

The Margins of Global Sourcing: Theory and Evidence from U.S. Firms*

Pol Antràs
Harvard University

Teresa C. Fort
Tuck School at Dartmouth

Felix Tintelnot
University of Chicago

July 3, 2014

(PRELIMINARY and INCOMPLETE)

Abstract

This paper studies the extensive and intensive margins of firms' global sourcing decisions. We develop a quantifiable multi-country sourcing model in which heterogeneous firms self-select into importing based on their productivity and country-specific variables (wages, trade costs, and technology). The model delivers a simple closed-form solution for firm profits as a function of the countries from which a firm imports, as well as those countries' characteristics. A key feature of the profit function is that the marginal change in profits from adding a country to the firm's set depends on the other countries in the set. These interdependencies make the analysis of the extensive margin of sourcing more complicated than in models of exporting, where entry is typically assumed to be independent across markets. Under plausible parametric restrictions, however, selection into importing features complementarity across markets. In this case, firms' sourcing strategies follow a strict hierarchical structure also predicted by exporting models. Our quantitative analysis exploits these complementarities to distinguish between a country's potential as a marginal cost-reducing source of inputs and the fixed cost associated with sourcing from this country. Counterfactual exercises suggest that a shock to the potential benefits of sourcing from a country leads to significant and heterogeneous changes in sourcing across both countries and firms.

*Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. We are grateful to Treb Allen, Andy Bernard, Arnaud Costinot, Paul Grieco, Gene Grossman, Elhanan Helpman, Sam Kortum, Marc Melitz, and Eduardo Morales for useful conversations and to BooKang Seol and Linh Vu for research assistance. We have also benefited from very useful feedback from seminar audiences at Barcelona GSE, Boston University, Dartmouth, ERWIT in Oslo, Harvard, National Bank of Belgium, MIT, Northwestern, Princeton, Tsinghua, UBC, Virginia, and Yale.

1 Introduction

In recent years, theoretical research in international trade has adopted a decidedly granular approach. The workhorse models in the literature now derive aggregate trade and multinational activity flows by aggregating the individual decisions of firms in the world economy. This novel approach is empirically anchored in a series of studies in the 1990s that demonstrate the existence of a large degree of intraindustry heterogeneity in revenue, productivity, factor inputs, and export participation across firms. With regards to export behavior, the studies show that only a small fraction of firms, those appearing to be particularly productive, engage in exporting, and that most exporting firms sell only to a few markets. This so-called extensive margin of trade has been shown to be important in understanding variation in aggregate exports across destination markets.

The typical model in this literature focuses on the export decisions of heterogeneous firms producing differentiated final goods that are demanded by consumers worldwide. In the real world, however, a significant share of the volume of international trade – possibly up to two-thirds – is accounted for by shipments of intermediate inputs (see Johnson and Noguera, 2012). As a result, and to the extent that the firms systematically select into importing, it is likely that firm-level import decisions are at least as important as firm-level export decisions in explaining aggregate trade patterns.

Given that every international trade transaction involves an exporter and an importer, one might wonder: why should one care about whether the extensive margin of trade is shaped by the export or import decisions of firms? Or, in other words, is the export versus import distinction relevant for the aggregate implications of models with intraindustry heterogeneity? This paper develops a new framework to analyze the margins of global sourcing in a multi-country environment and uses the model to highlight, both theoretically as well as empirically, several differential features of the determination of the margins of trade relative to their determination in canonical models of exporting.

Although the margins of trade have been much more systematically studied on the export side than on the import side, it is well-known that the extensive margins of imports (the number of importing firms and the number of imported products) are important in explaining aggregate imports, and that import-market participation varies systematically with firm characteristics. For instance, Bernard et al. (2009) find that about 65 percent of the cross-country variation in U.S. imports is accounted for the extensive margins of imports, while Bernard et al. (2007) show that U.S. importers are on average more than twice as large and about 12 percent more productive than non-importers.¹


Figure 1 provides a graphical illustration of the size premium of importers and how it varies with the number of countries from which a firm sources.² The figure indicates that firms that

¹We obtain very similar findings when replicating these analyses for the sample of U.S. manufacturing firms used in our empirical analysis (see the Online Appendix).

²To construct the figure, we regress the log of firm sales on cumulative dummies for the number of countries from which a firm sources and industry controls. The omitted category is non-importers, so the premia are interpreted as the difference in size between non-importers and firms that import from at least one country, at least two countries,

import from one country are more than twice the size of non-importers, firms that source from 13 are about four log points larger, and firms sourcing from 25 or more countries are over six log points bigger than non-importers. These enormous differences are suggestive of the empirical relevance of country-specific fixed costs of sourcing, which limit the ability of small firms to select into importing from a large number of countries.³

Figure 1: Sales premia and minimum number of sourcing countries in 2007



Not only are fixed costs of sourcing sizable, but they also seem to vary significantly across countries. To see this, consider Table 1, which lists the top ten sourcing countries for U.S. manufacturers in 2007, based on the number of importing firms. These countries account for 93 percent of importers in our sample and 74 percent of imports. The first two columns provide the country rank by number of firms and by import value. Canada ranks number one for manufacturers in both dimensions. For other countries, however, there are significant differences in these ranks. China is number two for firms but only number three for value. Mexico, the number two country in terms of value, ranks eighth in terms of the number of importing firms. Columns 3-4 provide details on the number of firms that import from each country and the fraction of total importers they represent. Columns 5-6 give similar information for import values. It is clear that there are significant differences across countries in extensive and intensive margins. For example, the U.K. and Taiwan account for only three and two percent of total imports respectively, but 18 percent of all importers source from the U.K. and 16 percent source from Taiwan. These considerable divergences between the intensive and extensive margins of trade are suggestive of the importance of heterogeneity in the fixed costs of sourcing from particular countries. Under this interpretation, Table 1 appears

etc. The horizontal axis denotes the number of countries from which a firm sources, with 1 corresponding to firms that use only domestic inputs.

³These premia are robust to controlling for the number of products a firm imports and the number of products it exports and thus do not merely capture the fact that larger firms import more products. Consistent with selection into importing, the same qualitative pattern is also evident among firms that did not import in 2002 and when using employment or productivity rather than sales. See the Online Appendix for additional details.

to indicate that relative to the sourcing potential of these countries, fixed costs of sourcing are disproportionately high in Mexico and Japan and low in Italy and Taiwan.

Table 1: Top 10 source countries for U.S. firms, by number of firms

	Rank by:		Number of Importers		Value of Imports	
	Firms	Value	Firms	% of Total	Imports	% of Total
Canada	1	1	37,800	59	145,700	16
China	2	3	21,400	33	121,980	13
Germany	3	5	13,000	20	62,930	7
United Kingdom	4	6	11,500	18	30,750	3
Taiwan	5	11	10,500	16	16,630	2
Italy	6	13	8,500	13	13,230	1
Japan	7	4	8,000	12	112,250	12
Mexico	8	2	7,800	12	125,960	14
France	9	9	6,100	9	22,980	3
Korea, South	10	10	5,600	9	20,390	2

Notes: Sample is U.S. firms with some manufacturing activity in 2007. Number of firms rounded to nearest 100 for disclosure avoidance. Imports in millions of \$s, rounded to nearest 10 million for disclosure avoidance.

Motivated by these patterns, in section 2 we develop a quantifiable multi-country sourcing model with heterogeneous firms in which firms self-select into importing based on their productivity and country-specific variables (wages, trade costs, and technology). Firms can, in principle, buy intermediate inputs from any country in the world. Nevertheless, adding a country to the set of countries a firm is able to import from requires incurring a market-specific fixed cost. As a result, relatively unproductive firms naturally opt out of importing from certain countries that are not particularly attractive sources of inputs.

Once a firm has determined the set of countries from which it has secured the ability to source inputs – which we refer to as its *global sourcing strategy* – it then learns the various firm-specific efficiency levels with which each input can be produced in each of these ‘active’ countries. These efficiency levels are assumed to be drawn from an extreme-value Fréchet distribution, as in Eaton and Kortum (2002). Factoring labor costs in those countries as well as transport costs, firms then decide the optimal source for each of the inputs used in production.

The model delivers a simple closed-form solution for the profits of the firm as a function of its *sourcing capability*, which in turn is a function of the set of countries from which a firm has paid the fixed costs to import and those countries’ characteristics (wages, trade costs, and average technology). The sourcing capability of a firm increases when a country is added to its sourcing strategy. Intuitively, by enlarging that set, the firm benefits from greater competition among suppliers, and thereby lowers the effective cost of its intermediate input bundle. The choice of a sourcing strategy therefore trades off lower variable cost of production against the greater fixed

costs associated with a more complex global sourcing strategy. Quite intuitively, we show that more productive firms necessarily choose strategies that give them (weakly) higher sourcing capabilities, which implies that their cost advantage is magnified by their sourcing decisions, thus generating an increased skewness in the size distribution of firms.

A key feature of the derived profit function is that the marginal change in profits from adding a country to the firm's set depends on the set of countries from which a firm imports, as well as those countries' characteristics. This contrasts with standard models of selection into exporting featuring constant marginal costs, in which the decision to service a given market is independent of that same decision in other markets. Whether the decisions to source from different countries are complements or substitutes crucially depends on a parametric restriction involving the elasticity of demand faced by the final-good producer and the Fréchet parameter governing the variance in the distribution of firm-specific input efficiencies across locations. Selection into importing features complementarity across markets whenever demand is relatively elastic (so profits are particularly responsive to variable cost reductions) and whenever input efficiency levels are relatively heterogeneous across markets (so that the reduction in expected costs achieved by adding an extra country in the set of active locations is relatively high).

Conversely, when demand is inelastic or input efficiency draws are fairly homogeneous, the addition of country to a firm's global sourcing strategy instead reduces the marginal gain from adding other locations. In such a situation, the problem of a firm optimally choosing its sourcing strategy is extremely hard to characterize, both analytically as well as quantitatively, since it boils down to solving a combinatorial problem with 2^J elements (where J is the number of countries) with little guidance from the model.

The case with complementary sourcing decisions turns out to be much more tractable and delivers sharp results rationalizing the monotonicity in the sales premia observed in Figure 1. In particular, we use standard tools from the monotone comparative statics literature to show that the sourcing strategies of firms follow a strict hierarchical structure in which the number of countries in a firm's sourcing strategy is (weakly) increasing in the firm's core productivity level. The attractiveness of a sourcing location is being shaped not only by wages, technology, transport costs, but also by fixed costs of offshoring associated with that sourcing location.

We also use the model to show how the aggregation of firms' sourcing decisions shapes aggregate input flows across countries. As in heterogeneous firms models of exporting, the extensive margin of global sourcing tends to amplify the response of trade flows to changes in variable trade costs. In fact, in the knife-edge case in which parameter values are such that the sourcing decisions are independent across markets, and when assuming a Pareto distribution of core productivities across final-good producers, our model delivers a gravity equation for input flows that is identical to those derived by Chaney (2008), Arkolakis et al. (2008), or Helpman et al. (2008) in models of exporting. When we depart from that knife-edge case, however, our model gives rise to an *extended* gravity equation featuring third market effects, which invalidates the estimation of the aggregate trade elasticity with standard approaches.

Our quantitative analysis enables separate identification of the sourcing potential (a function of technology, trade cost, and wages) of a country and the fixed cost of sourcing from that country. In a first step, we use the structure of the model to recover the sourcing potential of 64 foreign countries from U.S. firm-level data on the intensive margin of intermediate input purchases. In the second step, we use the estimated sourcing potentials and data on average markups in U.S. manufacturing to estimate the Fréchet parameter governing the dispersion in productivity of intermediate good producers and the elasticity of final-good demand. We find robust evidence suggesting that the extensive margin sourcing decisions of U.S. firms are complements, consistent with the pattern documented in Figure 1. Furthermore, our estimates imply that a firm sourcing from all foreign countries faces 7 to 11 percent lower variable costs, with the exact value depending on the preferred specification for the estimation of the dispersion parameter.⁴ In our third and final step, we estimate firm-country-specific fixed costs that depend on country characteristics using simulated method of moments. To do so, we adopt the iterative algorithm developed by Jia (2008), which exploits the complementarities in the ‘entry’ decisions of firms, and uses lattice theory to greatly reduce the dimensionality of the firm’s optimal sourcing strategy problem. We estimate a median fixed cost of sourcing from a country of 40-60 thousand U.S. dollars and find around 30 percent lower fixed cost of sourcing for countries with a common language. Our estimated model matches moments of the size distributions of U.S. firms’ input purchases and the shares of firms importing from foreign countries. We are in the process of using our estimates from the model to perform a series of counterfactual exercises.

Our paper is related to several strands of the literature. We build on the approach in Tintelnot (2012) of embedding the Eaton and Kortum (2002) stochastic representation of technology into the problem of a firm optimally choosing a production location in a multi-country world. The focus of our paper is nevertheless quite distinct. While Tintelnot (2012) studies the location of final-good production of multi-product multinational firms, we instead develop a model of global sourcing of inputs. Rodríguez-Clare (2010) and Garetto (2013) also adapt the Eaton and Kortum (2002) framework to the study of global sourcing decisions but in two-country models and with very different goals in mind.⁵

Several recent papers have explored the implications of foreign intermediate inputs on firm performance and aggregate productivity. These include the studies by Amiti and Konings (2007) for Indonesia, Goldberg et al. (2010) and Loecker et al. (2012) for India, Halpern et al. (2011) for Hungary, and Gopinath and Neiman (2013) for Argentina. Because the focus of these papers is largely empirical, these authors develop simple two-country models in which domestic and foreign inputs are assumed to be imperfectly substitutable, so that the productivity improvements resulting from an increase in imports of intermediate inputs are associated with love-for-variety effects, as

⁴Quantitatively, this is comparable to the findings of Halpern et al. (2011) for Hungarian firms. Using a two-country model and a method similar to Olley and Pakes (1996), they find that importing all foreign varieties would increase firm productivity of a Hungarian firm by 12 percent.

⁵See also Ramondo and Rodríguez-Clare (2013) and Arkolakis et al. (2013) for recent related contributions to quantitative models of multinational activity.

in Ethier (1982). Instead, in our framework, inputs are not differentiated by country of origin, and the gains in profitability associated with an expansion in the extensive margin of imports are brought about by an increase in competition across input sources. Moreover, our model can flexibly accommodate an arbitrary number of countries and is thus a more reliable tool for quantitative and counterfactual analysis.

Our paper is not the first one to describe the inherent difficulties in solving for the extensive margin of imports in a multi-country model with multiple intermediate inputs. Yeaple (2003) and Grossman et al. (2006) obtain partial characterizations of this problem when focusing on models with at most three countries and at most two inputs. In a model with an arbitrary number of countries and inputs, Blaum et al. (2013) discuss the existence of interdependencies across sourcing decisions analogous to those in our model (although with a love-for-variety model in mind), and conclude that their model provides no general predictions for the extensive margin of importing. Instead, we show that under certain plausible parametric restrictions, the problem can actually be characterized with standard optimization techniques.⁶ In addition, we adopt an algorithm from the industrial organization literature that exploits features of the profit function to solve what would otherwise be a computationally infeasible problem.

Finally, at a broader level, our paper naturally relates to the vast literature on offshoring, and particularly to the work of Antràs and Helpman (2004, 2008), Antràs et al. (2006), Grossman and Rossi-Hansberg (2008, 2012), and Fort (2013) who all emphasize the importance of heterogeneity for understanding the margins of offshoring. Relative to those papers, our focus is on the role of heterogeneous productivity (rather than heterogeneous skills or offshoring costs), and we also provide a tractable multi-country model, while those papers are based on two-country frameworks.⁷

The rest of the paper is structured as follows. We describe the assumptions of our multi-country model of global sourcing in section 2 and solve for the equilibrium in section 3. In section 4 we describe the firm-level data and provide reduced-form evidence that supports our theoretical framework. We estimate the model structurally in section 5 and perform several counterfactual exercises in section 6. Section 7 concludes.

2 Theoretical Framework

In this section we develop our quantifiable multi-country model of global sourcing.

⁶The seminal applications of the mathematics of complementarity in the economics literature are Vives (1990) and Milgrom and Roberts (1990). Grossman and Maggi (2000) and Costinot (2009) are particularly influential applications of these techniques in international trade environments.

⁷The focus in our paper is on international trade transactions in which firms are both on the buying and selling side of the transaction. As such, our work is somewhat related to a burgeoning literature that studies firm-to-firm interactions, including the work of Blum et al. (2010), Eaton et al. (2012), Eaton et al. (2013), Carballo et al. (2013), Bernard et al. (2013), Dragusanu (2014), and Monarch (2013).

2.1 Preferences and Endowments

Consider a world consisting of J countries where individuals value the consumption of differentiated varieties of manufactured goods according to a standard symmetric CES aggregator

$$U_M = \left(\int_{\omega \in \Omega_j} q(\omega)^{(\sigma-1)/\sigma} d\omega \right)^{(\sigma/\sigma-1)}, \quad \sigma > 1, \quad (1)$$

where Ω_j is the set of manufacturing varieties available to consumers in country $j \in J$ (with some abuse of notation we denote by J both the number as well as the set of countries). These preferences are assumed to be common worldwide and give rise to the following demand for variety ω in country j :

$$q_j(\omega) = E_j P_j^{\sigma-1} p_j(\omega)^{-\sigma}, \quad (2)$$

where $p_j(\omega)$ is the price of variety ω , P_j is the standard ideal price index associated with (1), and E_j is aggregate spending on manufacturing goods in country j . For what follows it will be useful to define a (manufacturing) market demand term for market j as follows

$$B_j = \frac{1}{\sigma} \left(\frac{\sigma}{\sigma-1} \right)^{1-\sigma} E_j P_j^{\sigma-1}. \quad (3)$$

There is a unique factor of production, labor, which commands a wage w_j in country j . When we close the model in general equilibrium, we later introduce an additional non-manufacturing sector into the economy. This non-manufacturing sector captures a constant share of the economy's spending, also employs labor, and is large enough to pin down wages in terms of that 'outside' sector's output. Although these assumptions are not important for the qualitative results of the paper, they are necessary for tractability in performing counterfactual exercises in our multi-country sourcing environment.

2.2 Technology and Market Structure

There exists a measure N_j of final-good producers in each country $j \in J$ and each of these producers owns a blueprint to produce a single differentiated variety. The market structure of final good production is characterized by monopolistic competition and there is free entry into the industry.

Production of final-good varieties requires the assembly of a bundle of *intermediates*. Intermediates can be offshored and a key feature of the equilibrium will be determining the location of production of different intermediates. The bundle of intermediates contains a continuum of measure one of inputs, assumed to be imperfectly substitutable with each other, with a constant and symmetric elasticity of substitution equal to ρ . Very little will depend on the particular value of ρ .

We index final-good firms by their 'core productivity', which we denote by φ , and which governs the mapping between the bundle of inputs and final-good production. All intermediates are produced with labor under *firm-specific* technologies featuring constant returns to scale. We denote by

$a_j(v, \varphi)$ firm φ 's unit labor requirement associated with the production of intermediate $v \in [0, 1]$ in country $j \in J$. Intermediates are produced by a competitive fringe of suppliers in each country and thus the free-on-board price paid by firm φ for input v in country j is given by $a_j(v, \varphi) w_j$.⁸

With the above assumptions and notation at hand, we can express the marginal cost for firm φ based in country i of producing a unit of a final-good variety as

$$c_i \left(\{j(v)\}_{v=0}^1, \varphi \right) = \frac{1}{\varphi} \left(\int_0^1 (\tau_{ij(v)} a_{j(v)}(v, \varphi) w_{j(v)})^{1-\rho} dv \right)^{1/(1-\rho)}, \quad (4)$$

where $\{j(v)\}_{v=0}^1$ corresponds to the infinitely-dimensional vector of locations of intermediate input production and $\tau_{ij(v)}$ denotes the iceberg trade costs between the base country i and the production location $j(v)$. For simplicity, we assume that final-good varieties are prohibitively costly to trade across borders, although we will later relax this assumption and briefly study the joint determination of the extensive margins of both exports and imports.

Final-good producers are heterogeneous in their core productivity level φ . Following Melitz (2003), we assume that firms draw their value of φ from a country-specific distribution $g_i(\varphi)$, with support in $[\underline{\varphi}_i, \infty)$, and with an associated continuous cumulative distribution $G_i(\varphi)$.

Building on Eaton and Kortum (2002), we treat the (infinite-dimensional) vector of firm-specific intermediate input efficiencies $a_j(v, \varphi)$ as the realization of an extreme value distribution. More specifically, firm φ draws the value of $a_j(v, \varphi)$ for a given location j from the Fréchet distribution

$$\Pr(a_j(v, \varphi) \leq a) = e^{-T_j a^{-\theta}}, \quad \text{with } T_j > 0. \quad (5)$$

These firm-specific draws are assumed to be independent across locations and inputs. Note that we specify these draws as being independent of the firm's core productivity, φ , but a higher core productivity enables firms to transform the bundle of intermediates into more units of final goods. As in Eaton and Kortum (2002), T_j governs the state of technology in country j , while θ determines the variability of productivity draws across inputs, with a lower θ fostering the emergence of comparative advantage *within* the range of intermediates sector across countries.⁹ It will be convenient to denote by $\mathbf{a}_j(\varphi) \equiv \{a_j(v, \varphi)\}_{v=0}^1$ the vector of of unit labor requirements drawn by firm φ in country j .

Although firms can potentially draw a vector $\mathbf{a}_j(\varphi)$ for *any* country $j \in J$, we assume that a firm from country i only acquires the capability to offshore intermediates to j and learns this vector after incurring a fixed cost equal to f_{ij} units of labor in country i (at a cost $w_i f_{ij}$). We denote by $\mathcal{J}_i(\varphi) \subseteq J$ the set of countries for which a firm based in i with productivity φ has paid the associated fixed cost of offshoring $w_i f_{ij}$. For brevity, we will often refer to $\mathcal{J}_i(\varphi)$ as the *sourcing strategy* of that firm.

⁸Implicitly, we assume that contracts between final-good producers and suppliers are perfectly enforceable, so that the firm-specificity of technology is irrelevant for the prices at which inputs are transacted.

⁹For technical reasons, we assume $\theta > \rho - 1$. Apart from satisfying this restriction, the value of ρ does not matter for any outcomes of interest and will be absorbed into a constant.

Apart from these fixed costs of offshoring, we consider an equilibrium with free entry of final-good producers, in which entry entails a fixed cost equal to f_{ei} units of labor in the country where the bundle of intermediate goods is assembled into the final good (country i in our notation above). As in Melitz (2003), final-good producers only learn their productivity φ after paying the entry cost, but are assumed to choose their sourcing strategy with knowledge of that core productivity level. We will sometimes refer to workers employed in input production as *production workers* and to those employed to cover fixed costs as *non-production workers*. This will facilitate the structural estimation of model in section 5, but is immaterial for the equilibrium, since these different occupations provide the same remuneration for workers.

This completes the description of the key assumptions of the model.

3 Equilibrium

We solve for the equilibrium of the model in three steps. First, we describe optimal firm behavior conditional on a given sourcing strategy $\mathcal{J}_i(\varphi)$. Second, we characterize the choice of this sourcing strategy and relate our results to some of the stylized facts discussed in the Introduction. Third, we aggregate the firm-level decisions and solve for the general equilibrium of the model. We conclude this section by outlining the implications of our framework for bilateral trade across countries and by briefly discussing two extensions of our framework.

3.1 Firm Behavior Conditional on a Sourcing Strategy

Consider a firm based in country i with productivity φ that has incurred all fixed costs associated with a given sourcing strategy $\mathcal{J}_i(\varphi)$. In light of the cost function in (4), it is clear that after drawing the vector $\mathbf{a}_j(\varphi)$ for each country $j \in \mathcal{J}_i(\varphi)$, the firm will choose the location of production of any input v that solves $\min_j \{\tau_{ij} a_j(v, \varphi) w_j\}$. Although such a problem might seem hard to characterize at first glance, we can follow closely the analysis in Eaton and Kortum (2002) to provide a simple characterization of it. In particular, using the properties of the Fréchet distribution in (5), one can show that the firm will source a positive measure of intermediates from each country in its sourcing strategy set $\mathcal{J}_i(\varphi)$. Furthermore, the share of intermediate input purchases sourced from any country j (including the home country i) is simply given by

$$\chi_{ij}(\varphi) = \frac{T_j (\tau_{ij} w_j)^{-\theta}}{\Theta_i(\varphi)} \quad \text{if } j \in \mathcal{J}_i(\varphi) \quad (6)$$

and $\chi_{ij}(\varphi) = 0$ otherwise, where

$$\Theta_i(\varphi) \equiv \sum_{k \in \mathcal{J}_i(\varphi)} T_k (\tau_{ik} w_k)^{-\theta}. \quad (7)$$

The term $\Theta_i(\varphi)$ summarizes the *sourcing capability* of firm φ from i . Note then that each country j 's market share in the firm's purchases of intermediates corresponds to this country's contribution

to this sourcing capability $\Theta_i(\varphi)$. Countries in the set $\mathcal{J}_i(\varphi)$ with lower wages w_j , more advanced technologies T_j , or lower distance from country i are predicted to have higher market shares in the intermediate input purchases of firms based in country i . We shall refer to the term $T_j(\tau_{ij}w_j)^{-\theta}$ as the *sourcing potential* of country j from the point of view of firms in i .

It may seem surprising that the dependence of country j 's market share $\chi_{ij}(\varphi)$ on wages and trade costs is shaped by the Fréchet parameter θ and not by the substitutability across inputs, as governed by the parameter ρ in equation (4). The reason for this, as in Eaton and Kortum (2002), is that variation in market shares is explained exclusively by a *product-level* extensive margin.¹⁰

After choosing the least cost source of supply for each input v , the overall marginal cost faced by firm φ from i can be expressed, after some nontrivial derivations, as

$$c_i(\varphi) = \frac{1}{\varphi} (\gamma \Theta_i(\varphi))^{-1/\theta}, \quad (8)$$

where $\gamma = \left[\Gamma\left(\frac{\theta+1-\rho}{\theta}\right) \right]^{\theta/(1-\rho)}$ and Γ is the gamma function.¹¹ Note that in light of equation (7), the addition of a new location to the set $\mathcal{J}_i(\varphi)$ increases the sourcing capability of the firm and necessarily lowers its effective marginal cost. Intuitively, an extra location grants the firm an additional cost draw for all varieties $v \in [0, 1]$, and it is thus natural that this greater competition among suppliers will reduce the expected minimum sourcing cost per intermediate. In fact, the addition of a country to $\mathcal{J}_i(\varphi)$ lowers the expected price paid for *all* varieties v , and not just for those that are ultimately sourced from the country being added to $\mathcal{J}_i(\varphi)$. This feature of the model distinguishes our framework from Armington-style love-for-variety models, in which the addition of an input location also decreases costs and increases revenue-based productivity, but in which the price paid for a country's variety is unaffected by the inclusion of other countries in the set $\mathcal{J}_i(\varphi)$.¹²

Using the demand equation (2) and the derived marginal cost function in (8), we can express the firm's profits conditional on a sourcing strategy $\mathcal{J}_i(\varphi)$ as

$$\pi_i(\varphi) = \varphi^{\sigma-1} (\gamma \Theta_i(\varphi))^{(\sigma-1)/\theta} B_i - w_i \sum_{j \in \mathcal{J}_i(\varphi)} f_{ij}, \quad (9)$$

where B_i is given in (3). As is clear from equation (9), when deciding whether to add a new country j to the set $\mathcal{J}_i(\varphi)$, the firm trades off the reduction in costs associated with the inclusion of that

¹⁰More specifically, firm φ from i sources a disproportionate measure of inputs from locations in its sourcing strategy $\mathcal{J}_i(\varphi)$ featuring particularly favorable production costs (as summarized by $T_j(\tau_{ij}w_j)^{-\theta}$). This oversampling in turn puts upwards pressure on the average price that firm φ from i ends up paying for inputs purchased from those attractive locations. As it turns out, these two effects on prices (low costs but oversampling) exactly cancel each other to the point at which the distribution of input prices paid by firm φ from i in any country $j \in \mathcal{J}_i(\varphi)$ is identical. As a result, variation in market shares simply reflects the fact that firms buy more inputs from more attractive sourcing locations.

¹¹These derivations are analogous to those performed by Eaton and Kortum (2002) to solve for the aggregate price index in their model of trade in final goods.

¹²See, among other, Halpern et al. (2011), Goldberg et al. (2010), and Gopinath and Neiman (2013). The cost function in (8) and profit function below in (9) are actually isomorphic (up to a scalar) to those derived from an Armington model in which inputs are differentiated by country of origin and their elasticity of substitution is constant and given by $(1 + \theta)/\theta$. We will empirically defend our modeling choices in section 4.1.

country in the set $\mathcal{J}_i(\varphi)$ – which increases the sourcing capability $\Theta_i(\varphi)$ – against the payment of the additional fixed cost $w_i f_{ij}$. We next turn to studying the optimal determination of the sourcing strategy $\mathcal{J}_i(\varphi)$.

3.2 Optimal Sourcing Strategy

Each firm’s optimal sourcing strategy is a combinatorial optimization problem in which a set $\mathcal{J}_i(\varphi) \subseteq J$ of locations is chosen to maximize the firm’s profits $\pi_i(\varphi)$ in (9). We can alternatively express this problem as

$$\max_{I_{ij} \in \{0,1\}_{j=1}^J} \pi_i(\varphi, I_{i1}, I_{i2}, \dots, I_{iJ}) = \varphi^{\sigma-1} \left(\gamma \sum_{j=1}^J I_{ij} T_j (\tau_{ij} w_j)^{-\theta} \right)^{(\sigma-1)/\theta} B_i - w_i \sum_{j=1}^J I_{ij} f_{ij}, \quad (10)$$

where the indicator variable I_{ij} takes a value of 1 when $j \in \mathcal{J}_i(\varphi)$ and 0 otherwise. In theory, with knowledge of the small number of parameters that appear in (10), this problem could be solved computationally by calculating firm profits for different combinations of locations and picking the unique strategy yielding that highest level of profits. Nevertheless, this would amount to computing profits for 2^J possible strategies, which is clearly infeasible unless one chooses a small enough set J of candidate countries. In section 5, we will discuss alternative approaches to circumvent this dimensionality problem when estimating some of the parameters of the model. For the time being, we focus on providing a characterization of some key analytic features of the solution to this optimal sourcing strategy problem.

The problem in (10) is not straightforward to solve because the decision to include a country j in the set $\mathcal{J}_i(\varphi)$ depends on the number and characteristics of the other countries in this set. That dependence is in turn crucially shaped by whether $(\sigma - 1)/\theta$ is higher or lower than one. When $(\sigma - 1)/\theta > 1$, the profit function $\pi_i(\varphi)$ features increasing differences in (I_{ij}, I_{ik}) for $j, k \in \{1, \dots, J\}$ and $j \neq k$, and as a result the marginal gain from adding a new location to $\mathcal{J}_i(\varphi)$ is increasing in the number of elements in that set. This case is more likely to apply whenever demand is elastic and thus profits are particularly responsive to variable cost reductions (high σ), and whenever input efficiency levels are relatively heterogeneous across markets (low θ), so that the expected reduction in costs achieved by adding an extra country into the set of active locations is relatively high. In the converse case in which $(\sigma - 1)/\theta < 1$, we instead have that firm profits feature decreasing differences in (I_{ij}, I_{ik}) , and the marginal gain from adding a new location to $\mathcal{J}_i(\varphi)$ is decreasing in the number of elements in that set. We shall refer to the case $(\sigma - 1)/\theta > 1$ as the *complements* case, and to the case $(\sigma - 1)/\theta < 1$ as the *substitutes* case.

Regardless of parameter values, we can use the fact that the profit function (10) is supermodular in φ and $\Theta_i(\varphi)$ to establish that:

Proposition 1. The solution $I_{ij}(\varphi) \in \{0,1\}_{j=1}^J$ to the optimal sourcing problem (10) is such that a firm’s sourcing capability $\Theta_i(\varphi) = \sum_{j=1}^J I_{ij}(\varphi) T_j (\tau_{ij} w_j)^{-\theta}$ is nondecreasing in φ .

Although this is not the focus of our analysis, it is worth pointing out that Proposition 1 has interesting implications for the size distribution of firms as implied by our model. Note that firm sales are given by a multiple σ of operating profits or

$$R_i(\varphi) = \sigma \varphi^{\sigma-1} (\gamma \Theta_i(\varphi))^{(\sigma-1)/\theta} B_i. \quad (11)$$

If the sourcing capability of firms were independent of φ , then the size distribution of firms would be governed by the distribution of the term $\varphi^{\sigma-1}$. Nevertheless, with $\Theta_i(\varphi)$ being nondecreasing in φ , the equilibrium size distribution of firms will feature more positive skewness than as implied by the distribution of $\varphi^{\sigma-1}$. For instance, if $G_i(\varphi)$ is a log-normal distribution with mean μ and variance σ^2 , $\varphi^{\sigma-1}$ is also distributed log-normal with mean $(\sigma-1)\mu$ and variance $(\sigma-1)^2\sigma^2$. Yet, Proposition 1 suggests that the actual firm size distribution will exhibit a fatter right tail than as predicted by a log-normal distribution estimated on the whole population of firms. As a result, the size distribution of the relatively large firms may be better approximated by distributions featuring higher positive skewness than the log-normal, such as the Pareto distribution. These implications of the model are broadly consistent with available empirical evidence on the size distribution of firms (see, for instance, Rossi-Hansberg and Wright, 2007, particularly their Figure 1).

It is important to emphasize that the result in Proposition 1 that a firm's sourcing capability is increasing in productivity does *not* imply that the extensive margin of sourcing at the firm level (i.e., the number of elements of $\mathcal{J}_i(\varphi)$) is necessarily increasing in firm productivity as well. For example, a highly productive firm from i might pay a large fixed cost to be able to offshore to a country j with a particularly high sourcing potential (i.e., a high value of $T_j(\tau_{ij}w_j)^{-\theta}$) – thus greatly increasing Θ_i – after which the marginal incentive to add further locations might be greatly diminished in the substitutes case, that is whenever $(\sigma-1)/\theta < 1$. Under these parameter conditions, it may not be profitable for a low productivity firm to pay a large fixed cost to access a high sourcing potential country. Instead, the low productivity firm may source from two countries with relatively low measures of sourcing potential but also with lower fixed costs.¹³

The evidence in Figure 1 in the Introduction suggests, however, that the extensive margin of offshoring *does* appear to respond positively to productivity. Furthermore, as we shall show in section 4.1 below, there also exists some evidence of the existence of a hierarchical structure in the sourcing location decisions of U.S. firms. In terms of our model, this would amount to the sourcing strategy $\mathcal{J}_i(\varphi)$ of relatively unproductive firms being a strict subset of the sourcing strategy of relatively productive firms, or $\mathcal{J}_i(\varphi_L) \subseteq \mathcal{J}_i(\varphi_H)$ for $\varphi_H \geq \varphi_L$. As we show in the Appendix, this strong pattern is in fact implied by our model in the complements case, and thus we have that:

Proposition 2. Whenever $(\sigma-1)/\theta > 1$, the solution $I_{ij}(\varphi) \in \{0, 1\}_{j=1}^J$ to the optimal sourcing problem (10) is such that $\mathcal{J}_i(\varphi_L) \subseteq \mathcal{J}_i(\varphi_H)$ for $\varphi_H \geq \varphi_L$, where $\mathcal{J}_i(\varphi) = \{j : I_{ij}(\varphi) = 1\}$.

¹³The profit function exhibits decreasing differences in a firm's sourcing strategy whenever $\sigma-1 < \theta$. As a result, this restriction on the parameter values may lead to non-hierarchical sourcing patterns (i.e., low productivity firms source from countries from high productivity firms do not source) when countries with high sourcing potential also have relatively high fixed costs. The evidence presented in Table 1 suggests this is case, at least for some countries.

In the complements case, the model thus delivers a ‘pecking order’ in the extensive margin of offshoring which is reminiscent of the one typically obtained in models of exporting with heterogeneous firms, such as Eaton et al. (2011). Another implication of Proposition 2 is that, for firms with a sufficiently low value of core productivity φ , the set $\mathcal{J}_i(\varphi)$ may be a singleton, and the associated unique profitable location j of input production will necessarily be the one associated with the highest ratio $T_j(\tau_{ij}w_j)^{-\theta}/f_{ij}$. This in turn implies that if fixed costs of offshoring are disproportionately large relative to fixed costs of domestic sourcing, so that $f_{ii} \ll f_{ij}$ for any $j \neq i$, the model is consistent with the observed superior performance of firms engaged in offshoring relative to firms sourcing exclusively in the domestic economy.

The characterization results in Propositions 1 and 2 are useful for interpreting some of the stylized facts regarding the extensive margin of offshoring discussed in the Introduction. It should also be clear that if countries only differ along one dimension (either sourcing potential or fixed costs), then they can be uniquely ranked in terms of their ‘sourcing appeal’ for all firms. When countries differ along these two dimensions – as suggested by Table 1 – then a particular country’s rank may be different for high versus low productivity firms. It follows from Proposition 2, however, that under the complements case, if one were to be able to rank countries by an index of their ‘sourcing appeal’, the measure of firms from i sourcing from a given country would be increasing in that index, and the computationally intensive problem (10) could be greatly simplified. Although it is obvious from (10) that such an index of sourcing appeal for a country j should be increasing in $T_j(\tau_{ij}w_j)^{-\theta}$ and decreasing in f_{ij} , it is less clear how *exactly* such an index should aggregate these variables.

We can obtain sharper characterizations of the solution to the sourcing strategy problem in (10) by making additional assumptions. Consider first a situation in which the fixed costs of offshoring are common for all foreign countries, so $f_{ij} = f_{iO}$ for all $j \neq i$. In such a case, and regardless of the value of $(\sigma - 1)/\theta$, one could then rank foreign locations $j \neq i$ according to their sourcing potential $T_j(\tau_{ij}w_j)^{-\theta}$ and denote by $i_r = \{i_1, i_2, \dots, i_{J-1}\}$ the country with the r -th highest value of $T_j(\tau_{ij}w_j)^{-\theta}$. Having constructed i_r , it then follows that for any firm with productivity φ from i that offshores to at least one country, $i_1 \in \mathcal{J}_i(\varphi)$; for any firm that offshores to at least two countries, we have $i_2 \in \mathcal{J}_i(\varphi)$, and so on. In other words, not only does the extensive margin increase monotonically with firm productivity, but it does so in a manner uniquely determined by the ranking of the $T_j(\tau_{ij}w_j)^{-\theta}$ sourcing potential terms. It is important to emphasize that this result holds both in the complements case as well as in the substitutes case, though again it relies on the assumption of identical offshoring fixed costs across sourcing countries, an assumption that appears particularly dubious in light of the evidence documented in Table 1 in the Introduction.

Even in the presence of cross-country differences in the fixed costs of offshoring, a similar sharp result emerges in the knife-edge case in which $(\sigma - 1)/\theta = 1$. In that case, the addition of an element to the set $\mathcal{J}_i(\varphi)$ has no effect on the decision to add any other element to the set, and the same pecking order pattern described in the previous paragraph applies, but when one ranks foreign locations according to the ratio $T_j(\tau_{ij}w_j)^{-\theta}/f_{ij}$ rather than $T_j(\tau_{ij}w_j)^{-\theta}$. This result is

analogous to the one obtained in standard models of selection into exporting featuring constant marginal costs, in which the decision to service a given market is independent of that same decision in other markets.

We can also use the properties of the profit function to derive a few firm level comparative statics that hold constant the market demand level B_i . First, and quite naturally, a reduction in any iceberg trade cost τ_{ij} or fixed cost of sourcing f_{ij} (weakly) increases the firm's sourcing capability $\Theta_i(\varphi)$ and thus firm-level profits. Second, in the complements case, a reduction of any τ_{ij} or f_{ij} also (weakly) increase the extensive margin of global sourcing, in the sense that the set $\mathcal{J}_i(\varphi)$ is nondecreasing in τ_{ij} and f_{ij} for any j . Finally, and perhaps more surprisingly, in the complements case it is also the case that a reduction of *any* τ_{ij} or f_{ij} increases (weakly) firm-level bilateral input purchases from *all* countries. Intuitively, in such a case, complementarities are strong enough to dominate the direct substitution effect related to market shares shifting towards the locations whose costs of sourcing have been reduced. It should be emphasized, however, that these results apply when holding the level of B_i fixed, when in fact it would be affected by these same parameter changes. As we shall see in our counterfactual exercises in section 6, the endogenous response of B_i appears to be quantitatively important in our estimation, and thus the implications we derive from changes in trade costs are much more nuanced than those discussed above.

3.3 Industry and General Equilibrium

Consider now the general equilibrium of the model. As mentioned before, we will simplify matters by assuming that consumers spend a constant share (which we denote by η) of their income in manufacturing. The remaining share $1 - \eta$ of income is spent on a perfectly competitive non-manufacturing sector that competes for labor with manufacturing firms. Technology in that sector is linear in labor and we assume that $1 - \eta$ is large enough to guarantee that the wage rate w_i in each country i is pinned down by labor productivity in that sector. We thus can treat wages as exogenous in solving for the equilibrium in each country's manufacturing sector. For simplicity, we also assume that this 'outside' sector's output is homogeneous, freely tradable across countries, and serves as a numeraire in the model.

We next turn to describing the equilibrium in the manufacturing sector and in particular of the market demand term B_i . Given our assumption that final-good producers only observe their productivity after paying the fixed cost of entry, we can use equation (9) to express the free-entry condition in manufacturing as

$$\int_{\tilde{\varphi}_i}^{\infty} \left[\varphi^{\sigma-1} (\gamma \Theta_i(\varphi))^{(\sigma-1)/\theta} B_i - w_i \sum_{j \in \mathcal{J}_i(\varphi)} f_{ij} \right] dG_i(\varphi) = w_i f_{ei}. \quad (12)$$

In the lower bound of the integral, $\tilde{\varphi}_i$ denotes the productivity of the least productive active firm in country i . Firms with productivity $\varphi < \tilde{\varphi}_i$ cannot profitably source from any country and thus exit upon observing their productivity level. Note that B_i affects expected operating profits both

directly via the explicit term on the left-hand-side of (12), but also indirectly through its impact on the determination of $\tilde{\varphi}_i$, $\mathcal{J}_i(\varphi)$ and $\Theta_i(\varphi)$. Despite these rich effects (and the fact that the set $\mathcal{J}_i(\varphi)$ is not easily determined), in the Appendix we show that one can appeal to monotone comparative statics arguments to prove that:

Proposition 3. Equation (12) delivers a unique market demand level B_i for each country $i \in J$.

This result applies both in the complements case as well as in the substitutes case and ensures the existence of a unique industry equilibrium. In particular, the firm-level combinatorial problem in (10) delivers a unique solution given a market demand B_i and exogenous parameters (including wages). Furthermore, the equilibrium measure N_i of entrants in the industry is easily solved from equations (3) and (12), by appealing to the marginal cost in (8), to constant-mark-up pricing, and to the fact that spending E_j in manufacturing is a share η of (labor) income. This delivers:

$$N_i = \frac{\eta L_i}{\sigma \left(\int_{\tilde{\varphi}_i}^{\infty} \sum_{j \in \mathcal{J}_i(\varphi)} f_{ij} dG_i(\varphi) + f_{ei} \right)}. \quad (13)$$

With this expression at hand, the equilibrium number of active firms is simply given by $N_i [1 - G_i(\tilde{\varphi}_i)]$.¹⁴

3.4 Gravity

In this section we explore the implications of our model for the aggregate volume of bilateral trade in manufacturing goods across countries. Because we have assumed that final goods are nontradable, we can focus on characterizing aggregate intermediate input trade flows between any two countries i and j . Given that firm spending on intermediate inputs constitutes a share $(\sigma - 1)/\sigma$ of revenue for all firms, we can use equation (11) and aggregate across firms, to obtain aggregate manufacturing imports from country j by firms based in i :

$$M_{ij} = (\sigma - 1) N_i B_i \int_{\tilde{\varphi}_{ij}}^{\infty} \chi_{ij}(\varphi) (\gamma \Theta_i(\varphi))^{(\sigma-1)/\theta} \varphi^{\sigma-1} dG_i(\varphi). \quad (14)$$

In this expression, $\tilde{\varphi}_{ij}$ denotes the productivity of the least productive firm from i offshoring to j , while $\chi_{ij}(\varphi)$ is given in (6) for $j \in \mathcal{J}_i(\varphi)$ and by $\chi_{ij}(\varphi) = 0$ otherwise. We next re-express equation (11) so that it is comparable to gravity equations used in empirical analyses. With that goal in mind, we begin with two special cases that permit a comparison of the implications of our model for the structure of trade flows with those of some recent models of trade with productivity heterogeneity.

Universal Importing Consider first consider the case in which the fixed costs of offshoring are low enough to ensure that *all* firms acquire the capability to source inputs from *all* countries.

¹⁴In the Online Appendix, we show that in the complements case, and when φ is distributed Pareto with shape parameter κ , we can further reduce equation (13) to $N_i = (\sigma - 1) \eta L_i / (\sigma \kappa f_{ei})$. In such a case, the measure of entrants is independent of trade costs.

This is obviously counterfactual in light of the stylized facts in the Introduction, but studying this unrealistic benchmark environment will prove to be useful below. To simplify matters, we abstract from fixed costs of sourcing altogether and set them to zero.

When all firms import from everywhere, the optimal sourcing strategy of all firms in i is simply given by $\mathcal{J}_i(\varphi) = \{1, 2, \dots, J\}$, and both $\Theta_i = \sum_{k \in J} T_k (\tau_{ik} w_k)^{-\theta}$ and thus χ_{ij} are independent of φ . This allows one to greatly simplify the equilibrium values for B_i and N_i in (12) and (13). Plugging these values into (14), we can express aggregate manufacturing imports from country j by firms based in i as

$$M_{ij} = \frac{E_i}{\Theta_i} (\tau_{ij})^{-\theta} \frac{Q_j}{\sum_k \frac{E_k}{\Theta_k} (\tau_{jk})^{-\theta}}. \quad (15)$$

where remember that E_i equals country i 's total spending in manufacturing goods (which itself is a multiple $\sigma/(\sigma - 1)$ of country i 's worldwide absorption of intermediate inputs) and $Q_j = \sum_k M_{kj}$ denotes the total production of intermediate inputs in country j . Equation (15) is a standard gravity equation relating bilateral trade flows to bilateral trade barriers τ_{ij} , the importer country's total absorption in manufacturing, the exporter country's total production of tradable manufacturing goods, and multilateral resistance terms. The latter are reflected in Θ_i (which is a negative transform of the importing country's ideal price index) and in the summation term in the denominator of the second term.¹⁵

Notice that equation (15) structurally justifies the use of an empirical log-linear specification for bilateral trade flows with importer and exporter asymmetric fixed effects and measures of bilateral trade frictions τ_{ij} . Furthermore, the model indicates that, in the absence of selection into importing, the elasticity of trade flows with respect to changes in these bilateral trade frictions is shaped by the Fréchet parameter θ , just as in the Eaton and Kortum (2002) framework. The intuition for this result is analogous to our earlier discussion of the effects of bilateral trade frictions on the cross-section of import purchases at the firm level. This should not be surprising since, in the absence of selection into offshoring, all firms buy inputs from all markets according to the same market shares χ_{ij} .

How does the introduction of fixed costs of sourcing and selection into offshoring affect the gravity specification in (15)? In order to gain intuition it is convenient to first consider the knife-edge case in which $\sigma - 1 = \theta$, and thus the extensive margin of global sourcing can be studied independently across countries.

Independent Entry Decisions Whenever $\sigma - 1 = \theta$, one can use the formula for χ_{ij} in (6) to simplify (14) to ¹⁶

$$M_{ij} = (\sigma - 1) N_i B_i \gamma T_j (\tau_{ij} w_j)^{-\theta} \int_{\tilde{\varphi}_{ij}}^{\infty} \varphi^{\sigma-1} dG_i(\varphi). \quad (16)$$

¹⁵The ideal price index in country i is given by $P_i^{1-\sigma} = N_i \int_{\tilde{\varphi}_i}^{\infty} p_i(\varphi)^{1-\sigma} dG_i(\varphi)$. In the absence of selection into importing, this reduces to $P_i = \left(\frac{\sigma}{\sigma-1}\right) (N_i)^{1/(1-\sigma)} (\gamma \Theta_i)^{-1/\theta}$, where N_i in (13) simplifies to $N_i = \eta L_i / \sigma f_{ei}$.

¹⁶To derive this equation we appeal to the fact that the extensive margin of sourcing features a hierarchical structure in this case (see the discussion following Proposition 2), and thus $\chi_{ij} > 0$ for all $\varphi > \tilde{\varphi}_{ij}$.

To the extent that a reduction in bilateral trade costs between i and j generates an increase in the measure of firms from i sourcing in j , this expression illustrates that the elasticity of bilateral trade flows with respect to τ_{ij} will now be higher than the firm-level one, i.e., θ . In fact, as we show in the Online Appendix, when assuming that firms draw their core productivity from a Pareto distribution with shape parameter $\kappa > \sigma - 1 = \theta$, we can express aggregate manufacturing imports from country j by firms based in i as

$$M_{ij} = \frac{(E_i)^{\kappa/(\sigma-1)}}{\Psi_i} (\tau_{ij})^{-\kappa} (f_{ij})^{1-\kappa/(\sigma-1)} \frac{Q_j}{\sum_k \frac{(E_k)^{\kappa/(\sigma-1)}}{\Psi_k} (\tau_{kj})^{-\theta} (f_{kj})^{1-\kappa/(\sigma-1)}}, \quad (17)$$

where

$$\Psi_i = \frac{f_{ei}}{L_i} \varphi_i^{-\kappa} P_i^{-\kappa} w_i^{\kappa/(\sigma-1)-1}.$$

Although equation (17) differs in some respects from equation (15), it continues to be a well-defined gravity equation in which the ‘trade elasticity’ (i.e., the elasticity of trade flows to variable trade costs) can be recovered from a log-linear specification that includes importer and exporter fixed effects. A key difference in (17) relative to (15) is that this trade elasticity κ is now predicted to be higher than the one obtained when the model features no extensive margin of importing at the country level.

It may be surprising that the Fréchet parameter θ , which was key in governing the ‘trade elasticity’ (i.e., the elasticity of trade flows to variable trade costs) at the firm level, is now irrelevant when computing that same elasticity at the aggregate level. To understand this result it is useful to relate our framework to the multi-country versions of the Melitz model in Chaney (2008), Arkolakis et al. (2008) or Helpman et al. (2008). In those models, firms pay fixed costs of exporting to obtain additional operating profit flows proportional to $\varphi^{\sigma-1}$ that enter linearly and separably in the firm’s profit function. Even though in our model, selection into offshoring increases firm profits by reducing effective marginal costs, whenever $\sigma - 1 = \theta$, the gain from adding a new market is strictly separable in the profit function and also proportional to $\varphi^{\sigma-1}$. Hence, this effect is isomorphic to a situation in which the firm obtained additional sale revenue by selecting into exporting. It is thus not surprising that the gravity equation we obtain in (17) is essentially identical to those obtained by Chaney (2008) or Arkolakis et al. (2008).

General Case We finally discuss the implications of our framework for bilateral trade flows in the presence of fixed costs of offshoring and interdependencies in the extensive margin of sourcing. Following the same steps as in the derivation of (15), we find that in the general case the volume of imports from country j by firms from i is given by

$$M_{ij} = \frac{E_i}{P_i^{1-\sigma}/N_i} \tau_{ij}^{-\theta} \Lambda_{ij} \frac{Q_j}{\sum_k \frac{E_k}{P_k^{1-\sigma}/N_k} \tau_{kj}^{-\theta} \Lambda_{kj}} \quad (18)$$

where

$$\Lambda_{ij} = \int_{\tilde{\varphi}_{ij}}^{\infty} I_{ij}(\varphi) (\Theta_i(\varphi))^{(\sigma-1-\theta)/\theta} \varphi^{\sigma-1} dG_i(\varphi). \quad (19)$$

As in the previous two gravity equations in (15) and (17), the first and last terms in equation (18) once again constitute importer and an exporter ‘fixed effects’, which relate to the importer’s total absorption and the exporter’s total tradable production in manufacturing. Notice, however, that even after partialling out importer and exporter fixed effects, in equation (18) we are left with both τ_{ij} as well as Λ_{ij} in (19) varying across exporters even when fixing the importer.

The presence of the term Λ_{ij} has two implications for standard estimates of the gravity equation. First, and analogously to the independent entry decisions case analyzed above, the aggregate trade elasticity will no longer coincide with the firm-level one, given by θ . Intuitively, changes in variable trade costs will not only affect firm-level sourcing decisions conditional on a sourcing strategy, but will also affect these same sourcing strategies. In the plausible case in which a reduction in τ_{ij} disproportionately enhances the extensive margin of imports from country j , the aggregate trade elasticity will thus be higher than θ . Our derivations above have demonstrated this to be the case whenever $\sigma - 1 = \theta$. Unfortunately, a general proof of this magnification result for arbitrary parameters σ and θ and for a general distribution of productivity $G_i(\varphi)$ is intricate due to the difficulties in the characterization of $\Theta_i(\varphi)$ in the substitutes case and due to industry equilibrium effects.

A second distinguishing feature of equation (18) is that, as long as $(\sigma - 1) \neq \theta$, Λ_{ij} in (19) will be a function of $\Theta_i(\varphi)$ for $\varphi > \tilde{\varphi}_{ij}$, and will thus depend on technology, trade costs and wages in *all* countries, and not just i and j . Hence, equation (18) is an *extended* gravity equation – to use the term in Morales et al. (2014) – featuring third market effects. Holding constant the sourcing strategy of all firms (and thus $\tilde{\varphi}_{ij}$ and $I_{ij}(\varphi)$ in equation (19)), it appears that the sign of these third-market effects depends crucially on whether one is in the complements or the substitutes case. Nevertheless, changes in trade costs naturally affect the extensive margin of sourcing and also lead to rich industry equilibrium effects, thereby thwarting a sharp characterization of these extended gravity effects in our model.

3.5 Extensions

In this section, we briefly outline two extensions of the model that illustrate how our framework can accommodate additional prominent patterns of the involvement of U.S. firms in international trade transactions. Because we will not incorporate these features into the structural estimation and quantitative analysis in the next sections, we will limit ourselves to studying the effects of these extensions on firm behavior, and not on the aggregate implications of the model.

A. Tradable Final Goods: Exporting and Importing

We have assumed so far that final-good varieties are prohibitively costly to trade across borders. We have done so to focus our analysis on the determinants and implications of selection into

global sourcing. In this section, we briefly relax this assumption and demonstrate the existence of intuitive complementarities between the extensive margin of exporting and that of importing at the firm level.

Suppose then that trade in final-varieties is only partially costly and involves both iceberg trade costs τ_{ij}^X as well as fixed costs f_{ij}^X of exporting. Firm behavior conditional on a sourcing strategy is largely analogous to that in section 3.1. In particular, after observing the realization of its supplier-specific productivity shocks, each final-good producer will continue to choose the location of production for each input to minimize costs, which will lead to the same marginal cost function $c_i(\varphi)$ obtained above in equation (8). The main novelty is that the firm will now produce output not only for the domestic market but also for a set of endogenously chosen foreign markets, which constitute the firm's 'exporting strategy'. We can then express the problem of determining the optimal exporting and sourcing strategies of a firm from country i with core productivity φ as:

$$\begin{aligned} \max_{\substack{I_{ij}^M \in \{0,1\}_{j=1}^J \\ I_{ik}^X \in \{0,1\}_{k=1}^J}} \pi_i(\varphi, \mathbf{I}^M, \mathbf{I}^X) &= \varphi^{(\sigma-1)} \left(\gamma \sum_{j=1}^J I_{ij}^M T_j (\tau_{ij} w_j)^{-\theta} \right)^{(\sigma-1)/\theta} \sum_{k=1}^J I_{ik}^X (\tau_{ik}^X)^{1-\sigma} B_k \\ &\quad - w_i \sum_{j=1}^J I_{ij}^M f_{ij} - w_i \sum_{k=1}^J I_{ik}^X f_{ij}^X, \end{aligned}$$

Note that \mathbf{I}^M and \mathbf{I}^X denote the vector of extensive margin import and export decisions, respectively. It is straightforward to see that, whenever $(\sigma - 1) / \theta > 1$, this more general profit function continues to feature increasing differences in (I_j^M, I_k^M) for $j, k \in \{1, \dots, J\}$ with $j \neq k$, and also features increasing differences in (I_j^M, φ) for any $j \in \{1, \dots, J\}$. As a result, Proposition 2 continues to apply here and we obtain a 'pecking order' in the extensive margin of offshoring in the complements case.

The key new feature of the above profit function $\pi_i(\varphi, \mathbf{I}^M, \mathbf{I}^X)$ is that it also exhibits increasing differences in (I_j, X_k) for any $j, k \in \{1, \dots, J\}$ and increasing differences in (X_j, φ) for any $j \in \{1, \dots, J\}$. This has at least two implications. First, regardless of whether $\sigma - 1 > \theta$ or $\sigma - 1 < \theta$, any change in parameters that increases the sourcing capability $\Theta_i(\varphi)$ of the firm – such as reduction in any τ_{ij} or an increase in any T_j – will necessarily lead to a (weak) increase in the vector \mathbf{I}^X , and thus (weakly) increase the export margin of exporting. Second, restricting attention to the complements case $(\sigma - 1) / \theta > 1$, the model delivers a complementarity between the exporting and importing margins of firms. For instance, holding constant the vector of residual demand parameters B_i , reductions in the costs of trading final goods across countries will not only increase the participation of firms in export markets, but will also increase the extensive margin of sourcing, in the sense that vector \mathbf{I}^M is non-increasing in τ_{ik}^X . Furthermore, as firm productivity increases, the participation of firms in both export and import markets increases, and at a faster rate than when the export margin is shut down. This complementarity result is useful in interpreting the

fact that, as Bernard et al. (2007) indicate, 41 percent of U.S. exporting firms also import while 79 percent of importers also export. Although it would be interesting to explore the aggregate implications of this extended framework, we shall not do so in this paper because our quantitative exercises in the next section do not incorporate these features into the analysis.

B. Endogenous Input Variety

Our benchmark model assumes that all final good producers use a measure one of inputs. We next briefly outline how our results extend and generalize to the case in which the final-good producer is allowed to choose the complexity of production, as captured by the measure of inputs used in production (see Acemoglu et al. (2007)). As we shall see, this ends up producing an equilibrium essentially identical to the one we have described above but with additional implications for how the measure of inputs purchased by firms changes with firm productivity.

The formal details of this extension are as follows. Final-good production continues to combine inputs according to a CES technology but we now let the measure of inputs be firm-specific and given by $n_i(\varphi)$. More specifically, we generalize the marginal cost function in (4) as follows:

$$c_i(\{j(v)\}_{v=0}^1, \varphi) = \frac{1}{\varphi} n_i(\varphi)^{1/(\rho-1)-\lambda} \left(\int_0^{n_i(\varphi)} (\tau_{ij(v)} a_{j(v)}(v, \varphi) w_{j(v)})^{1-\rho} dv \right)^{1/(1-\rho)}.$$

A higher value of $n_i(\varphi)$ enhances productivity via an input variety effect. As in Benassy (1998) and Acemoglu et al. (2007), we introduce the term $n_i(\varphi)^{1/(\rho-1)-\lambda}$ in front of the integral in order to control the importance of variety effects for productivity via a parameter λ disentangled from the elasticity substitution between inputs ρ . In order to create a check on the optimal degree of complexity, we assume that firms face a fixed cost equal to $w_i n_i(\varphi) f_i^n$ when combining $n_i(\varphi)$ inputs in production. As in our benchmark model, in each of the countries in which the final-good producer incurred the fixed cost of sourcing, there is a competitive fringe of potential suppliers that can provide differentiated inputs to the firm with a firm-specific intermediate input efficiencies drawn from a Fréchet distribution.

With a continuum of inputs, the equilibrium measure of inputs used in production by a final-good producer has no implications for the distribution of input prices faced by that producer. Exploiting this feature, we can use derivations analogous to those in the benchmark model and in Eaton and Kortum (2002), to write the marginal cost of production as

$$c_i(\varphi) = \frac{1}{\varphi} (n_i(\varphi))^{-\lambda} (\gamma \Theta_i(\varphi))^{-1/\theta}, \quad (20)$$

and the firm's profits conditional on a sourcing strategy $\mathcal{J}_i(\varphi)$ as

$$\pi_i(\varphi) = \varphi^{\sigma-1} (n_i(\varphi))^{(\sigma-1)\lambda} (\gamma \Theta_i(\varphi))^{(\sigma-1)/\theta} B_i - w_i \sum_{j \in \mathcal{J}_i(\varphi)} f_{ij} - w_i n_i(\varphi) f_i^n,$$

where B_i is again given in (3). It is clear that conditional on a sourcing strategy $\mathcal{J}_i(\varphi)$ – and thus a value of $\Theta_i(\varphi)$ –, this profit function is supermodular in productivity and the measure of inputs $n_i(\varphi)$.¹⁷ Hence, a novel prediction from this extension is that more productive firms will tend to source more inputs from all sources combined (domestic and foreign) than less productive firms, even when these firms share a common sourcing strategy.¹⁸ In the complements case with $\sigma - 1 > \theta$, this variant of the model also predicts that more productive firms will tend to buy (weakly) more inputs from *any* source than less productive firms.

Although the inclusion of endogenous input variety and fixed costs of sourcing at the input level might help rationalize certain features in the data, it is important to emphasize that they do not serve as a substitute for country-specific fixed costs of sourcing. More specifically, in the absence of the latter type of fixed costs, our framework would not be able to account for the key facts motivating our benchmark model, since in such a case, all firms would source inputs from all countries, thus violating the patterns in Figure 1 and Table 1 in the Introduction.

4 Data Sources and Descriptive Evidence

In the theory section, we provide a parsimonious model that characterizes the margins of firms’ global sourcing decisions. When there are complementarities in the firm’s extensive margin sourcing decisions, the model is consistent with the strong, increasing relationship between firm size and the number of source countries depicted in Figure 1. The model also provides a framework for distinguishing between country-level fixed costs and country sourcing potential – two key dimensions along which Table 1 suggests that countries differ. Before turning to the structural estimation, we describe the data used in the paper and provide several novel empirical facts that support the theoretical framework.

4.1 Data Description

The primary data used in the paper are from the U.S. Census Bureau’s 2007 Economic Censuses (EC), Longitudinal Business Database (LBD), and Import transaction database. The LBD uses administrative record data to provide employment and industry for every private, non-farm, employer establishment in the U.S. The ECs supplement this information with additional establishment-level variables, such as sales, value-added, and input usage.¹⁹ The import data, collected by U.S. Customs facilities, are based on the universe of import transactions into the U.S. They contain information

¹⁷For the choice of $n_i(\varphi)$ to satisfy the second-order conditions for a maximum, we need to impose that the efficiency gains from input variety are small enough to guarantee that $(\sigma - 1)\lambda < 1$ holds.

¹⁸Although our benchmark model is also consistent with more productive firms importing more inputs than less productive firms, with a common measure of inputs, this could only be rationalized by having more productive firms sourcing less inputs domestically than less productive firms.

¹⁹The Census of Manufactures (CM) has been widely used in previous work. The other censuses are for Construction, Finance, Insurance and Real Estate, Management of Companies, Professional and Technical Services, Retail Trade, Transportation and Warehousing, and Wholesale Trade. The variables available differ across these censuses. This coverage ensures that we provide an accurate depiction of the entire firm compared to studies that rely solely on the CM.

on the products, values, countries, and related party status of firms' imports. We match these data at the firm level to LBD and the EC data.

The focus of this paper is on firms involved in the production of goods. We therefore limit the analysis to firms with at least one manufacturing establishment. Because we envision a production process entailing physical transformation activities (manufacturing) as well as headquarter activities (design, distribution, marketing, etc.), we include firms with activities outside of manufacturing.²⁰ We also limit the sample to firms with positive sales and employment and exclude all mineral imports from the analysis since they do not represent offshoring. Firms with at least one manufacturing plant account for five percent of firms, 23 percent of employment, 38 percent of sales, and 65 percent of non-mineral imports. In terms of explaining aggregate U.S. sourcing patterns, it is critically important to include firms with manufacturing and other activities. They account for 60 percent of U.S. imports, while manufacturing-only firms account for just five percent. The import behavior of the firms in our sample is consistent with patterns documented in past work on heterogeneous firms in trade. Only one quarter of U.S. manufacturing firms have positive imports in 2007. Additional details on the sample and data construction are in the Data Appendix.

Equation (6) provides a clear formulation of the share of intermediate inputs sourced by a firm in country i from country j . In the next section, we exploit this relationship to estimate country sourcing potential. To do so, we construct a measure of a firm's total intermediate input purchases based on the difference between the firm's sales and its value added, adjusting for changes in its inventories of final goods and materials. This approach ensures a more complete metric of a firm's inputs than traditional measures based purely on manufacturers' use of materials because it takes into account the input usage of both the manufacturing as well as the wholesale establishments of U.S. firms.²¹ The model does not take a stance on whether intermediate inputs are sourced within or across firm boundaries. For the purposes of this paper, this is of little relevance for international transactions, but it might lead to important biases in our measure of overall input use if a significant share of domestic inputs is produced within the firm and is recorded as value added. For this reason, we add a firm's total production-worker wage bill to the difference between the firm's sales and value added. In terms of our model, this corresponds to assuming that the final-good producer employs production workers to manufacture internally any inputs produced by

²⁰We recognize that focusing on firms with positive manufacturing activity will miss some offshoring, for example by factoryless goods producers (FGPs) in the wholesale sector that have offshored all physical transformation activities (see Bernard and Fort, 2013, for details). Unfortunately, there is no practical way to distinguish all FGPs from traditional wholesale establishments. Furthermore, data on value-added and input usage, which is crucial for our structural estimation is less complete for firms outside manufacturing. We also note that we cannot identify manufacturing firms that use inputs imported by intermediaries.

²¹The wholesale sector includes a significant number of plants that design goods and coordinate production, often by offshoring, but do not perform physical transformation activities (see Bernard and Fort, 2013, for a description). Ignoring these plants' inputs could severely understate multi-sector firms' total inputs. For example, Feenstra and Jensen (2012) find that a significant fraction of some manufacturing firms' imports are not reported as input purchases. We address this issue by including a firm's wholesale plants' inputs. Although there is no way to measure inputs for establishments outside the manufacturing and wholesale sectors, those plants are much less likely to be involved in production or importing. An alternative approach would be to use an estimate of the demand elasticity σ and exploit the CES structure of our model to back out input usage from sales data.

the firm, while it uses the other factors of production (nonproduction workers, physical capital, and land) to combine intermediate inputs and cover all fixed costs. This approach is also motivated by the notion that the services typically provided by production workers are particularly offshorable. This new measure of total intermediate input purchases is highly correlated with traditional input measures for manufacturing firms based on reported inputs of materials and parts. A firm’s share of inputs from country j , χ_{ij} , is computed as imports from j divided by total input purchases. A firm’s share of domestic inputs, χ_{ii} , is simply the difference between its total input purchases and imports, divided by total input purchases.

The model predicts an important role for country characteristics in determining country-level fixed costs and sourcing potential. We compile a dataset with the key country characteristics—technology, wages, and tariffs, as well as other controls— from various sources. Country R&D data for 2007 are from the World Bank Development Indicators. Wage data are from the ILO data described by Oostendorp (2005). Tariffs are the simple average of country tariffs from the World Bank WITS database. Distance and language are from CEPII. Physical capital is based on the methodology in Hall and Jones (1999), but constructed using the most recent data from the Penn World Tables. Control of corruption is from the World Bank’s Worldwide Governance Indicators. We measure country human capital using the Hall and Jones (1999) methodology with updated data from Barro and Lee (2010). We also obtain country population from Barro and Lee (2010).

4.2 Descriptive Evidence

In this subsection, we provide reduced-form evidence that supports the notion that firms source multiple inputs from multiple countries. In addition, we show that firms tend to source multiple inputs per country and that they generally source a specific input from a single location. Finally, we provide evidence that firms follow a hierarchical pattern in their sourcing decisions.

4.2.1 Multiple inputs and countries

Two key assumptions that drive our theoretical approach are that firms source multiple inputs and that they may source these inputs from multiple countries. While the Census data do not provide detailed information about the total number of inputs used by a firm, the linked import data can shed light on the number of foreign inputs firms use. To exploit these data, we define a product as a distinct Harmonized Schedule ten-digit code, of which there are nearly 17,000 categories in the U.S. import data. Table 2 presents the firm-level statistics for importers on the number of countries from which a firm sources, as well as the number of unique HS10 products that it imports. The first column shows that, on average, firms import 12 products from about three foreign countries. There is considerable heterogeneity in these statistics across firms. The 95th percentile is 11 for the number of source countries and 41 for the number of imported products.

Table 2: Firm-level statistics on the number of source countries and imported inputs

	Mean	Std. Dev.	25th Ptile	Median	95th Ptile
Country Count	3.26	5.09	1	2	11
Product Count	11.91	48.89	1	3	41

Notes: The first row reports on the number of countries from which a firm imports. The second row reports on the number of unique HS10 products a firm imports.

4.2.2 The number of products per country and countries per product

In contrast to an Armington framework in which products are differentiated by country of origin, in our model we assume that each input could be sourced from any country. As a result, when a firm profit maximizes, it will only source each input from its lowest cost location. It is important to recognize that in either framework there will be interdependencies in a firm’s extensive margin decision about which countries to include in its sourcing strategy. In other words, the Armington assumption cannot be used to simplify the firm’s extensive margin sourcing decision. The two frameworks do lead to different predictions, however, in terms of how firms respond to shocks. In an Armington model, a shock to trade costs will only affect the amount that a firm sources of a particular product, while firms in our model will respond by changing the set of products sourced, as well as how much they source of each product.

Consistent with our model, we document that while firms tend to source multiple inputs per country, they seldom buy the same product from more than one country. The left panel of Table 3 presents statistics on the firm-level mean, median, and maximum number of products that a firm imports from a particular country. We report the mean, median, and 95th percentile of these firm-level measures.²² The average of the firm-level mean is 2.78 products imported per country and the 95th percentile of the firm-level mean is 8.23. Column 3 shows that the average of the maximum number of products a firm imports from a particular country is 7.21 and the 95th percentile is 25 products per country.

The right panel of Table 3 presents the same firm-level statistics for the number of countries from which a firm imports the same HS10 product. Almost every statistic reported in this table is about one. The median firm imports a single product from an average of only 1.03 countries. The median number of countries per product for firms is always 1.00, even for the 95th percentile of firms. Finally, the maximum number of countries per product for the median firm is still just one, while firms in the 95th percentile import the same product from a maximum of four countries. These results provide strong support for the premise that firms source each input from one location. In the Appendix we show that this pattern is still evident when the sample of importers is limited to firms that source from at least three countries. We also provide the statistics at the HS6 level

²²Data confidentiality protection rules preclude us from disclosing exact percentiles. Statistics for all percentiles are therefore the average for all firms that are within +/- one percent of a given percentile.

and show that every statistic on the number of countries from which a firm sources a given product is lower than the comparable statistic for the number of countries to which a firm exports a given product.

Table 3: Firm-level statistics on the number of imported products per source country and the number of source countries per imported product

	Products Per Country			Countries Per Product		
	Firm-level			Firm-level		
	Mean	Median	Max	Mean	Median	Max
Mean	2.78	2.18	7.21	1.11	1.00	1.61
Median	2.00	2.00	2.00	1.03	1.00	1.00
95%tile	8.23	5.00	25.00	1.78	1.00	4.00

Notes: The left panel reports on the number of unique HS10 products that a firm imports from a particular country. The right panel reports on the number of countries from which a firm imports the same HS10 product.

4.2.3 Hierarchies in firm sourcing patterns

We conclude this section by assessing the extent to which firms follow a hierarchical pecking order in their import behavior. Specifically, we count the number of firms that import from Canada (the top destination by firm rank) and no other countries, the number that import from Canada and China (the top two destinations) and no others, the number that import from Canada, China, and Germany and no others, and so on. When calculating these numbers, we limit the analysis to the top ten countries by firm rank. Table 4 presents the results. The first column shows the number of firms that import from each string of countries. To assess the significance of these numbers, column three provides the number of firms that would import from each string if sourcing decisions were independent and the fraction of actual firms importing from a given country represented that independent probability. The total fraction of firms that follows a pecking order is 36.0, almost twice the 19.6 percent that would be observed under independence.

This pattern is reminiscent of the results found by Eaton et al. (2011) in their study of French exporters. While it is certainly suggestive of a pecking order in which country characteristics make some countries particularly appealing for all U.S. firms, it also points to a high degree of firm-specific idiosyncracies in the selection of a firm’s sourcing strategy. We will incorporate this feature of the data in our structural analysis by extending the theory to allow for firm-country-specific fixed costs.

Table 4: U.S. firms importing from strings of top 10 countries

String	Data		Under Independence	
	Firms	% of Importers	Firms	% of Importers
CA	17,980	29.82	6,760	11.21
CA-CH	2,210	3.67	3,730	6.19
CA-CH-DE	340	0.56	1,030	1.71
CA-CH-DE-GB	150	0.25	240	0.40
CA-CH-DE-GB-TW	80	0.13	50	0.08
CA-CH-DE-GB-TW-IT	30	0.05	10	0.02
CA-CH-DE-GB-TW-IT-JP	30	0.05	0	0.00
CA-CH-DE-GB-TW-IT-JP-MX	50	0.08	0	0.00
CA-CH-DE-GB-TW-IT-JP-MX-FR	160	0.27	0	0.00
CA-CH-DE-GB-TW-IT-JP-MX-FR-KR	650	1.08	0	0.00
TOTAL Following Pecking Order	21,680	36.0	11,820	19.6

Notes: The string CA means importing from Canada but no other among the top 10; CA-CH means importing from Canada and China but no other, and so forth. % of Importers shows percent of each category relative to all firms that import from top 10 countries.

5 Structural Analysis

Having described the firm-level data, we now use it in conjunction with country-level data to estimate the key parameters of the model. We distinguish country sourcing potential from the fixed costs of sourcing and we quantify the extent to which the latter depend upon distance, common language, and control of corruption. The parameter estimates obtained here are also critical for performing the counterfactual exercises in the next section.

The structural analysis is performed in three distinct steps. In the first step, we use a simple linear regression to obtain an estimate of each country's sourcing potential $T_j (\tau_{ij} w_j)^{-\theta}$ from a U.S. perspective (i.e., $i = U.S.$). In the second step, we estimate the productivity dispersion parameter, θ , by projecting the estimated sourcing potential values on observed cost shifters and other controls. We also measure the elasticity of demand, σ , from observed variable mark-ups. In the third and final step, we estimate the fixed costs of sourcing and other distributional parameters via simulated method of moments. To make the firm's problem computationally feasible, we apply the technique in Jia (2008), originally designed to estimate an entry game among chains and other discount retailers in a large number of markets, to the sourcing strategy choices of U.S. firms.

Because we use data on the sourcing strategies of firms from a single country, in what follows, we often drop the subscript i from the notation, with the understanding that the unique importing country is the U.S. We also denote a firm by superscript n . To facilitate the estimation, we include only those countries that have at least 200 U.S. firms importing from them. This criterion leaves us with a total of 64 foreign sourcing options for firms. We include firms that import from these countries in the estimation, but adjust their input usage by subtracting their imports from the

excluded countries.²³

5.1 Step 1: Estimation of a Country’s Sourcing Potential

The first step in our structural analysis is to estimate each country’s sourcing potential. To do so, we take the firm’s sourcing strategy \mathcal{J}_i as given and exploit differences in its share of sourcing across countries. Recall from Equation (6) in the model that a firm’s share of inputs sourced from country j , χ_{ij} , is simply that country’s contribution to the firm’s sourcing capability, Θ_i . Country j ’s sourcing potential – from the perspective of country i – is therefore summarized by the term $\xi_j \equiv T_j (\tau_{ij} w_j)^{-\theta}$. Rearranging equation (6) by taking logs and normalizing the share of inputs purchased from country j by the firm’s share of domestic inputs leads to

$$\log \chi_{ij}^n - \log \chi_{ii}^n = \log \xi_j + \log \epsilon_j^n, \tag{21}$$

where n denotes firm. In order to turn the model’s equilibrium condition (6) into an empirical specification, note that this equation includes a firm-country-specific shock ϵ_j^n (relative to the firm-domestic-specific shock which we set at $\epsilon_i^n = 1$).²⁴

Intuitively, this specification allows us to identify a country’s average sourcing potential ξ_j by observing how much a firm imports from that country relative to the same firm’s domestic input purchases, restricting attention to countries included in the firm’s sourcing strategy. For this measurement strategy to be consistent, it is important that there is no selection based on the errors in this regression. This condition will be satisfied if firms only learn the country-specific efficiency shocks ϵ_j^n after their sourcing strategy is selected, or if the term ϵ_j^n simply represents measurement error.²⁵

Table 5 provides summary statistics from estimating equation 21 via OLS, using country fixed effects to capture the ξ_j terms. The estimated coefficients on these fixed effects represent each country’s sourcing potential. All sourcing potential fixed effects are significant at the 99 percent level. We have also estimated these sourcing potential measures controlling for industry effects. The estimates are highly correlated (0.996) with our baseline results and retain their statistical significance. To simplify the interpretation of the fixed effects, we estimate equation 21 without a constant. The adjusted R^2 from this specification is 0.85. The same specification with a constant yields an adjusted R^2 of 0.09. Including industry controls raises the R^2 to 0.15. We note that these relatively low R^2 s are typical when using firm-level data with high levels of idiosyncratic noise.

²³We note that a very small fraction of firms has negative values for its domestic input purchases. This occurs when a firm’s total input purchases are less than its imports. Likely explanations for this are measurement error and imports of capital equipment. We drop these firms from the estimation since we cannot use observations with negative input purchases. Since these firms are a small fraction of our total sample, we do not report how many there are to avoid future disclosure problems. We will provide details on this in the final version of the paper.

²⁴When normalizing by the domestic share, we set domestic trade costs to 1. All other country variables are then measured relative to the U.S. value.

²⁵However, we need to rule out measurement error related to a firm’s global sourcing strategy. In other words, we assume throughout that the set of countries from which the firm imports is correctly observed by the researcher and that a firm has positive imports from all countries for which it has paid a fixed cost of sourcing.

according to equation (8) on the dispersion parameter θ of the intermediates productivities.

Remember that this parameter θ , together with the elasticity of demand σ faced by the U.S. importer, are crucial for determining the qualitative implications of our model for the optimal choice of a global sourcing strategy. We next turn to discussing how we estimate these two parameters.

5.2 Step 2: Estimation of the Elasticity of Demand and Input Productivity Dispersion

5.2.1 Estimation of Elasticity of Demand

It is simpler to start by discussing how we recover σ from the data. Note that with CES preferences and monopolistic competition, the ratio of sales to variable input purchases (including intermediates and basic factors of production) is $\sigma/(\sigma - 1)$. We exploit this relationship to obtain a parameter value for σ by calculating a measure of average mark-ups from the establishment-level data in the Census of Manufactures. Specifically, the mark-up is the ratio of sales to variable inputs, where inputs are the sum of an establishment’s materials, wages, capital expenditures, and total expenses. The mark-up for the median establishment is 35 percent, with a bootstrapped standard error of 0.88. This implies an estimate for the elasticity of demand, σ , of 3.85.

5.2.2 Estimation of Dispersion of Input Productivity Shocks

A second key parameter of our model is the dispersion of the productivity shocks of the intermediate inputs. Conditional on the firm’s sourcing strategy, θ represents the firm-level trade elasticity in our model. We next use data on cost-shifters – wages and tariffs – to identify this elasticity. Recall that the sourcing potential ξ_j which we estimated in the previous section, is a function of a country’s technology parameter, trade costs, and wages (apart from the parameter θ). We thus project the estimated sourcing potential on proxies for all these terms, including R&D stock, capital per worker, a measure of control of corruption, wages, distance, tariffs, and common language. Specifically, we estimate the following equation:

$$\begin{aligned} \log \xi_j = & \beta_r \log \text{R\&D}_j + \beta_k \log \text{capital}_j + \beta_C \text{control of corruption} - \theta \log w_j \\ & - \theta (\log \beta_c + \beta_d \log \text{distance}_{ij} + \text{language}_{ij} \log \beta_l + \log(1 + \text{tariff}_{ij})) + \iota_j. \end{aligned} \quad (22)$$

Notice that the parameter θ can be recovered from the coefficient associated with wages or tariffs in that regression. A potential issue with the use of country wage data is the fact that variation in wages partly reflects differences in worker productivity and skill across countries. Since firms’ sourcing decisions are based on the cost of an efficiency unit of labor, we follow Eaton and Kortum (2002) and use a human capital-adjusted wage. Even adjusting for skill differences across workers, there are other country-level factors that are likely correlated with the average wage, such as infrastructure, that will lead to an upward bias on the wage coefficient. To address this issue, as well the potential for measurement error, we instrument for a country’s wage using its population.

The first column of Table 6 presents the OLS estimate of θ obtained when running (22) while constraining the same coefficient to apply to wages and tariffs. Column 2 provides the analogous IV estimate of θ using population as an instrument. Notice that the IV estimate (-1.785) is, as expected, larger in absolute value than the OLS estimate. It is interesting that, in line with our discussion in section 3.4, the data on firm-level trade flows suggest a much larger dispersion in productivities across countries than is typically obtained with aggregate trade data. For example, Eaton and Kortum (2002) estimate a coefficient of -3.8 using data on wages. Similarly to them, we find a coefficient of -4.763 when using the same specification as in equation (22) but with aggregate imports as a left-hand-side variable. These results are displayed in columns 4 and 5. It is noteworthy that our estimate of θ is sensitive to whether a measure of control of corruption is included in the projection or not. Corruption may be correlated with the extent to which incomplete contracting affects U.S. firms' sourcing decisions (see Antràs (2014) for some evidence). The results without control of corruption lead to a lower estimate for θ of -1.083 , as displayed in column 3. The aggregate trade elasticity in this specification (see column 6) is also lower with an estimate of -2.399 . We think that controlling for corruption in this regression makes sense, but since this is not a commonly used control in the literature which estimates the trade elasticity, we highlight the difference if a specification without this control is used.

Table 6: Estimation of with-in firm and aggregate trade elasticity

	log ξ			log aggregate import		
	OLS	IV	IV	OLS	IV	IV
log(1+tariff) + log wage	-0.471 (0.175)	-1.785 (0.651)	-1.083 (0.320)	-0.516 (0.389)	-4.763 (1.824)	-2.399 (0.788)
log distance	-0.326 (0.182)	-0.730 (0.311)	-0.514 (0.217)	-1.025 (0.404)	-2.329 (0.870)	-1.614 (0.534)
common language	0.291 (0.209)	0.199 (0.288)	0.434 (0.219)	0.489 (0.463)	0.189 (0.806)	0.965 (0.539)
log R&D	0.372 (0.0500)	0.490 (0.0875)	0.445 (0.0556)	0.668 (0.111)	1.052 (0.245)	0.859 (0.137)
log KL	-0.158 (0.170)	0.506 (0.384)	0.415 (0.291)	-0.266 (0.378)	1.879 (1.076)	1.528 (0.717)
Control of corruption	0.117 (0.147)	0.626 (0.308)		0.323 (0.327)	1.966 (0.864)	
Constant	-7.034 (0.852)	-11.44 (2.340)	-9.925 (1.523)	6.282 (1.891)	-7.969 (6.556)	-2.680 (3.752)
Observations	57	57	58	57	57	58

Notes: Standard errors in parentheses. IVs are population and tariff.

These estimates imply that a firm that sources from all countries faces between seven per-

cent ($1.137^{(-1/1.78)}$) and 11 percent ($1.137^{(-1/1.08)}$) lower input costs than a firm sourcing purely domestically, and consequently its sales are between 22 percent ($1.137^{(-2.85/1.78)}$) and 40 percent ($1.137^{(-2.85/1.08)}$) larger.

Across various specifications, we find that the ratio of elasticity of demand, $\sigma - 1$, to the dispersion of intermediate good efficiencies, θ , is always larger than one. As argued in section 2, this implies that the profit function has increasing differences in the firm's sourcing strategy. As explained below, in the third step of our estimation, we exploit this feature of our model when numerically solving the firm's problem in order to estimate the fixed cost of sourcing associated with different markets. For the following, we use the θ estimate of 1.78.

5.3 Step 3: Estimation of Fixed Costs of Sourcing

In this section, we estimate the fixed cost of sourcing, along with other distributional parameters, via Simulated Method of Moments. In order to match the model outcomes to moments generated from firm-level data, we extend our model to incorporate, aside from core productivity differences, two additional sources of firm heterogeneity: (i) firm-country-specific variation in fixed cost of sourcing, and (ii) firm-country-specific shocks to the sourcing potential, ϵ_j^n (which we already introduced in the estimation of sourcing potentials above). We assume that the firm initially observes its set of fixed cost draws, $f_{ij}^n, j = 1, \dots, J$ and knows the vector of sourcing potential, ξ_j , but learns about the firm-country-specific shock ϵ_j^n to the sourcing potential of country j only after paying the fixed cost to source from country j . We assume that firms have rational expectations when making their sourcing strategy decision.²⁶ The expected profits of selecting a sourcing strategy \mathcal{J} for a firm with core productivity φ and a vector of fixed costs of sourcing f_{ij}^n are given by:

$$\Pi(\mathcal{J}, \varphi, f_{ij}^n) = \varphi^{\sigma-1} B E_\epsilon \left((\gamma \Theta_i(\mathcal{J}, \epsilon))^{(\sigma-1)/\theta} \right) - \sum_{j \in \mathcal{J}} f_{ij}^n, \quad (23)$$

In a setting with a large number of countries, the firm faces a very large discrete choice problem to solve for its optimal sourcing strategy; if there are 65 countries, the firm selects between 2^{65} , which is roughly 10^{19} , possible sourcing strategies. Clearly, calculating the profits of each of these strategies for every firm is infeasible. Instead, we will apply an algorithm first developed by Jia (2008) to solve the firm's problem.

This algorithm works as follows. Given a core productivity φ and a sourcing strategy \mathcal{J} , the expected marginal benefit of adding country j to a firm's sourcing strategy is:

$$\varphi^{\sigma-1} \gamma^{(\sigma-1)/\theta} B E_\epsilon \left((\Theta_i(\mathcal{J}, \epsilon))^{(\sigma-1)/\theta} - (\Theta_i(\mathcal{J} \setminus j, \epsilon))^{(\sigma-1)/\theta} \right) - f_{ij}^n.$$

We define a mapping, $V_j(\mathcal{J})$ which takes a value of one if the marginal benefit of including country j in the sourcing strategy \mathcal{J} is positive, and takes a value of zero otherwise. Because of increasing

²⁶We calculate the expectation by assuming that the shocks $\log \epsilon_j^n$ are distributed according to a normal distribution with standard deviation equal to the root mean squared error of the OLS regression based on equation (21), and use 100 Halton draws to evaluate the expectation.

differences in the profit function, this mapping is an increasing function itself. Jia (2008) shows that when starting from the set \mathcal{J}^0 (which contains no country in the sourcing strategy), an iterative application of the V-operator leads to a lower bound of the sourcing strategy. That is, the optimal sourcing strategy contains at least the countries contained in this set. Similarly, when starting from the set \mathcal{J}^1 (which contains all countries in the sourcing strategy), the iterative application of the V-operator leads to an upper bound for the optimal sourcing strategy. Should the two bounds not overlap, one then only needs to evaluate the profits resulting from all possible combinations between the two bounds. In the presence of a high degree of complementarity, there is the potential for this algorithm to lead to a large number of possible choices between the two bounds, hence rendering this approach infeasible. Intuitively, the iterative process might stall too quickly if it is optimal for firms to add or drop countries from the set \mathcal{J} only in pairs (or larger groups).

Fortunately, in our application, this approach leads to completely overlapping lower and upper bounds in the vast majority of simulations. In addition, only a small number of countries differ in the bounds in those cases in which the bounds did not completely overlap. In principle, the algorithm could still be useful even if a sizable number of location sets need to be evaluated; for example, one could assume that the firm evaluates the lower and upper bounds and a random vector of alternative sourcing strategies that are contained in the two bounds.²⁷

Before describing the estimation method we complete the parameterization of the model. Specifically, we assume that the firm-country-specific fixed costs, f_{ij}^n , are drawn from a log-normal distribution with dispersion parameter, β_{disp}^f and scale parameter $\log \beta_c^f + \beta_d^f \log \text{distance}_{ij} + \log \beta_l^f \text{language}_{ij}$. Hence, fixed costs of sourcing are specified to be a function of distance and common language. We set the fixed cost of domestic sourcing to 0 since all firms in our sample feature positive purchases of domestic inputs. Some firms may be better at screening for foreign suppliers than other firms and thus we assume that the fixed cost draws are rank-correlated across countries, with a correlation coefficient of 0.9. Without rank-correlation, as the number of countries becomes large, it becomes very likely that each firm has a very low fixed cost of sourcing from at least one country, which would imply that the number of importers would go to one.²⁸ In addition to the fixed cost parameters, we estimate at this step the firm core productivity levels' dispersion parameter, κ , and the scale parameter B , which also includes the CES price index and the mass of firms.

Overall, we are left with the following six parameters to be estimated: $\delta = [B, \kappa, \beta_c^f, \beta_d^f, \beta_l^f, \beta_{\text{disp}}^f]$.²⁹ We simulate $S = 50,000$ U.S. firms, that is we draw for each firm a core-productivity shock from a uniform distribution (which, given a parameter guess κ , can be inverted to yield the Pareto distributed firm core productivity level), and a J -dimensional vector of fixed cost shocks from a

²⁷In two important aspects the sourcing strategy problem here is simpler than the original problem analyzed by Jia (2008), who studied the location choices of Walmart and Kmart. First, by assuming monopolistic competition we abstract away from strategic interactions between firms. Second, while the firm's choice in her problem manifests geographic dependence, in contrast, in our setting only the source countries' distances to the U.S. matter, but not the distances between the countries themselves.

²⁸We use a Gaussian copula and the marginal distributions of fixed cost for a particular country are distributed lognormal.

²⁹Note that we also set $\varphi_{\text{US}} = 1$, as it scales input purchases equivalently to an increase in B .

standard normal distribution (which, given a parameter guess β^f , can be used to calculate the firm-country specific fixed cost level).³⁰ Note there is no relationship between the number of simulated firms and the number of actual firms in the data. The model assumes that we have a continuum of firms whose core efficiency, fixed cost draws, and country-specific efficiency shocks follow particular distributions, and we use the simulated firms as evaluation points of these distributions.

We use the simulated firms to construct the following three sets of moments, which are directly affected by the level of the fixed costs, β^f , the scale term, B , and the dispersion of core productivity levels, κ :

1. The share of importing firms (about 24 percent in the data). This is simply a scalar, and we label this moment in the data as m_1 and the simulated moment as $\hat{m}_1(\delta)$.
2. The share of firms that sources from a particular country. We label this $J \times 1$ vector of moments in the data as m_2 and the simulated moment vector as $\hat{m}_2(\delta)$.
3. The share of firms that purchases inputs from a country less or more than the q -th percentile of input purchases in the data, where $q = (25, 50, 90)$. We label this $4 \cdot J \times 1$ vector of moments in the data as m_3 and the simulated moment vector as $\hat{m}_3(\delta)$.

We describe the difference between the moments in the data and in the simulated model by $\hat{y}(\delta)$:

$$\hat{y}(\delta) = m - \hat{m}(\delta) = \begin{bmatrix} m_1 - \hat{m}_1(\delta) \\ m_2 - \hat{m}_2(\delta) \\ m_3 - \hat{m}_3(\delta) \end{bmatrix}$$

The following moment condition is assumed to hold at the true parameter value δ_0 :

$$E[\hat{y}(\delta_0)] = 0 \tag{24}$$

The method of simulated moments selects the model parameters that minimize the following objective function:

$$\hat{\delta} = \arg \min_{\delta} [\hat{y}(\delta)]^\top \mathbf{W} [\hat{y}(\delta)], \tag{25}$$

where \mathbf{W} is a weighting matrix. At this point we use as weights simply the identity matrix, but we intend to use the optimal weighting matrix in future iterations of this paper.

The parameter estimates are displayed in Table 7 below.

We find that the fixed cost of sourcing are only slightly increasing in distance and that sourcing from countries with a common language reduces the fixed cost by about 30 percent. The fixed cost of sourcing are reasonable in magnitude. The median fixed cost estimate ranges from 40,000 to 60,000 USD. As a result of assuming a lognormal distribution, the fixed cost can be quite large for

³⁰We use Halton draws for the fixed cost shocks, which have better coverage properties than usual pseudo-random draws. See Chapter 6 in Train (2009) for a discussion.

Table 7: Simulated Method of Moments Estimates

B	0.0001
κ	1.5773
β_c^f	0.0592
β_d^f	0.0027
β_l^f	0.7134
β_{disp}^f	2.2022
Objective value	1.09

Notes: Standard errors yet to be calculated.

some individual firm-country combinations. The shape parameter of the Pareto distribution, κ , is lower than the elasticity of substitution, $\sigma - 1$, and hence we need to truncate the support of the core efficiencies at a very large positive number to guarantee a finite price index. Other studies also have estimated Pareto size distributions that require truncation (see the discussion in Antràs and Yeaple (2013)). We proceed by describing the fit of the data by the estimated model.


5.4 Fit of the model

Fit of the share of importers In the data around 24 percent of US firms import. In the simulated model we slightly under-predict this fraction and predict that 21 percent of firms are importers.

Fit of the share of importers by country Our fit of the share of importers by country is displayed in Figure 3. Overall, the model does a pretty good job at predicting which countries will have a large or small number of importers. We under-predict the number of firms that import from Canada (around 15 percent in the data and around 5 percent in the model) and over-predict the number of firms that import from China (around 20 percent in the model and 9 percent in the data). Hence, in our calibrated model China is the most popular destination country. Overall, it appears that we generally tend to overpredict the share of importers in countries where these shares are small in the data.

Fit of percentiles by country Our model also provides a reasonably good fit of the percentiles of imports by country. In Figure 4 we display the Median and 90th percentile import purchases by country. The median input purchase from Canada is around 30,000 USD and around 80,000 USD from China in both model and data (all these figures are conditional on buying a positive amount from the country). The 90 percentile purchase by firm from China is around 4 million USD and less than 2 million from Canada - again, the model fits those numbers well. However, the model

Figure 3: Share of importers by country in model and data




has a more difficult time fitting the distribution of U.S. input purchases. The median domestic input purchase is 560,000 USD in the data and only 1,000 USD in the estimated model. The 90th percentile of the domestic input purchases is 7.8 million USD in the data and 22,000 USD in the model.³¹

We are currently exploring how to improve the match of the size distribution of domestic input purchases. In particular, the current mis-match between the data and the model for U.S. sourcing appears to be driven by the fact that we allow for a foreign firm-country-specific shock to the sourcing potential but not for a firm-specific shock to the domestic sourcing potential. Because the firm chooses the products to source from each country after it observes the realization of the firm-country-specific shocks, and we assume that the shocks are distributed log-normal, this leads to relatively more sourcing from foreign countries since they can have positive shocks but domestic potential cannot. We plan to overcome this asymmetry by estimating country sourcing potential in Step 1 with firm random effects. These random effects incorporate a shock to firms' domestic sourcing potential and we can use the estimated dispersion of the random effects as a direct measure of the dispersion of domestic sourcing potential. Since we find that treating the firm-country-specific shocks to the sourcing potential as measurement error in Step 1, and not including these shocks in Step 3, closely matches the fit of the domestic percentiles without reducing the fit of the foreign percentiles, we are optimistic about this approach.

³¹In the simulated model, a domestic input purchase of 560,000 is reached by the 98th percentile, and a purchase of 7.8 million is reached and exceeded by the top .6 percent of firms.

Figure 4: Median and 90th percentile imports in model and data



Notes: In logarithmic scales. Line: $x=x$.

6 Counterfactuals

In this section, we use the parameter estimates from section 5 to assess how firm imports and the firm size distribution are affected by a shock to China’s sourcing potential. This shock to China’s sourcing potential could be due to a decrease in U.S tariffs on Chinese goods or to an increase in China’s productivity.

6.1 A 10 percent shock to China’s potential

We begin by assessing how a ten percent shock to China’s sourcing potential affects U.S. firms’ sourcing decisions when total market potential (and therefore the price index) is held constant. Table 8 shows that 1.5 percent of firms start importing from China in response to the shock. These firms increase their domestic sourcing by 22.1 percent and their sourcing from other countries by 8.5 percent. Almost 20 percent of firms continue sourcing from China, and these firms also increase both domestic and other foreign sourcing, though only by a small amount less than 0.1 percent. For both entrants and continuers, the relatively larger increase in domestic sourcing is attributable to the fact that a fraction of these firms do not source from countries other than China or the U.S. Finally, firms that do not source from China before or after the shock are unaffected by the change.

We next consider how a ten percent shock to China’s sourcing potential affects firms’ sourcing decisions when the price index adjusts. Table 9 shows that, in this case, only 0.8 percent of firms start importing from China. While these firms increase their domestic sourcing by 13.5 percent, they decrease their sourcing from foreign countries by 17 percent. Firms that continue sourcing from China still comprise 19.5 percent of firms, however, these firms now decrease both domestic and foreign sourcing by 6.4 and 6.7 percent respectively. The most notable difference from allowing the price index to adjust is evident in the fact that firms that do not source from China before or after the shock shrink substantially. The last row of Table 9 shows that these firms decrease their

Table 8: Firm responses to 10% shock to China’s sourcing potential, constant price index

China Import Status	Share of Firms	Relative Sourcing from	
		U.S.	Other Countries
Entrants	0.015	1.221	1.085
Exiters	0.000	NaN	NaN
Continuers	0.195	1.006	1.002
Others	0.790	1.000	1.000

domestic sourcing by 6.9 percent and the foreign sourcing by 24.8 percent.

Table 9: Firm responses to 10% shock to China’s sourcing potential, endogenous price index


China Import Status	Share of Firms	Relative Sourcing from	
		U.S.	Other Countries
Entrants	0.008	1.135	0.830
Exiters	0.000	NaN	NaN
Continuers	0.195	0.936	0.933
Others	0.796	0.931	0.752

We finally consider how the same ten percent shock to China’s sourcing potential affects the size distribution of U.S. firms, when the price index adjusts. Figure 5 depicts the percent growth in sales by firm size decile. Although not evident in the figure, we first note that firms in the top decile experience a very small percentage decrease in their sales. Firms in deciles six to nine grow by just under five percent to almost twenty percent. In contrast, firms in the bottom five deciles see their sales shrink. These results highlight the fact that extensive margin sourcing decisions lead to heterogeneous responses across firms to trade or other shocks.

7 Conclusion

To be written.

Figure 5: Size distribution of firms in response to 10% shock to China's sourcing potential



A Theory Appendix

Proof of Proposition 1

Consider two firms with productivities φ_H and φ_L , with $\varphi_H > \varphi_L$. Denote by $\mathcal{J}_i(\varphi_H) = \{j : I_{ij}(\varphi_H) = 1\}$ and $\mathcal{J}_i(\varphi_L) = \{j : I_{ij}(\varphi_L) = 1\}$ the optimal sourcing strategies of these firms, and suppose that $\mathcal{J}_i(\varphi_H) \neq \mathcal{J}_i(\varphi_L)$ (when $\mathcal{J}_i(\varphi_H) = \mathcal{J}_i(\varphi_L)$ the result in the Proposition holds trivially). For firm φ_H to prefer $\mathcal{J}_i(\varphi_H)$ over $\mathcal{J}_i(\varphi_L)$, we need

$$\varphi_H^{\sigma-1} (\gamma \Theta_i(\mathcal{J}_i(\varphi_H)))^{(\sigma-1)/\theta} B_i - \sum_{j \in \mathcal{J}_i(\varphi_H)} f_{ij} > \varphi_H^{\sigma-1} (\gamma \Theta_i(\mathcal{J}_i(\varphi_L)))^{(\sigma-1)/\theta} B_i - \sum_{j \in \mathcal{J}_i(\varphi_L)} f_{ij},$$

while φ_L preferring $\mathcal{J}_i(\varphi_L)$ over $\mathcal{J}_i(\varphi_H)$ requires

$$\varphi_L^{\sigma-1} (\gamma \Theta_i(\mathcal{J}_i(\varphi_H)))^{(\sigma-1)/\theta} B_i - \sum_{j \in \mathcal{J}_i(\varphi_H)} f_{ij} < \varphi_L^{\sigma-1} (\gamma \Theta_i(\mathcal{J}_i(\varphi_L)))^{(\sigma-1)/\theta} B_i - \sum_{j \in \mathcal{J}_i(\varphi_L)} f_{ij}.$$

Combining these two conditions, we find

$$[\varphi_H^{\sigma-1} - \varphi_L^{\sigma-1}] \left[\Theta_i(\mathcal{J}_i(\varphi_H))^{(\sigma-1)/\theta} - \Theta_i(\mathcal{J}_i(\varphi_L))^{(\sigma-1)/\theta} \right] \gamma^{(\sigma-1)/\theta} B_i > 0.$$

Given $\varphi_H > \varphi_L$, this necessarily implies $\Theta_i(\varphi_H) > \Theta_i(\varphi_L)$.

Proof of Proposition 2

As noted in the main text, when $(\sigma - 1)/\theta > 1$, the profit function in (10) features increasing differences in (I_{ij}, I_{ik}) for $j, k \in \{1, \dots, J\}$ with $j \neq k$. Furthermore, it also features increasing differences in (I_{ij}, φ) for any $j \in J$. Invoking Topkis's monotonicity theorem, we can then conclude that for $\varphi_H \geq \varphi_L$, we must have $(I_{i1}(\varphi_H), I_{i2}(\varphi_H), \dots, I_{iJ}(\varphi_H)) \geq (I_{i1}(\varphi_L), I_{i2}(\varphi_L), \dots, I_{iJ}(\varphi_L))$. Naturally, this rules out a situation in which $I_{ij}(\varphi_H) = 0$ but $I_{ij}(\varphi_L) = 1$, and thus we can conclude that $\mathcal{J}_i(\varphi_L) \subseteq \mathcal{J}_i(\varphi_H)$ for $\varphi_H \geq \varphi_L$.

Proof of Proposition 3

Given a vector of wages, equations (12) and (13) determine the equilibrium values of B_i and N_i . Notice that the firm-level global sourcing problem depends only on B_i , w_i and exogenous parameters, and not directly on N_i . As a result, if a unique solution for B_i exists, all thresholds $\tilde{\varphi}_{ij}$ for any pair of countries (i, j) will be pinned down uniquely, given wages. Hence, if a unique solution for B_i in equation (12) exists, we can ensure that there will be a unique value of N_i solving (13). Let us then focus on studying whether (12) indeed delivers a unique solution for B_i .

For given wages, the equilibrium condition (12) can be rearranged as follows

$$w_i f_e = B_i \int_{\tilde{\varphi}_{i\vartheta(i)}}^{\infty} (\gamma \Theta_i(\varphi))^{(\sigma-1)/\theta} \varphi^{\sigma-1} dG_i(\varphi) - w_i \int_{\tilde{\varphi}_{i\vartheta(i)}}^{\infty} \sum_{j \in \mathcal{J}_i(\varphi)} f_{ij} dG_i(\varphi), \quad (26)$$

where remember that $\vartheta(i)$ is defined as $\vartheta(i) = \{j \in J : \tilde{\varphi}_{ij} \leq \tilde{\varphi}_{ik} \text{ for all } k \in J\}$ and thus satisfies

$$(\tilde{\varphi}_{i\vartheta(i)})^{\sigma-1} B_i \left(\gamma T_{\vartheta(i)} (\tau_{i\vartheta(i)} w_{\vartheta(i)})^{-\theta} \right)^{(\sigma-1)/\theta} = w_i f_{i\vartheta(i)}. \quad (27)$$

Remember also that $\Theta_i(\varphi) \equiv \sum_{k \in \mathcal{J}_i(\varphi)} T_k (\tau_{ik} w_k)^{-\theta}$, and $\mathcal{J}_i(\varphi) \subseteq J$ is the set of countries for which a firm based in i with productivity φ has paid the associated fixed cost of offshoring $w_i f_{ij}$.

Computing the derivative of the right-hand-side of (26) with respect to B_i , and using (27) to eliminate the effects working through changes in $\tilde{\varphi}_{i\vartheta(i)}$, we can write this derivative as simply

$$\int_{\tilde{\varphi}_{i\vartheta(i)}}^{\infty} \frac{\partial \left(\varphi^{\sigma-1} (\gamma \Theta_i(\varphi))^{(\sigma-1)/\theta} B_i - w_i \sum_{j \in \mathcal{J}_i(\varphi)} f_{ij} \right)}{\partial B_i} dG_i(\varphi) > 0. \quad (28)$$

The fact that this derivative is positive follows directly from the firm's global sourcing problem in 10. In particular, holding constant the firm's sourcing strategy $\mathcal{J}_i(\varphi)$ – and thus $\Theta_i(\varphi)$ –, it is clear that an increase in B_i will increase firm level profits $\varphi^{\sigma-1} (\gamma \Theta_i(\varphi))^{(\sigma-1)/\theta} B_i - w_i \sum_{j \in \mathcal{J}_i(\varphi)} f_{ij}$. Now such an increase in B_i might well affect the profit-maximizing choice of $\mathcal{J}_i(\varphi)$ – and thus $\Theta_i(\varphi)$ –, but firm profits could not possibly be reduced by those changes, since the firm can always decide not to change the global sourcing strategy in light of the higher B_i and still obtain higher profits.³² We can thus conclude that the right-hand-side of (26) is monotonically increasing in B_i .

It is also clear that when $B_i \rightarrow \infty$, all firms will find it optimal to source everywhere and the right-hand-side of (26) becomes

$$B_i \left(\gamma \sum_{k \in J} T_k (\tau_{ik} w_k)^{-\theta} \right)^{(\sigma-1)/\theta} \int_{\underline{\varphi}_i}^{\infty} \varphi^{\sigma-1} dG_i(\varphi) - w_i \sum_{j \in J} f_{ij}$$

and thus goes to ∞ . Conversely, when $B_i \rightarrow 0$, no firm can profitably source to any location, given the positive fixed costs of sourcing, and thus the right-hand-side of (26) goes to 0.

It thus only remains to show that the right-hand-side of (26) is a *continuously* non-decreasing function of B_i . This may not seem immediate because firm-level profits jump discontinuously with B_i when such

³²Following the same steps as in the proof of Proposition 1 we can show that both $\Theta_i(\varphi)$ and $\sum_{j \in \mathcal{J}_i(\varphi)} f_{ij}$ are actually non-decreasing in B_i . This result is immaterial for the proof of existence and uniqueness in the case of free entry, but can be used to proof the same result for the case of an exogenous number of firms N_i .

changes in B_i lead to changes in the global sourcing strategy of firms. It can be shown, however, that

$$\int_{\tilde{\varphi}_i(\varphi)}^{\infty} \frac{\partial \left((\Theta_i(\varphi))^{(\sigma-1)/\theta} B_i \varphi^{\sigma-1} \right)}{\partial B_i} dG_i(\varphi)$$

is continuously differentiable in B_i . To see this, one can first follow the same steps as in the proof of Proposition 1 to show that $\Theta_i(\varphi; B_i)$ must be non-decreasing not only in φ , but also in B_i and $B_i \varphi^{\sigma-1}$. We can then represent $(\Theta_i(\varphi))^{(\sigma-1)/\theta} B_i \varphi^{\sigma-1}$ as a non-decreasing step function in φ , in which the jumps occur at different levels of $B_i \varphi^{\sigma-1}$. This is analogous to writing

$$(\Theta_i(\varphi))^{(\sigma-1)/\theta} B_i \varphi^{\sigma-1} = \begin{cases} \theta_1 B_i \varphi^{\sigma-1} & \text{if } \varphi < a_1/B_i^{1/(\sigma-1)} \\ \theta_2 B_i \varphi^{\sigma-1} & \text{if } a_1/B_i^{1/(\sigma-1)} \leq \varphi < a_2/B_i^{1/(\sigma-1)} \\ \vdots & \vdots \\ \theta_J B_i \varphi^{\sigma-1} & \text{if } a_{J-1}/B_i^{1/(\sigma-1)} \leq \varphi \end{cases}. \quad (29)$$

Hence, we have

$$\begin{aligned} \int_{\tilde{\varphi}_i(\varphi)}^{\infty} (\Theta_i(\varphi))^{(\sigma-1)/\theta} B_i \varphi^{\sigma-1} dG_i(\varphi) &= \int_{\tilde{\varphi}_i(\varphi)}^{a_1/B_i^{1/(\sigma-1)}} \theta_1 B_i \varphi^{\sigma-1} dG_i(\varphi) + \\ &\int_{a_1/B_i^{1/(\sigma-1)}}^{a_2/B_i^{1/(\sigma-1)}} \theta_2 B_i \varphi^{\sigma-1} dG_i(\varphi) + \dots + \int_{a_{J-1}/B_i^{1/(\sigma-1)}}^{\infty} \theta_J B_i \varphi^{\sigma-1} dG_i(\varphi). \end{aligned}$$

It is then clear that the derivative of this expression with respect to B_i is a sum of continuous functions of B_i , and thus is continuous in B_i itself. Using similar arguments we can next show that

$$\int_{\tilde{\varphi}_i(\varphi)}^{\infty} \frac{\partial \left(w_i \sum_{j \in \mathcal{J}_i(\varphi)} f_{ij} \right)}{\partial B_i} dG_i(\varphi) \quad (30)$$

is also continuously differentiable in B_i . First, a simple proof by contradiction can be used to show that $\sum_{j \in \mathcal{J}_i(\varphi)} f_{ij}$ is non-decreasing in $B_i \varphi^{\sigma-1}$. More specifically, suppose that for $(B_i \varphi^{\sigma-1})_H > (B_i \varphi^{\sigma-1})_L$ we also had $\sum_{j \in \mathcal{J}_{iH}} f_{ij} < \sum_{j \in \mathcal{J}_{iL}} f_{ij}$. Given the non-decreasing dependence of $\Theta_i(\varphi)$ on $B_i \varphi_i^{\sigma-1}$, we would then have

$$(\gamma \Theta_{iH}(\varphi))^{(\sigma-1)/\theta} (B_i \varphi^{\sigma-1})_L - \sum_{j \in \mathcal{J}_{iH}} f_{ij} > (\gamma \Theta_{iL}(\varphi))^{(\sigma-1)/\theta} (B_i \varphi^{\sigma-1})_L - \sum_{j \in \mathcal{J}_{iL}(\varphi)} f_{ij},$$

which clearly contradicts \mathcal{J}_{iL} being optimal given $B_i \varphi^{\sigma-1} = (B_i \varphi^{\sigma-1})_L$. With this result, $\sum_{j \in \mathcal{J}_i(\varphi)} f_{ij}$ can then be expressed as a step function analogous to that in (29), in which the position of the steps is continuously differentiable in B_i . This in turn ensures that (26) is continuous in B_i and concludes the proof that there exists a unique B_i that solves equation (12).

Online Appendix (Not for Publication)

Equilibrium in the Complements-Pareto Case

In Proposition 2, we have established that whenever $\sigma - 1 > \theta$, the model delivers a ‘pecking order’ in the extensive margin of offshoring. For each country i , we can then rank foreign countries in terms of some index of sourcing appeal. Suppose for simplicity that $\tilde{\varphi}_i = \tilde{\varphi}_{ii}$, so that all firms that source (i.e., all active firms) do so, at least in part, from Home. Denote by r the r -th least appealing country from which firms from i source from, so Home is $r = 1$. Define also

$$\Theta_{ir} = \sum_{j=1}^r T_j (\tau_{ij} w_j)^{-\theta}.$$

Note also that Proposition 2 implies that the set of productivity thresholds $\tilde{\varphi}_{ir}$ defined in the main text will be such that any firm with productivity above that threshold $\tilde{\varphi}_{ir}$ necessarily sources from country r , or in terms of the notation in equation (18), $I_{ir}(\varphi) = 1$ for all $\varphi > \tilde{\varphi}_{ir}$.

In light of the profit function in 9, these thresholds are given by

$$\begin{aligned} \tilde{\varphi}_{i1}^{\sigma-1} &= \frac{w_i f_{ir}}{\gamma^{(\sigma-1)/\theta} B_i \left(T_i (w_i)^{-\theta} \right)^{(\sigma-1)/\theta}}; \\ \tilde{\varphi}_{ir}^{\sigma-1} &= \frac{w_i f_{ir}}{\gamma^{(\sigma-1)/\theta} B_i \left(\Theta_{ir}^{(\sigma-1)/\theta} - \Theta_{ir-1}^{(\sigma-1)/\theta} \right)} \quad \text{for } r > 1. \end{aligned} \quad (31)$$

Consider now the industry equilibrium. Using the above notation, we can write the free entry condition (12) as

$$\gamma^{(\sigma-1)/\theta} B_i \sum_{r=1}^{J-1} \Theta_{ir}^{(\sigma-1)/\theta} \int_{\tilde{\varphi}_{ir}}^{\tilde{\varphi}_{ir+1}} \varphi^{\sigma-1} dG_i(\varphi) - w_i \sum_{r=1}^{J-1} f_{ir} \int_{\tilde{\varphi}_{ir}}^{\infty} dG_i(\varphi) = w_i f_e.$$

Next, invoking the Pareto distribution, $G_i(\varphi) = 1 - (\underline{\varphi}_i/\varphi)^\kappa$, and solving for the integrals, we obtain:

$$\gamma^{(\sigma-1)/\theta} B_i \sum_{r=1}^{J-1} \Theta_{ir}^{(\sigma-1)/\theta} \kappa \left(\frac{\underline{\varphi}_i}{\tilde{\varphi}_{ir}} \right)^\kappa \frac{(\tilde{\varphi}_{ir})^{\sigma-\kappa-1} - (\tilde{\varphi}_{ir+1})^{\sigma-\kappa-1}}{\kappa - \sigma + 1} - w_i \sum_{r=1}^{J-1} f_{ir} \left(\frac{\underline{\varphi}_i}{\tilde{\varphi}_{ir}} \right)^\kappa = w_i f_e.$$

Plugging the thresholds in (31) delivers

$$\begin{aligned} \gamma^{(\sigma-1)/\theta} B_i \kappa \left(\frac{\underline{\varphi}_i}{\tilde{\varphi}_{ir}} \right)^\kappa \sum_{r=1}^{J-1} \Theta_{ir}^{(\sigma-1)/\theta} \frac{(\tilde{\varphi}_{ir})^{-\kappa} \frac{w_i f_{ir}}{\gamma^{(\sigma-1)/\theta} B_i \left(\Theta_{ir}^{(\sigma-1)/\theta} - \Theta_{ir-1}^{(\sigma-1)/\theta} \right)} - (\tilde{\varphi}_{ir+1})^{-\kappa} \frac{w_i f_{ir+1}}{\gamma^{(\sigma-1)/\theta} B_i \left(\Theta_{ir+1}^{(\sigma-1)/\theta} - \Theta_{ir}^{(\sigma-1)/\theta} \right)}}{\kappa - \sigma + 1} \\ - w_i \sum_{r=1}^{J-1} f_{ir} \left(\frac{\underline{\varphi}_i}{\tilde{\varphi}_{ir}} \right)^\kappa = w_i f_e. \end{aligned}$$

Expanding the summation involving the terms Θ_{ir} , canceling the terms in $\Theta_{ir}^{(\sigma-1)/\theta} - \Theta_{ir-1}^{(\sigma-1)/\theta}$, and simplifying, we finally obtain

$$\frac{\sigma - 1}{\kappa - \sigma + 1} \sum_{r=1}^{J-1} \left(\frac{\underline{\varphi}_i}{\tilde{\varphi}_{ir}} \right)^\kappa f_{ir} = f_{ei}. \quad (32)$$

It is worth emphasizing that this equation holds regardless of the relative values of $\sigma - 1$ and θ as long as

these parameters and the degree of heterogeneity in fixed costs are such that a hierarchy in sourcing decisions exists. The key insight of Proposition 2 is that $\sigma - 1 > \theta$ is a sufficient condition for this hierarchical structure to emerge regardless of the values of the fixed costs of offshoring f_{ij} .

Note that equation (32) implies that

$$\int_{\tilde{\varphi}_i}^{\infty} \sum_{j \in \mathcal{J}_i(\varphi)} f_{ij} dG_i(\varphi) + f_{ei} = \left(\frac{\sigma - 1}{\kappa - \sigma + 1} + 1 \right) f_{ei},$$

and thus plugging this expression in (13), we can conclude that

$$N_i = \frac{(\sigma - 1) \eta L_i}{\sigma \kappa f_{ei}}, \quad (33)$$

as claimed in footnote 14.

Some of the above expressions are useful in deriving the gravity equation in (17) characterizing bilateral manufacturing trade flows in the case of independent entry decisions (or $\sigma - 1 = \theta$). To see this, begin with equation (16) and plug the formula for the Pareto distribution to obtain

$$M_{ij} = (\sigma - 1) N_i B_i \gamma T_j (\tau_{ij} w_j)^{-\theta} \kappa \underline{\varphi}_i^\kappa \frac{(\tilde{\varphi}_{ij})^{\sigma-1-\kappa}}{\kappa - \sigma + 1}.$$

With independent entry decisions, the threshold in (31) simplifies to

$$\tilde{\varphi}_{ij}^{\sigma-1} = \frac{w_i f_{ij}}{\gamma B_i T_j (\tau_{ij} w_j)^{-\theta}}.$$

Plugging this expression for $\tilde{\varphi}_{ij}^{\sigma-1}$ into the previous one for M_{ij} , imposing $\theta = \sigma - 1$, and manipulating the resulting expression in a manner analogous to the derivation of the benchmark Eaton-Kortum gravity equation in (15), we obtain

$$M_{ij} = N_i (B_i)^{\frac{\kappa}{\sigma-1}} (\tau_{ij})^{-\kappa} \underline{\varphi}_i^\kappa (w_i f_{ij})^{1-\frac{\kappa}{\sigma-1}} \frac{Q_j}{\sum_k N_k (B_k)^{\frac{\kappa}{\sigma-1}} (\tau_{kj})^{-\kappa} (\tilde{\varphi}_k)^\kappa (w_k f_{kj})^{1-\frac{\kappa}{\sigma-1}}}.$$

Using (3) and (33) and defining

$$\Psi_i = \frac{f_{ei}}{L_i} \underline{\varphi}_i^{-\kappa} P_i^{-\kappa} w_i^{\kappa/(\sigma-1)-1},$$

we thus obtain equation (17) in the main text.

B Data Appendix

B.1 Sample

Table B.1 provides details of all firms in the Economic Censuses with positive sales and employment. The first row corresponds to firms that consist only of manufacturing establishments (“M” firms). The second row presents information for all firms with one or more manufacturing establishments and at least one establishment outside of manufacturing (“M+” firms). Together, these two types of firms comprise our sample.

Table B.1: Sample of firms

Firm Type	Firms	Imports \$millions	Empl \$millions	Sales \$billions	Fraction Importers
Manufacturing Only (M)	238,800	75,938	5,866	1,230	0.23
Manufacturing Plus (M+)	11,500	829,594	20,573	9,180	0.77
Other (O)	4,001,000	99,937	77,204	12,553	0.03
Wholesale Only (W)	300,000	240,654	3,484	2,300	0.31
Wholesale and Other (WO)	7,500	141,740	6,426	2,252	0.51
Total	4,561,700	1,387,863	113,553	27,515	0.06

Notes: Table provides information on firms in the Economic Census with positive sales and employment. Analysis in paper based on all M and M+ firms. Numbers rounded for disclosure avoidance. Imports exclude products classified under mining.

B.2 Premia figures

In the introduction, we plot the relationship between the log of firm sales and the minimum number of countries from which a firm sources. To construct the figure, we regress the log of firm sales on cumulative dummies for the number of countries from which a firm sources and industry controls. The omitted category is non-importers, so the premia are interpreted as the difference in size between non-importers and firms that import from at least one country, at least two countries, etc. The horizontal axis denotes the number of countries from which a firm sources, with 1 corresponding to firms that use only domestic inputs. The introduction figure controls for firm industry with variables that measure the share of a firm’s employment in four-digit NAICS industries. (These are simply industry fixed effects for all firms that span only one industry.) Here we show that the patterns depicted in the introduction are robust when considering a firm’s size prior to importing and when controlling for the products that a firm imports or exports. Figure B.1 plots the relationship between a firm’s log sales in 2002 and the number of countries from which it sources in 2007, for firms that did *not* import in 2002. Figure B.2 depicts the relationship when controlling for the number of products a firm imports (left panel) and the number of products the firm exports (right panel). In additional undisclosed results, available upon request, we show similar patterns when using firm employment and the log of value-added labor productivity.

Figure B.1: Importer premia for firm’s 2002 sales, limited to firms that do not import in 2002,





Figure B.2: Importer premia with product controls



(a) Controlling for number of products imported by the firm (b) Controlling for number of products exported by the firm

B.3 Countries per product counts

In section 4.2.2, we show that most firms source most products from a single location. Table B.2 shows that this pattern persists for firms that source from at least three foreign countries. We also compare these firm-level statistics to the same numbers for exporting. To make a valid comparison, we must first aggregate to the HS6 level. This ensures the same number of product categories for imports and exports. Table B.3 shows that, even at the HS6 level, most firms source most products from one location. This is in contrast to firms’ exporting decisions, where we see that the median firm sells at least one product to three destinations and the 95th percentile sells to 21 countries.

Table B.2: Number of countries from which a firm imports the same HS10 product, for firms that import from at least 3 countries

	Firm Level		
	Mean	Median	Max
Mean	1.28	1.05	3.18
Median	1.19	1.00	2.00
95%tile	1.96	1.00	9.00

Notes: Table reports statistics on the firm-level mean, median, and maximum of the number of countries from which a firm imports the same HS10 product.

Table B.3: Number of countries per HS6 product traded by a firm

	Firm Level Imports			Firm Level Exports		
	Mean	Median	Max	Mean	Median	Max
Mean	1.15	1.00	1.92	1.76	1.00	4.26
Median	1.05	1.00	1.00	1.33	1.00	3.00
95%tile	1.92	1.00	5.00	4.87	2.00	21.00

Notes: Table reports statistics on the firm-level mean, median, and maximum of the number of countries from which a firm imports or exports the same HS6 product.

References

- Acemoglu, Daron, Pol Antràs, and Elhanan Helpman**, “Contracts and Technology Adoption,” *American Economic Review*, 2007, 97 (3), 916–943.
- Amiti, Mary and Jozef Konings**, “Trade Liberalization, Intermediate Inputs, and Productivity: Evidence from Indonesia,” *American Economic Review*, December 2007, 97 (5), 1611–1638.
- Antràs, Pol**, *Global Production: A Contracting Perspective* 2014. in-progress book manuscript available at <http://scholar.harvard.edu/antras/scholar.harvard.edu/antras/pages/books>.
- **and Elhanan Helpman**, “Global Sourcing,” *Journal of Political Economy*, 2004, 112 (3).
- **and –**, “Contractual Frictions and Global Sourcing,” in “The Organization of Firms in a Global Economy,” Cambridge, MA: Harvard University Press, 2008.
- **and Stephen Yeaple**, “Multinational Firms and International Trade Structure,” in “Handbook of International Economics,” Vol. 4, Elsevier, 2013.
- **, Luis Garicano, and Esteban Rossi-Hansberg**, “Offshoring in a Knowledge Economy,” *The Quarterly Journal of Economics*, 2006, 121 (1), 31–77.

- Arkolakis, Costas, Natalia Ramondo, Andrés Rodríguez-Clare, and Stephen Yeaple**, “Innovation and Production in the Global Economy,” Working Paper 18972, National Bureau of Economic Research April 2013.
- , **Svetlana Demidova, Peter J Klenow, and Andrés Rodríguez-Clare**, “Endogenous Variety and the Gains from Trade,” *American Economic Review*, April 2008, *98* (2), 444–450.
- Barro, Robert and Jong Wha Lee**, “A New Data Set of Educational Attainment in the World, 1950-2010,” 2010. NBER Working Paper No. 15902.
- Benassy, Jean-Pascal**, “Is there always too little research in endogenous growth with expanding product variety?,” *European Economic Review*, January 1998, *42* (1), 61–69.
- Bernard, Andrew B. and Teresa C. Fort**, “Factorless Goods Producers in the US,” Working Paper 19396, NBER June 2013.
- , **Andreas Moxnes, and Karen Helene Ulltveit-Moe**, “Two-Sided Heterogeneity and Trade,” SSRN Scholarly Paper ID 2341044, Social Science Research Network, Rochester, NY October 2013.
- Bernard, Andrew B, J. Bradford Jensen, Stephen J Redding, and Peter K Schott**, “Firms in International Trade,” *Journal of Economic Perspectives*, August 2007, *21* (3), 105–130.
- , —, —, and —, “The Margins of US Trade,” *American Economic Review*, April 2009, *99* (2), 487–493.
- Blaum, Joaquin, Claire Lelarge, and Michael Peters**, “Non-Homothetic Import Demand: Firm Productivity and Quality Bias,” *Unpublished paper*, 2013.
- Blum, Bernardo S, Sebastian Claro, and Ignatius Horstmann**, “Facts and Figures on Intermediated Trade,” *American Economic Review*, May 2010, *100* (2), 419–423.
- Carballo, Jerónimo, Gianmarco I. P. Ottaviano, and Christian Volpe Martincus**, “The Buyer Margins of Firms’ Exports,” CEP Discussion Paper dp1234, Centre for Economic Performance, LSE 2013.
- Chaney, Thomas**, “Distorted Gravity: The Intensive and Extensive Margins of International Trade,” *American Economic Review*, August 2008, *98* (4), 1707–1721.
- Costinot, Arnaud**, “An Elementary Theory of Comparative Advantage,” *Econometrica*, 2009, *77* (4), 1165–1192.
- Dragusanu, Raluca**, “Firm-to-Firm Matching Along the Global Supply Chain,” 2014. mimeo Harvard University.
- Eaton, Jonathan and Samuel Kortum**, “Technology, Geography, and Trade,” *Econometrica*, 2002, *70* (5), 1741–1779.
- , **Marcela Eslava, David Jinkins, C. Krizan, and James Tybout**, “A Search and Learning Model of Export Dynamics,” 2012. mimeo Penn State University.
- , **Samuel Kortum, and Francis Kramarz**, “An Anatomy of International Trade: Evidence From French Firms,” *Econometrica*, 2011, *79* (5), 1453–1498.
- , —, —, and **Raul Sampognaro**, “Firm-to-Firm Trade: Imports, Exports, and the Labor Market,” 2013. unpublished manuscript.
- Ethier, Wilfred J.**, “National and International Returns to Scale in the Modern Theory of International Trade,” *The American Economic Review*, June 1982, *72* (3), 389–405.
- Feenstra, Robert C. and J. Bradford Jensen**, “Evaluating Estimates of Materials Offshoring from U.S. Manufacturing,” *Economic Letters*, 2012, *117*, 170–173.

- Fort, Teresa C.**, “Breaking Up Is Hard To Do: Why Firms Fragment Production Across Locations,” Working Paper 13-35, Center for Economic Studies, U.S. Census Bureau 2013.
- Garetto, Stefania**, “Input Sourcing and Multinational Production,” *American Economic Journal: Macroeconomics*, April 2013, 5 (2), 118–151.
- Goldberg, Pinelopi Koujianou, Amit Kumar Khandelwal, Nina Pavcnik, and Petia Topalova**, “Imported Intermediate Inputs and Domestic Product Growth: Evidence from India,” *The Quarterly Journal of Economics*, November 2010, 125 (4), 1727–1767.
- Gopinath, Gita and Brent Neiman**, “Trade Adjustment and Productivity in Large Crises,” *American Economic Review*, 2013, p. Forthcoming.
- Grossman, Gene M. and Esteban Rossi-Hansberg**, “Trading Tasks: A Simple Theory of Offshoring,” *American Economic Review*, 2008, 98 (5), 1978–97.
- and —, “Task Trade Between Similar Countries,” *Econometrica*, 2012, 80 (2), 593–629.
- Grossman, Gene M and Giovanni Maggi**, “Diversity and Trade,” *American Economic Review*, December 2000, 90 (5), 1255–1275.
- Grossman, Gene M., Elhanan Helpman, and Adam Szeidl**, “Optimal integration strategies for the multinational firm,” *Journal of International Economics*, 2006, 70 (1), 216–238.
- Hall, Robert E. and Charles I. Jones**, “Why do some countries produce so much more output per worker than others?,” *Quarterly Journal of Economics*, 1999, 114, 83–116.
- Halpern, László, Miklós Koren, and Adam Szeidl**, “Imported inputs and productivity,” *Working Paper, CEU*, 2011, 8, 28.
- Helpman, Elhanan, Marc Melitz, and Yona Rubinstein**, “Estimating Trade Flows: Trading Partners and Trading Volumes,” *The Quarterly Journal of Economics*, May 2008, 123 (2), 441–487.
- Jia, Panle**, “What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry,” *Econometrica*, November 2008, 76 (6), 1263–1316.
- Johnson, Robert C. and Guillermo Noguera**, *Fragmentation and trade in value added over four decades*, Working Paper, Dartmouth College, January 2012.
- Loecker, Jan De, Pinelopi K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik**, “Prices, Markups and Trade Reform,” Working Paper 17925, National Bureau of Economic Research March 2012.
- Melitz, Marc J.**, “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 2003, 71 (6), 1695–1725.
- Milgrom, Paul and John Roberts**, “The Economics of Modern Manufacturing: Technology, Strategy, and Organization,” *The American Economic Review*, June 1990, 80 (3), 511–528.
- Monarch, Ryan**, “It’s not you it’s me: Breakups in U.S.-China Trade Relationships,” Working Paper, University of Michigan 2013.
- Morales, Eduardo, Gloria Sheu, and Andrés Zahler**, “Gravity and Extended Gravity: Using Moment Inequalities to Estimate a Model of Export Entry,” Working Paper 19916, National Bureau of Economic Research February 2014.
- Olley, G Steven and Ariel Pakes**, “The Dynamics of Productivity in the Telecommunications Equipment Industry,” *Econometrica*, 1996, 64 (6), 1263–97.
- Oostendorp, Remco H.**, “The Standardized ILO October Inquiry 1983-2003,” mimeo, Amsterdam Institute for International Development 2005.

- Ramondo, Natalia and Andrés Rodríguez-Clare**, “Trade, Multinational Production, and the Gains from Openness,” *Journal of Political Economy*, April 2013, *121* (2), 273–322.
- Rodríguez-Clare, Andrés**, “Offshoring in a Ricardian World,” *American Economic Journal: Macroeconomics*, April 2010, *2* (2), 227–258.
- Rossi-Hansberg, Esteban and Mark L. J Wright**, “Establishment Size Dynamics in the Aggregate Economy,” *American Economic Review*, December 2007, *97* (5), 1639–1666.
- Tintelnot, Felix**, “Global Production with Export Platforms,” 2012. mimeo Penn State University.
- Train, Kenneth**, *Discrete Choice Methods with Simulation*, 2nd ed., Cambridge University Press, 2009.
- Vives, Xavier**, “Nash equilibrium with strategic complementarities,” *Journal of Mathematical Economics*, 1990, *19* (3), 305–321.
- Yeaple, Stephen Ross**, “The complex integration strategies of multinationals and cross country dependencies in the structure of foreign direct investment,” *Journal of International Economics*, 2003, *60* (2), 293–314.