s denotes the industry in which the firm operates. Additionally, we include ξ_{ist} as an unobserved demand shock. We follow De Loecker (2011) and De Loecker et al. (2016) and decompose this demand shock into two components: $\xi_{ist} = \xi_{st} + \tilde{\xi}_{it}$. Here, ξ_{st} will be captured by industry specific time dummies, and the residual shocks $\tilde{\xi}_{it}$ are assumed to be i.i.d. across firms and time. Firm revenue can therefore be written as

$$Rev_{ist} = p_{ist} \cdot y_{ist} = y_{ist}^{\frac{\sigma_s-1}{\sigma_s}} Y_s^{\frac{1}{\sigma_s}} P_s \times \exp(\xi_{st} + \tilde{\xi}_{it}).$$

Taking logs and substituting (B3) yields the estimating equation

$$(B4) \quad \ln(Rev_{ist}) = \frac{\sigma_s-1}{\sigma_s} \ln(\varphi_{it}) + \tilde{\phi}_{ks} \ln(k_{ist}) + \tilde{\phi}_{ls} \ln(l_{ist}) \\ + \tilde{\gamma}_s \ln(m_{ist}) - \frac{\tilde{\gamma}_s}{\epsilon_s-1} \ln(s_{Dist}) + \ln(C_{st}) + \tilde{\xi}_{it} + \epsilon_{it},$$

where $\tilde{\gamma}_s = \gamma_s(\sigma_s - 1)/\sigma_s$, and $\tilde{\phi}_{ks}$ and $\tilde{\phi}_{ls}$ are defined accordingly. Furthermore, C_{st} collects all terms that are common to all firms in a same industry s (and are controlled for using time dummies).

Identification and estimation of this equation follows closely De Loecker (2011) and Edmond, Midrigan, and Xu (2015). The precise timing assumptions are as follows.³⁷ The firm:

- (i) observes the state variables $(k_{it}, l_{it}, \mathcal{L}_{it})$ as well as the realization of productivity $\varphi_{it}^{obs} = n(\varphi_{it-1}^{obs}) + \chi_{it}$;
- (ii) observes the prices of intermediate inputs (both domestic and internationally sourced) and makes optimal input choice;
- (iii) observe the shocks ϵ_{it} and ξ_{it} ;
- (iv) sets prices and produces output;
- (v) makes choices on future state variables $(k_{it+1}, l_{it+1}, \mathcal{L}_{it+1})$.

³⁷ These timing assumptions are very similar to Halpern, Koren, and Szeidl (2015).

Given this timing, the estimation procedure is as follows:

- (i) We follow Levinsohn and Petrin (2003) and rely on material demand to proxy for productivity φ_{it}^{obs} by inverting m :

$$(B5) \quad m_{it} = m_s(k_{it}, l_{it}, \mathcal{L}_{it}, \varphi_{it}^{obs}) \Leftrightarrow \ln(\varphi_{it}^{obs}) = h_s(m_{it}, k_{it}, l_{it}, \mathcal{L}_{it}).$$

In practice, we proxy firms' sourcing strategy by firms' import status and the number of varieties sourced internationally. Note that since the firm is a price taker on input markets, the price of inputs (conditional on the sourcing strategy) does not depend on φ_{it}^{obs} . This ensures that $m_s(\cdot)$ is invertible.

- (ii) We rely on h_s to proxy for firm productivity. Under our assumptions in (B2), χ_{it} is uncorrelated with the state variables as well as with lagged values of all inputs. Hence, we can estimate the production function technological parameters $\tilde{\phi}_{ks}$, $\tilde{\phi}_{ls}$, $\tilde{\gamma}_s$, and ϵ_s using the following moment conditions:

$$(B6) \quad \mathbb{E} \left[\chi_{it} \times \begin{pmatrix} k_{it} \\ l_{it} \\ k_{it-1} \\ l_{it-1} \\ m_{it-1} \\ m_{it-2} \\ z_{it} \end{pmatrix} \right] = 0,$$

where z_{it} is the Bartik type of instrument in (25). Note that this instrument is indeed exogenous under our timing assumptions if foreign supply shocks are unexpected. See also Hummels et al. (2014). We take advantage of these moments in the IV regression framework suggested by Petrin and Levinsohn (2012) and Wooldridge (2009). More specifically, we consider the equation

$$\begin{aligned} \ln(\text{Rev}_{ist}) &= \tilde{\phi}_{ks} \ln(k_{it}) + \tilde{\phi}_{ls} \ln(l_{it}) + \tilde{\gamma}_s \ln(m_{it}) - \frac{\tilde{\gamma}_s}{\epsilon_s - 1} \ln(s_{Dit}) \\ &\quad + \frac{\theta_s - 1}{\theta_s} n(h_s(X_{t-1})) + D_{st} + \chi_{it} + \tilde{\xi}_{it} + \epsilon_{it}, \end{aligned}$$

where

$$n(h_s(X_{t-1})) = n(h_s(m_{it-1}, k_{it-1}, l_{it-1}, \mathcal{L}_{it-1}))$$

TABLE B3—PRODUCTION FUNCTION ESTIMATION: PARAMETERS

Industry	ISIC	Cobb-Douglas			Translog in (k, l)			Observations
		$-\frac{\tilde{\gamma}_s}{\epsilon_s - 1}$	$\tilde{\gamma}_s$	ϵ_s	$-\frac{\tilde{\gamma}_s}{\epsilon_s - 1}$	$\tilde{\gamma}_s$	ϵ_s	
Mining	10–14	0.309 (0.177)	0.119 (0.076)	0.616 (0.324)	0.341 (0.184)	0.087 (0.075)	0.745 (0.254)	4,393
Food, tobacco, beverages	15–16	-0.358 (0.034)	0.459 (0.047)	2.285 (0.212)	-0.223 (0.031)	0.398 (0.046)	2.789 (0.381)	129,567
Textiles and leather	17–19	-0.226 (0.071)	0.233 (0.069)	2.031 (0.546)	-0.241 (0.071)	0.238 (0.069)	1.986 (0.500)	19,002
Wood and wood products	20	-0.252 (0.028)	0.352 (0.047)	2.397 (0.279)	-0.197 (0.026)	0.383 (0.046)	2.943 (0.399)	16,748
Paper, printing, publishing	21–22	-0.163 (0.042)	0.315 (0.058)	2.932 (0.709)	-0.141 (0.043)	0.314 (0.059)	3.233 (0.910)	34,301
Chemicals	24	0.111 (0.093)	0.767 (0.159)	-5.877 (5.244)	0.040 (0.088)	0.697 (0.150)	-16.38 (36.46)	7,502
Rubber and plastic products	25	-0.126 (0.048)	0.202 (0.094)	2.611 (1.190)	-0.170 (0.060)	0.081 (0.149)	1.478 (1.003)	11,989
Non-metallic mineral products	26	-0.383 (0.063)	0.311 (0.080)	1.813 (0.263)	-0.288 (0.061)	0.307 (0.079)	2.067 (0.382)	14,587
Basic metals	27	-0.678 (0.397)	-0.143 (0.519)	0.788 (0.655)	-0.697 (0.394)	-0.158 (0.513)	0.773 (0.623)	2,435
Metal products	28	-0.402 (0.023)	0.151 (0.026)	1.374 (0.0734)	-0.347 (0.023)	0.156 (0.025)	1.450 (0.0865)	61,017
Machinery and equipment	29	-0.191 (0.028)	0.323 (0.048)	2.688 (0.415)	-0.178 (0.028)	0.323 (0.048)	2.808 (0.459)	27,450
Office and computing machinery	30	-0.078 (0.134)	0.123 (0.189)	2.564 (3.823)	-0.059 (0.131)	0.118 (0.188)	2.996 (5.615)	655
Electrical machinery	31	-0.180 (0.055)	0.334 (0.084)	2.859 (0.910)	-0.201 (0.052)	0.334 (0.082)	2.659 (0.735)	8,326
Radio and communication	32	-0.301 (0.170)	0.238 (0.208)	1.790 (1.071)	-0.276 (0.177)	0.258 (0.209)	1.934 (1.279)	3,146
Medical and optical instruments	33	-0.243 (0.037)	0.306 (0.049)	2.261 (0.319)	-0.195 (0.040)	0.304 (0.048)	2.558 (0.454)	22,541
Motor vehicles, trailers	34	-0.203 (0.077)	0.599 (0.288)	3.958 (2.388)	-0.169 (0.072)	0.608 (0.281)	4.605 (2.972)	4,870
Transport equipment	35	-0.098 (0.150)	0.462 (0.129)	5.705 (7.770)	-0.106 (0.141)	0.477 (0.123)	5.515 (6.577)	3,949
Recycling, nec.	36–37	-0.386 (0.049)	0.303 (0.040)	1.786 (0.167)	-0.321 (0.047)	0.308 (0.039)	1.958 (0.216)	34,863

Notes: Robust standard errors are in parentheses. The table contains the results of estimating (B4) with the instruments given in (25). For non-importers, the instrument is set to zero in the full sample specifications. Estimation relies on data for the years 2004–2006 because two lags are required to build the appropriate instruments for the estimation of the production function. Standard errors for the estimates of ϵ_s are retrieved by the delta method.

is the control function for productivity (see (B5)) and $n(\cdot)$ stems from the productivity process (see (B2)). In practice we parametrize $n(h(\cdot))$ by a first-order polynomial and use the instruments in (B6).

Besides this parametrization with Cobb-Douglas technology, we also implement a more flexible translog specification, where we continue to assume a constant output elasticity for materials but allow for higher order terms in capital and labor.³⁸

³⁸ More specifically, we included the squared terms $\ln(k)^2$ and $\ln(l)^2$ and the interaction term $\ln(k) \times \ln(l)$.

Results.—We report the parameters of interest for both specifications in Table B3. For the majority of industries the point estimates are precisely estimated and in the same ballpark as the pooled estimate from the factor shares approach.³⁹ For a few other industries, we lack precision and we cannot reject that ε is below one: the existence of non-importers, however, implies that ε has to exceed unity. Overall and reassuringly, however, the estimated elasticities from the production function approach are within the sampling variation of the factor share estimates.

APPENDIX C: VARIABLE MARKUPS

In this section, we extend our analysis to a setting where competition among firms, and hence the distribution of markups, endogenously responds to changes in the trading environment. We follow Edmond, Midrigan, and Xu (2015) and consider an extension of the economy in Section I to allow for variable markups as in Atkeson and Burstein (2008). We first extend the sufficiency result in Proposition 1 to this environment with variable markups (Section A below). We then apply this result to estimate the gains from input trade in the French manufacturing sector (Section B).

A. Variable Markups: Theoretical Results

Demand is given by (8) above, together with

$$(C1) \quad C_s = \left(\int_0^{N_s} c_{js}^{\frac{\sigma_s-1}{\sigma_s}} dj \right)^{\frac{\sigma_s}{\sigma_s-1}} \quad \text{and} \quad c_{js} = \left(\sum_{i=1}^{N_{js}} c_{ijs}^{\frac{\theta_s}{\theta_s-1}} \right)^{\frac{\theta_s-1}{\theta_s}},$$

where $\theta_s \geq \sigma_s$.⁴⁰ Instead of each variety $j \in [0, N_s]$ being produced by a monopolistic firm, it is now given by another CES composite produced by a number N_{js} of firms. As in Atkeson and Burstein (2008), we assume that firms behave strategically vis-à-vis other firms producing their variety j , but take all other prices as given. In particular, we assume that firms compete à la Cournot. In this context, the change in consumer prices resulting from a shock to the import environment is given by Proposition 1 where Λ_s is now

$$(C2) \quad \Lambda_s^{VM} = \frac{1}{1-\sigma_s} \ln \left[\int_0^{N_s} \left(\sum_{i=1}^{N_{js}} \omega_i \left(\frac{SD_i}{S'_Di} \right)^{\frac{\gamma_s}{1-\varepsilon_s} (1-\theta_s)} \left(\frac{\mu'_i \left(\left[\omega_i, \frac{SD_i}{S'_Di} \right] \right)}{\mu_i(\omega_i)} \right)^{1-\theta_s} \right)^{\frac{1-\sigma_s}{1-\theta_s}} \omega_{js} dj \right].$$

³⁹ This suggests that the relatively low value for ε found in the pooled factor shares approach is not a result of the sectoral pooling of our data. This is in contrast to estimations on aggregate data, which find a downward bias (Imbs and Mejean 2015).

⁴⁰ As before, we assume that the domestic input aggregator in (9) takes the same form as consumer demand in (8) and (C1), where the Cobb-Douglas weights in (8) are given by $[c_i^x]$.

Here, ω_{js} denotes variety js 's share in total spending, ω_i denotes firm i 's share in variety spending, and μ_i, μ'_i denote firm i 's markup in the current and counterfactual equilibria. Note first that Λ_s^{VM} reduces to Λ_s in the constant markup case (i.e., when $\theta_s = \sigma_s$) as markups do not respond to the trade shock (i.e., $\mu_i = \mu'_i$).

Importantly, the notation in (C2) makes explicit that firms' counterfactual markups μ'_i are again fully determined from firms' domestic shares and size. To see how (C2) is calculated, note that current markups μ_i can be obtained from firms' observed revenue market shares via

$$(C3) \quad \mu_i = \frac{1}{1 - \left(\frac{1}{\sigma} \omega_i + \frac{1}{\theta} (1 - \omega_i)\right)}.$$

The distribution of counterfactual markups μ'_i can then be obtained from the following system of nonlinear equations,

$$(C4) \quad \frac{1}{\frac{1}{\sigma_s} - \frac{1}{\theta_s}} \left(\left(\frac{\theta_s - 1}{\theta_s} \right) - \frac{1}{\mu'_i} \right) = \frac{(\mu'_i)^{1-\theta_s} a_i}{\sum_{\nu} (\mu'_{\nu})^{1-\theta_s} a_{\nu}},$$

where

$$(C5) \quad a_i \equiv \mu_i^{-(1-\theta_s)} \omega_i \left(\frac{s_{Di}}{s'_{Di}} \right)^{\frac{\gamma_s}{1-\varepsilon_s} (1-\theta_s)}.$$

As $[a_i]$ in (C5) are a function of observables in the micro data (and parameters), so are firms' counterfactual markups $[\mu'_i]$ stemming from (C4). Once again, the micro data on firms' revenue and domestic expenditure shares $[\omega_i, s_{Di}/s'_{Di}]$ is sufficient to compute the change in consumer prices.

PROOF:

We now prove equations (C2)–(C5). Consumer preferences in (8) and (C1) imply that firm i 's demand is given by

$$(C6) \quad y_i = \left(\frac{p_i}{p_{js}} \right)^{-\theta_s} \left(\frac{p_{js}}{P_s} \right)^{-\sigma_s} Y_s,$$

where the variety and sector-level price indices are

$$(C7) \quad p_{js} = \left(\sum_{i=1}^{N_{js}} p_i^{1-\theta_s} \right)^{\frac{1}{1-\theta_s}},$$

$$(C8) \quad P_s = \left(\int_0^{N_s} p_{js}^{1-\sigma_s} dj \right)^{\frac{1}{1-\sigma_s}}.$$

Assuming that firms compete in à la Cournot, the profit maximization problem is given by

$$\max_{y_i} (p_i(y_i, y_{-i}) - c_i)y_i,$$

where c_i denotes firm i 's marginal cost and $p_i(y_i, y_{-i})$ is defined by (C6)–(C8). As in Atkeson and Burstein (2008), firms internalize the effect of their pricing decisions on y_{js} but not on Y_s . The profit-maximizing price can therefore be written as

$$(C9) \quad p_i = \frac{\left(\frac{1}{\sigma}\omega_i + \frac{1}{\theta}(1 - \omega_i)\right)^{-1}}{\left(\frac{1}{\sigma}\omega_i + \frac{1}{\theta}(1 - \omega_i)\right)^{-1} - 1} \times c_i \equiv \mu_i \times c_i,$$

where μ_i denotes the markup over marginal cost. Hence, (C10) provides a system of equations that characterizes optimal prices $[p_i]$ given marginal costs $[c_i]$. Using (C6), it is straightforward to show that

$$(C10) \quad \omega_i \equiv \frac{p_i y_i}{p_{js} y_{js}} = \left(\frac{p_i}{p_{js}}\right)^{1-\theta_s},$$

where ω_i denotes the variety-level expenditure share on firm i 's good. In particular, note that (C10) implies that $\mu_i = \mu(p_i/p_{js})$, i.e., markups are a function of relative prices. Thus, (C9) and (C10) establish (C3).

The change in sectoral price indices are

$$(C11) \quad \ln\left(\frac{P'_s}{P_s}\right) = \frac{1}{1-\sigma_s} \ln\left(\int_0^{N_s} \left(\frac{p'_{js}}{p_{js}}\right)^{1-\sigma_s} \omega_{js} dj\right),$$

where we used that

$$\omega_{js} = \frac{p_{js}^{1-\sigma_s}}{\int_0^{N_s} p_{js}^{1-\sigma_s} dj}.$$

In turn, the change in variety-level price indices is

$$(C12) \quad \frac{p'_{js}}{p_{js}} = \left(\sum_{i=1}^{N_{js}} \left(\frac{p'_i}{p_i}\right)^{1-\theta_s} \omega_i\right)^{\frac{1}{1-\theta_s}}.$$

Computing the change in the consumer price index therefore reduces to computing $[p'_i/p_i]$. Equation (C9) implies

$$(C13) \quad \frac{p'_i}{p_i} = \frac{\mu\left(\frac{p'_i}{p_{js}}\right)}{\mu_i} \times \frac{c'_i}{c_i},$$

where μ_i can be computed with data on the firm's revenue share ω_i via (C3). Using (6), we have that

$$(C14) \quad \frac{c'_i}{c_i} = \left(\frac{s'_{Di}}{s_{Di}} \right)^{\frac{\gamma_s}{\varepsilon_s - 1}} \left(\frac{p'_{Ds}}{p_{Ds}} \right)^{\gamma_s}.$$

It follows from (C13) and (C14) that computing p'_i/p_i reduces to computing the counterfactual relative prices p'_i/p'_{js} and p'_{Ds}/p_{Ds} . Combining (6) and (C9), we obtain

$$(C15) \quad p'_i = \mu \left(\frac{p'_i}{p'_{js}} \right) \mu_i^{-1} r_i^{\frac{1}{1-\theta_s}} \left(\frac{s'_{Di}}{s_{Di}} \right)^{\frac{\gamma_s}{\varepsilon_s - 1}} \left(\frac{p'_{Ds}}{p_{Ds}} \right)^{\gamma_s} \Phi_{js}^{-\frac{1}{1-\theta_s}},$$

where we substituted for unobserved efficiency:

$$\tilde{\varphi}_i = r_i^{\frac{1}{\theta_s - 1}} \mu_i \frac{\gamma_s}{\varepsilon_s - 1} p_{Ds}^{\gamma_s} \Phi_{js}^{\frac{1}{1-\theta_s}}.$$

Using (C7), (C15) can be expressed as

$$(C16) \quad \left(\frac{p'_i}{p'_{js}} \right)^{1-\theta_s} = \frac{\left(\mu \left(\frac{p'_i}{p'_{js}} \right) \mu_i^{-1} \right)^{1-\theta_s} \omega_i \left(\frac{s'_{Di}}{s_{Di}} \right)^{\frac{\gamma_s}{\varepsilon_s - 1} (1-\theta_s)}}{\sum_{\nu=1}^{N_{js}} \left(\mu \left(\frac{p'_\nu}{p'_{js}} \right) \mu_\nu^{-1} \right)^{1-\theta_s} \omega_\nu \left(\frac{s'_{D\nu}}{s_{D\nu}} \right)^{\frac{\gamma_s}{\varepsilon_s - 1} (1-\theta_s)}}.$$

Here, (C16) provides a system of equations that characterizes counterfactual relative prices $[p'_i/p'_{js}]$ given data $[\omega_i, s'_{Di}/s_{Di}]$. Given the expression for $\mu(p'_i/p'_{js})$ in (C9), we can derive equation (C4). Equations (C11) and (C12) together with (C13) and (C14) imply:

$$\ln \left(\frac{P'_s}{P_s} \right) = \gamma_s \ln \left(\frac{p'_{Ds}}{p_{Ds}} \right) + \Lambda_s,$$

where Λ_s is defined in (C2). This completes the proof. ■

B. Variable Markups: Empirical Results

We now implement this procedure for our specific application. According to (C2)–(C5), we require only two additional pieces of information. First, we need to make a choice regarding the level of aggregation of the “lower” market segment of the CES demand structure. Second, we need to estimate the additional demand elasticity θ_s . We then simply solve for the counterfactual distribution of markups from (C4) and calculate the change in consumer prices according to (C2) and Proposition 1.

TABLE C1—THE GAINS FROM INPUT TRADE IN FRANCE WITH VARIABLE MARKUPS: VARIATION ACROSS INDUSTRIES

	Constant markups		Variable markups					
	σ_s	Sectoral price gains	Case 1: $\theta_s = 2 \times \sigma_s$			Case 2: $\theta_s = 1.1 \times \sigma_s$		
			σ_s	θ_s	Sectoral price gains	σ_s	θ_s	Sectoral price gains
Mining	2.58	7.8	1.406	2.812	10.0	2.362	2.599	9.0
Food, tobacco, beverages	3.85	17.8	1.969	3.937	17.7	3.505	3.856	17.4
Textiles and leather	3.35	55.6	1.723	3.446	62.5	3.054	3.360	62.3
Wood and wood products	4.65	14.4	2.350	4.701	17.1	4.229	4.651	16.6
Paper, printing, publishing	2.77	20.1	1.400	2.801	21.1	2.518	2.769	20.8
Chemicals	3.29	45.1	1.711	3.422	43.2	3.002	3.302	43.0
Rubber and plastic products	4.05	38.4	2.141	4.282	39.5	3.705	4.076	39.1
Non-metallic mineral products	3.48	20.8	1.854	3.709	19.8	3.189	3.507	19.5
Basic metals	5.95	38.9	3.203	6.407	38.6	5.451	5.997	38.1
Metal products	3.27	18.3	1.653	3.306	20.5	2.972	3.269	20.2
Machinery and equipment	3.52	31.7	1.856	3.712	33.6	3.219	3.541	33.4
Office and computing machinery	7.39	44.6	4.108	8.216	68.7	6.790	7.469	73.6
Electrical machinery	4.49	36.1	2.443	4.886	38.4	4.120	4.532	38.7
Radio and communication	3.46	38.5	2.215	4.430	43.8	3.230	3.553	45.4
Medical and optical instruments	2.95	29.2	1.543	3.085	31.5	2.691	2.960	31.8
Motor vehicles, trailers	6.86	23.3	4.436	8.873	23.9	6.421	7.063	23.0
Transport equipment	1.87	22.9	1.067	2.134	22.5	1.724	1.896	22.5
Recycling, nec.	3.94	26.0	2.020	4.040	28.6	3.588	3.946	28.1

Note: The table contains the sectoral demand elasticities and sectoral price reductions for the baseline economy with constant markups (columns 1 and 2) and the economy with variable markups (columns 3–8).

TABLE C2—MARKUP REGRESSIONS

Industry	Ω_s	RCR_s	$q_{\omega,s}^{0.99} - q_{\omega,s}^{0.9}$	$q_{1/\mu}^{0.1} - q_{1/\mu}^{0.01}$	σ_s	θ_s
Mining	0.092	1.635	0.029	0.049	0.517	4.325
Food, tobacco, beverages	0.023	1.351	0.015	0.155	0.098	46.911
Textiles and leather	0.029	1.426	0.024	0.162	0.144	9.863
Wood and wood products	0.012	1.274	0.017	0.172	0.097	10.585
Paper, printing, publishing	0.013	1.566	0.013	0.124	0.103	4.108
Chemicals	0.041	1.437	0.048	0.167	0.274	6.157
Rubber and plastic products	0.056	1.328	0.016	0.168	0.095	-2.722
Non-metallic mineral products	0.064	1.402	0.078	0.142	0.502	5.874
Basic metals	0.077	1.202	0.132	0.141	0.868	11.574
Metal products	0.013	1.441	0.009	0.165	0.056	12.220
Machinery and equipment	0.054	1.397	0.060	0.186	0.311	8.575
Office and computing machinery	0.112	1.157	0.083	0.130	0.654	-24.235
Electrical machinery	0.088	1.286	0.040	0.170	0.246	-6.834
Radio and communication	0.282	1.407	0.039	0.175	0.282	-1.013
Medical and optical instruments	0.047	1.514	0.011	0.145	0.075	-3.270
Motor vehicles, trailers	0.293	1.171	0.342	0.200	1.789	-39.188
Transport equipment	0.141	2.150	0.267	0.176	0.908	2.264
Recycling, nec.	0.026	1.341	0.018	0.158	0.114	43.237

Notes: Columns 1 and 2 contain the measures of the aggregate Herfindahl index ($\Omega_s = \sum_{j \in s} \omega_{js} \sum_{i \in j} \omega_{ij}^2$) and the aggregate revenue cost ratio (see (C22)). Columns 3 and 4 contain the differences in the top 1 and 10 percent quantiles of the market share and (inverse) markup distribution. Columns 5 and 6 contain the implied elasticities σ_s and θ_s based on (C24) and (C25). In columns 7 and 8, we report the estimates and standard errors for the regression equation (C23) and the implied elasticities σ_s and θ_s based on (C19) and (C20).

We follow Gaubert and Itskhoki (2016) and model the lower segment of demand as a 3-digit industry. In our data, this leaves us with 106 subsectors. To discipline the two elasticities (θ_s, σ_s), we require the variable markup economy to be consistent

with the same moment as the constant markup economy, namely the aggregate revenue-cost ratio.⁴¹ The model implies that

$$(C17) \quad \frac{1}{RCR_s} = 1 - \frac{1}{\theta_s} (1 - \Omega_s) - \frac{1}{\sigma_s} \Omega_s,$$

where RCR_s denotes the aggregate revenue cost ratio in sector s and Ω_s is defined as $\Omega_s = \sum_{j \in s} \omega_{js} H_{js}$, where ω_{js} is the share of subsector j in sector s and H_{js} is the Herfindahl index of firms' market shares in subsector j . Equation (C17) defines a schedule of (θ_s, σ_s) that are consistent with RCR . To select among them, we consider two approaches: (i) selecting arbitrary values of θ_s to explore the sensitivity of our results to this parameter and (ii) choosing θ_s to match the dispersion of markups.

For the first approach, we consider two cases: one where θ_s is large relative to σ_s (we take $\theta_s = 2 \times \sigma_s$) and one where it is small (we take $\theta_s = 1.1 \times \sigma_s$). The changes in consumer prices are very similar in the two cases: in the high θ_s case, consumer prices in the manufacturing sector would be 28.9 percent higher in autarky; in the low θ_s case, the corresponding number is 28.7 percent. These estimates are quantitatively very similar to our earlier estimate of 27.5 percent in the constant markup economy. We note, however, that the importance of the markup channel varies substantially across sectors and, in particular, that input trade is pro-competitive in some industries and anticompetitive in others. This can be seen in Table C1, which contains the calibrated demand elasticities (σ_s, θ_s) and the price gains by sector for each of the two cases. For comparison, we also report the demand elasticity and the sectoral price gains of the constant markup economy (see Tables 1 and 4).

In our second approach, we bring an additional moment to identify (θ_s, σ_s) . It is natural to exploit the information about the dispersion in markups. The model implies that

$$(C18) \quad q_{1/\mu, s}^p - q_{1/\mu, s}^q = \left(\frac{1}{\sigma_s} - \frac{1}{\theta_s} \right) \times (q_{\omega, s}^{1-p} - q_{\omega, s}^{1-q}),$$

where $q_{1/\mu, s}^\tau$ ($q_{\omega, s}^\tau$) denotes the τ quantile of the distribution of the inverse of markups μ^{-1} (the within subsegment market shares ω_{ij}).⁴² Equations (C17) and (C18) are sufficient to identify the parameters (θ_s, σ_s) . Empirically, we find it challenging to find solutions to these equations that satisfy the theoretical restriction $1 \leq \sigma_s \leq \theta_s$. To see the source of the problem, consider Table C2. In this table we report—for all industries s —the data on RCR_s , Ω_s , $q_{1/\mu}^{0.1} - q_{1/\mu}^{0.01}$ and $q_{\omega, s}^{0.99} - q_{\omega, s}^{0.9}$.⁴³ To calculate the dispersion of markups, we rely on two measures of markups. To be consistent with our treatment of the RCR_s we consider the simple accounting equation $\mu_i^{ACC} = p_i y_i / (w l_i + m_i + R k_i)$. As an alternative we follow the approach pioneered by Jan De Loecker and measure markups from firms' cost-minimizing behavior as

⁴¹ The aggregate revenue-cost ratio is computed as the sector level average of firms' profit margins—see (20).

⁴² See below for a derivation of equations (C17) and (C18).

⁴³ Here we decided to focus on the tail of large firms as these are arguably the firms, where the Atkeson and Burstein (2008) model is most applicable. We tried many other combinations of quantiles, which all yielded qualitatively similar results. These results are available upon request.

$\mu_i^{CMIN} = \phi_l / (w_l i / p_i y_i)$, where ϕ_l is the output elasticity of labor in the production function. The results are contained in columns 2–5 in Table C2. In columns 6 and 7, we report the implied values for σ_s and θ_s according to (C17) and (C18). It is clearly seen that all but one σ_s is smaller than unity and that many θ_s are negative. The reason is that the dispersion in measured markups is large relative to the dispersion in market shares. Hence, $(q_{1/\mu}^{p_2} - q_{1/\mu}^{p_1}) / (q_{\omega,s}^{1-p_1} - q_{\omega,s}^{1-p_2})$ is large, which tends to reduce both σ_s and θ_s . One reason for this disconnect could be measurement error, if measured markups are more affected than market shares. While a possibility, the implied measurement error must be quite large. Even if $q_{1/\mu}^{p_2} - q_{1/\mu}^{p_1}$ was only half as large, most implied σ_s would still be below one.

We also pursued an alternative calibration strategy, which follows Edmond, Midrigan, and Xu (2015) closely. Using the observed markups, we could have estimated σ and θ from (C3). Edmond, Midrigan, and Xu (2015) measure markups from firms' labor shares (i.e., they use μ_i^{CMIN}). Markups and market shares are related via $(w_l i) / (p_i y_i) = \beta_0 + \beta_1 \omega_i + u_i$, where $\beta_0 = \phi_l (\theta - 1) / \theta$ and $\beta_1 = -\phi_l (1/\sigma - 1/\theta)$. Hence, the two demand elasticities are related via

$$(C19) \quad \theta = \frac{\beta_0 + \beta_1}{\frac{\beta_0}{\sigma} + \beta_1}.$$

It can be shown that together with (C18), this implies

$$(C20) \quad \sigma_s = \left(\frac{\frac{\beta_{0,s}}{\beta_{0,s} + \beta_{1,s}} (1 - \Omega_s) + \Omega_s}{\frac{RCR - 1}{RCR} - \frac{\beta_{1,s}}{\beta_{0,s} + \beta_{1,s}} (1 - \Omega_s)} \right),$$

which identifies σ_s directly from moments in the data. When we implement this approach, we find that our estimates for β_1 are consistently negative as required by the theory.⁴⁴ As before, we find estimates for σ_s , which are often below unity and implied values for θ_s , which are negative.⁴⁵ We can also compare our results with Edmond, Midrigan, and Xu (2015). They consider the exact same specification for Taiwanese firms but do not allow for their parameters to vary at the sector level. They estimate $\beta_0 = 0.64$ and $\beta_1 = -0.5$ and arrive at $\sigma = 1.24$ and $\theta = 10.5$.⁴⁶ Hence, their “upper” demand elasticity σ is also small. In particular, from

⁴⁴ All these results are available upon request.

⁴⁵ As a third possibility, we also allowed for β_0 to differ at the subsegment level (by including subsegment fixed effects) and only used $\beta_1 = -\phi_l (1/\sigma - 1/\theta)$ (together with an estimate of ϕ_l) for identification. This strategy gave qualitatively similar results, which are available upon request.

⁴⁶ Their strategy is to calibrate σ within the context of a model to match the aggregate trade elasticity. Given σ they then use their estimate for β_0 and β_1 to identify θ . Note that their notation is different. They use γ instead of θ and θ instead of σ .

$\theta = (\beta_0 + \beta_1)/((\beta_0/\sigma) + \beta_1)$ it is easy to see that θ can only be positive if $\sigma < 1.28$.⁴⁷

We conclude that the model is unable to rationalize the observed dispersion in markups for parameters (θ_s, σ_s) that are consistent with (C17) and satisfy the theoretical restriction $1 \leq \sigma_s \leq \theta_s$. The French data therefore seems to ask for a richer model of pricing than the Atkeson and Burstein (2008) model. While we do find that markups are increasing in subsector market shares, the relationship is quite noisy, i.e., there is ample cross-firm dispersion in markups that is not explained by the variation in market shares.

Derivation of Equations (C17) and (C18).—As we showed above, the equilibrium markup of firm i is given by

$$(C21) \quad \mu_i = \frac{\theta_s(1 - \omega_{ijs}) + \sigma_s \omega_{ijs}}{\theta_s(1 - \omega_{ijs}) + \sigma_s \omega_{ijs} - 1}.$$

Here ω_{ijs} is firm i 's sales share in segment j of sector s . Our calibration strategy is to calibrate the demand elasticities (θ_s, σ_s) to match the aggregate revenue-cost ratio, i.e., $RCR_s = (\sum_i p_i y_i) / (\sum_i Cost_i)$, which is the same moment as in our benchmark economy (see (20)) and a moment related to the dispersion in markups. Using (C21) it is easy to show that

$$(C22) \quad RCR_s = \frac{\sum_i p_i y_i}{\sum_i p_i y_i \mu_i^{-1}} = \frac{1}{\sum_i \frac{p_i y_i}{\sum_i p_i y_i} \mu_i^{-1}} = \frac{1}{\sum_i \frac{p_i y_i}{\sum_{i \in j} p_i y_i} \frac{\sum_{i \in j} p_i y_i}{\sum_i p_i y_i} \mu_i^{-1}}$$

$$= \left(\sum_i \omega_{ij} \omega_{js} \left(1 - \left(\frac{1}{\theta} (1 - \omega_{ij}) + \frac{1}{\sigma} \omega_{ij} \right) \right) \right)^{-1}$$

$$= \left(\left(1 - \frac{1}{\theta} \left(1 - \sum_{j \in s} \omega_{js} \sum_{i \in j} \omega_{ij}^2 \right) - \frac{1}{\sigma} \sum_{j \in s} \omega_{js} \sum_{i \in j} \omega_{ij}^2 \right) \right)^{-1}.$$

Here ω_{ij} denotes firm i 's market in subsector j and ω_{js} denotes the share of subsector j in industry s . Note that $H_{js} \equiv \sum_{i \in j} \omega_{ij}^2$ is simply the Herfindahl index of subsector j so that $\Omega_s = \sum_{j \in s} \omega_{js} H_{js}$ is the (weighted) average of the Herfindahl indices in sector s . Using the definition of Ω_s and (C22) yields (C17). To derive (C18), note that (C21) implies a linear relationship between inverse markups and market shares, i.e.,

$$(C23) \quad \mu_i^{-1} = \frac{\theta - 1}{\theta} - \left(\frac{1}{\sigma} - \frac{1}{\theta} \right) \omega_{ij}.$$

⁴⁷ When we pool all the data and estimate a common β_1 and β_0 , we find $\beta_1 = -0.517$ and $\beta_0 = 0.33$. Hence, we estimate the exact same slope parameter and a lower intercept. The difference in the intercept stems from the fact that we use revenue, while they use value added. According to the theory, this should not make a difference. Quantitatively, the results in Table C2 imply that $\sigma_s < 1$ for many sectors even if $\beta_0 = 0.6$.

Using the definition of quantiles, (C18) follows directly from (C23). Combining (C17) and (C18), we can solve for the parameters in terms of data moments. In particular, for any two quantiles p_1 and p_2 we get

$$(C24) \quad \sigma_s = \frac{RCR_s}{RCR_s - 1 + (1 - \Omega_s) RCR_s \frac{q_{1/\mu}^{p_2} - q_{1/\mu}^{p_1}}{q_{\omega,s}^{1-p_1} - q_{\omega,s}^{1-p_2}}},$$

$$(C25) \quad \theta_s = \left(\frac{1}{\sigma_s} - \frac{q_{1/\mu}^{p_2} - q_{1/\mu}^{p_1}}{q_{\omega,s}^{1-p_1} - q_{\omega,s}^{1-p_2}} \right)^{-1}.$$

REFERENCES

- Akerberg, Daniel A., Kevin Caves, and Garth Frazer. 2015. "Identification Properties of Recent Production Function Estimators." *Econometrica* 83 (6): 2411–51.
- Amiti, Mary, Mi Dai, Robert C. Feenstra, and John Romalis. 2017. "How Did China's WTO Entry Benefit U.S. Consumers?" National Bureau of Economic Research (NBER) Working Paper 23487.
- Amiti, Mary, Oleg Itskhoki, and Jozef Konings. 2014. "Importers, Exporters, and Exchange Rate Disconnect." *American Economic Review* 104 (7): 1942–78.
- Amiti, Mary, and Jozef Konings. 2007. "Trade Liberalization, Intermediate Inputs, and Productivity: Evidence from Indonesia." *American Economic Review* 97 (5): 1611–38.
- Antràs, Pol, Teresa C. Fort, and Felix Tintelnot. 2017. "The Margins of Global Sourcing: Theory and Evidence from US Firms." *American Economic Review* 107 (9): 2514–64.
- Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare. 2012. "New Trade Models, Same Old Gains?" *American Economic Review* 102 (1): 94–130.
- Atkeson, Andrew, and Ariel Burstein. 2008. "Pricing-to-Market, Trade Costs, and International Relative Prices." *American Economic Review* 98 (5): 1998–2031.
- Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott. 2012. "The Empirics of Firm Heterogeneity and International Trade." *Annual Review of Economics* 4: 283–313.
- Blaum, Joaquin, Claire Lelarge, and Michael Peters. 2017. "Firm Size and the Intensive Margin of Import Demand." Centre for Economic Policy Research (CEPR) Working Paper 12237.
- Blaum, Joaquin, Claire Lelarge, and Michael Peters. 2018. "The Gains from Input Trade with Heterogeneous Importers: Dataset." *American Economic Journal: Macroeconomics*. <https://doi.org/10.1257/mac.20160426>.
- Bøler, Esther Ann, Andreas Moxnes, and Karen Helene Ulltveit-Moe. 2015. "R&D, International Sourcing, and the Joint Impact on Firm Performance." *American Economic Review* 105 (12): 3704–39.
- Broda, Christian, and David E. Weinstein. 2006. "Globalization and the Gains from Variety." *Quarterly Journal of Economics* 121 (2): 541–85.
- Caliendo, Lorenzo, and Fernando Parro. 2015. "Estimates of the Trade and Welfare Effects of NAFTA." *Review of Economic Studies* 82 (1): 1–44.
- Costinot, Arnaud, and Andrés Rodríguez-Clare. 2014. "Trade Theory with Numbers: Quantifying the Consequences of Globalization." In *Handbook of International Economics*, Vol. 4, edited by Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, 197–261. Amsterdam: North-Holland.
- De Loecker, Jan. 2011. "Product Differentiation, Multiproduct Firms, and Estimating the Impact of Trade Liberalization on Productivity." *Econometrica* 79 (5): 1407–51.
- De Loecker, Jan, and Pinelopi Koujianou Goldberg. 2014. "Firm Performance in a Global Market." *Annual Review of Economics* 6: 201–27.
- De Loecker, Jan, Pinelopi K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik. 2016. "Prices, Markups and Trade Reform." *Econometrica* 84 (2): 445–510.
- De Loecker, Jan, and Frederic Warzynski. 2012. "Markups and Firm-Level Export Status." *American Economic Review* 102 (6): 2437–71.
- Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography, and Trade." *Econometrica* 70 (5): 1741–79.

- Eaton, Jonathan, Samuel Kortum, and Francis Kramarz.** 2011. "An Anatomy of International Trade: Evidence from French Firms." *Econometrica* 79 (5): 1453–98.
- Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu.** 2015. "Competition, Markups, and the Gains from International Trade." *American Economic Review* 105 (1): 3183–3221.
- Feenstra, Robert C.** 1994. "New Product Varieties and the Measurement of International Prices." *American Economic Review* 84 (1): 157–77.
- Fielor, Ana Cecilia, Marcela Eslava, and Daniel Yi Xu.** 2018. "Trade, Quality Upgrading, and Input Linkages: Theory and Evidence from Colombia." *American Economic Review* 108 (1): 109–46.
- Gaubert, Cecile, and Oleg Itskhoki.** 2016. "Granular Comparative Advantage." <http://www.princeton.edu/~itskhoki/papers/GranularCA.pdf>.
- Goldberg, Pinelopi Koujianou, Amit Jumar Khandelwal, Nina Pavcnik, and Petia Topalova.** 2010. "Imported Intermediate Inputs and Domestic Product Growth: Evidence from India." *Quarterly Journal of Economics* 125 (4): 1727–67.
- Gopinath, Gita, and Brent Neiman.** 2014. "Trade Adjustment and Productivity in Large Crises." *American Economic Review* 104 (3): 793–831.
- Griliches, Zvi, and Jerry A. Hausman.** 1986. "Errors in variables in panel data." *Journal of Econometrics* 31 (1): 93–118.
- Halpern, László, Miklóós Koren, and Adam Szeidl.** 2015. "Imported Inputs and Productivity." *American Economic Review* 105 (12): 3660–3703.
- Hummels, David, Rasmus Jørgensen, Jakob Munch, and Chong Xiang.** 2014. "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data." *American Economic Review* 104 (6): 1597–1629.
- Imbs, Jean, and Isabelle Mejean.** 2015. "Elasticity Optimism." *American Economic Journal: Macroeconomics* 7 (3): 43–83.
- Kasahara, Hiroyuki, and Joel Rodrigue.** 2008. "Does the use of imported intermediates increase productivity? Plant-level evidence." *Journal of Development Economics* 87 (1): 106–18.
- Kugler, Maurice, and Eric Verhoogen.** 2012. "Prices, Plant Size, and Product Quality." *Review of Economic Studies* 79 (1): 307–39.
- Levinsohn, James, and Amil Petrin.** 2003. "Estimating Production Functions Using Inputs to Control for Unobservables." *Review of Economic Studies* 70 (2): 317–41.
- Lu, Dan, Asier Mariscal, and Luis-Fernando Mejía.** 2017. "How Firms Accumulate Inputs: Evidence from Import Switching." Unpublished.
- Oberfield, Ezra, and Devesh Raval.** 2014. "Micro Data and Macro Technology." National Bureau of Economic Research (NBER) Working Paper 20452.
- Olley, G. Steven, and Ariel Pakes.** 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64 (6): 1263–97.
- Pavcnik, Nina.** 2002. "Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants." *Review of Economic Studies* 69 (1): 245–76.
- Petrin, Amil, and James Levinsohn.** 2012. "Measuring aggregate productivity growth using plant-level data." *RAND Journal of Economics* 43 (4): 705–25.
- Ramanarayanan, Ananth.** 2017. "Imported Inputs and the Gains from Trade." Unpublished.
- Simonovska, Ina, and Michael E. Waugh.** 2014a. "The elasticity of trade: Estimates and evidence." *Journal of International Economics* 92 (1): 34–50.
- Simonovska, Ina, and Michael E. Waugh.** 2014b. "Trade Models, Trade Elasticities, and the Gains from Trade." National Bureau of Economic Research (NBER) Working Paper 20495.
- Syerson, Chad.** 2011. "What Determines Productivity?" *Journal of Economic Literature* 49 (2): 326–65.
- Topalova, Petia, and Amit Khandelwal.** 2011. "Trade Liberalization and Firm Productivity: The Case of India." *Review of Economics and Statistics* 93 (3): 995–1009.
- Wooldridge, Jeffrey M.** 2009. "On estimating firm-level production functions using proxy variables to control for unobservables." *Economics Letters* 104 (3): 112–14.