

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

December 13, 2014

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. 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. In contrast to canonical models of exporting in which firm profits are additively separable across exporting markets, we show that global sourcing decisions naturally interact through the firm's cost function. In particular, the marginal change in profits from adding a country to the firm's set of potential sourcing locations depends on the number and characteristics of other countries in the set. Still, under plausible parametric restrictions, selection into importing features complementarity across markets and firms' sourcing strategies follow a hierarchical structure analogous to the one 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, Ariel Burstein, Arnaud Costinot, Pablo Fajgelbaum, Paul Grieco, Gene Grossman, Elhanan Helpman, Sam Kortum, Marc Melitz, Eduardo Morales, and Michael Peters for useful conversations, to Andrés Rodríguez-Clare for his comments while discussing the paper at the NBER, 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, Brown, Chicago Booth, Dartmouth, ECARES in Brussels, Econometric Society Meeting in Minneapolis, ERWIT in Oslo, Harvard, John Hopkins SAIS, National Bank of Belgium, Michigan, MIT, NBER Summer Institute, Northwestern, Princeton, SED in Toronto, Tsinghua, UBC, Urbana-Champaign, 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 demonstrates the existence of a large degree of intraindustry heterogeneity in revenue, productivity, factor inputs, and export participation across firms. With regard to export behavior, the studies show that only a small fraction of exceptionally productive firms engage in exporting, and that most of them 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

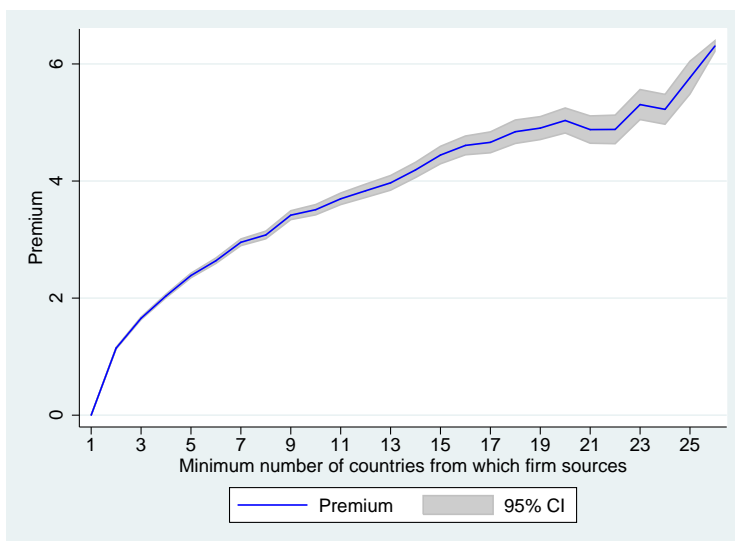
Figure 1 provides a graphical illustration of the size premia of importers and how they vary with the number of countries from which a firm sources.² The figure indicates that firms that

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 countries 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 the extensive and intensive margins of importing. 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.

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.

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
South Korea	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 from which a firm is able to import 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 costs 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 firm 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 requires 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.⁴ The attractiveness of a sourcing location is shaped not only by wages, technology, transport costs, but also by the fixed cost of offshoring to that location.

We also use the model to study how the aggregation of firms' sourcing decisions shapes aggregate input flows across countries. In doing so, we show that our model nests several key workhorse trade models developed in recent years. When we set all fixed costs of sourcing to zero, our model ceases to feature an extensive margin and delivers a gravity equation for aggregate trade flows identical to that in Eaton and Kortum (2002), with the elasticity of trade flows with respect to changes in variable trade frictions coinciding with the Fréchet parameter governing input productivity dispersion. As in heterogeneous firms models of exporting, in the presence of fixed costs of sourcing, the endogenous determination of 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, our model delivers a gravity

⁴The 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.

equation for input flows that is identical to those derived by Chaney (2008), Arkolakis et al. (2008), and Helpman et al. (2008) in their multi-country, Pareto extensions of the Melitz (2003) framework 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 hints at the existence of biases when estimating 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 66 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 cross-country wage data to estimate the Fréchet parameter governing the dispersion in productivity of intermediate good producers. In this second step, we also back out the elasticity of final-good demand from data on average mark-ups in U.S. manufacturing. 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 seven percent lower variable costs and achieves 24 percent higher sales than when sourcing exclusively from domestic suppliers.⁵ In our third and final step, we use the method of simulated moments to estimate firm-country-specific fixed costs that depend on country characteristics. 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. Our median fixed cost of sourcing estimates range from 9,000 to 29,000 USD, are around 40 percent lower for countries with a common language, and increase in distance with an elasticity of 0.34. Our estimated model targets moments related to the extensive margin of global sourcing, but we show that it also provides a very good fit of the shares of aggregate foreign sourcing by country, which are non-targeted in the estimation.

Another advantage of adopting Jia’s (2008) approach is that it not only provides a tractable way of solving a large discrete choice problem with many firms, but it also produces point estimates that allow one to conduct counterfactual exercises.⁶ We exploit this feature by studying the implications of a 100% increase in China’s sourcing potential (perhaps due to a reduction in bilateral trade costs between these two countries). We show that, consistent with other quantitative models of trade, such a shock tends to increase import competition from China, decrease the equilibrium industry-level U.S. price index, and drive some U.S. final good producers out of the market. Although the net result of these forces is a marked decrease in U.S. domestic sourcing (and U.S. employment in that sector), this net decline masks the existence of significant heterogeneity in how the shock affects the sourcing decisions of firms at different points in the size distribution. More specifically, a range

⁵Quantitatively, 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.

⁶This is in contrast to moment inequality methods, which were first adopted in an international trade context by Morales et al. (2014).

of U.S. firms is induced to select into sourcing from China as a result of the shock, and we find that these firms on average increase their input purchases not only from China, but also from the U.S. and other countries. The existence of gross changes in sourcing that operate in different directions is a distinctive feature of our framework and does not arise in the absence of fixed costs of offshoring (i.e., in the Eaton-Kortum limit of our model) or whenever entry decisions are independent across markets (i.e., in the Melitz-limit of our model). Furthermore, the quantitative implications of the increase in China’s sourcing potential appear to be quite distinct relative to a situation in which the extensive margin of offshoring is held constant following the shock.

Our paper is related to several strands of the literature. We build on the approach in Tintelnot (2014) 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 (2014) 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.⁷

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 De Loecker et al. (2012) for India, Halpern et al. (2011) for Hungary, and Gopinath and Neiman (2013) for Argentina.⁸ 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 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 due to 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. Two recent papers by Blaum et al. (2014) and Ramanarayanan (2014) develop multi-country quantitative models similar to ours to study the effect of imported inputs on firm-level and aggregate productivity, but those frameworks make restrictive assumptions to bypass a general analysis of the extensive margin of sourcing.

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 and cross-country heterogeneity in the fixed costs of sourcing. 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

⁷See also Ramondo and Rodríguez-Clare (2013) and Arkolakis et al. (2013) for recent related contributions to quantitative models of multinational activity.

⁸See De Loecker and Goldberg (2014) for an insightful overview of this literature.

our model, but they do not attempt to characterize the extensive margin of importing.

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 (2014) 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.⁹

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 our counterfactual exercises in section 6. Section 7 concludes.

2 Theoretical Framework

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

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

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

⁹The 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), Monarch (2013), and Dragusanu (2014).

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 $a_j(v, \varphi) w_j$.¹⁰

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

¹⁰Implicitly, 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.

intermediate input efficiencies $1/a_j(v, \varphi)$ as the realization of an extreme value distribution. More specifically, firm φ draws the value of $1/a_j(v, \varphi)$ for a given location j from the Fréchet distribution

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

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. It will be convenient to denote by $\mathbf{a}_j(\varphi) \equiv \{a_j(v, \varphi)\}_{v=0}^1$ the vector 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.

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.

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 $\min_j \{\tau_{ij} a_j(v, \varphi) w_j\}$. 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, in equation (6), 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 .¹¹

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 cumbersome 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.¹² 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

¹¹It 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.

¹²These 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. To ensure a well-defined marginal cost index, 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.

price paid for a country's variety is unaffected by the inclusion of other countries in the set $\mathcal{J}_i(\varphi)$.¹³

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 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 how we exploit key predictions of theory and evidence from the data 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 firm's sourcing capability. This case is more likely to apply whenever demand is elastic and thus profits are particularly responsive to variable cost reductions (high σ), and

¹³See, 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 empirically justify our modeling choices in section 4.1.

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 firm's sourcing capability. 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.¹⁴

Regardless of parameter values, we can use the fact that the profit function (10) is supermodular in φ and $\Theta_i(\varphi)$ and in the Appendix 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 $\varphi^{\sigma-1}$.

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 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.

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 subset of the sourcing strategy of relatively

¹⁴Readers familiar with the work of Eaton and Kortum (2002) might be led to believe that $\theta > \sigma - 1$ is in fact implied by the need for the firms marginal cost function to be well-defined. Note, however, that our parameter ρ plays the role of σ in the Eaton-Kortum setup, and thus this technical condition is instead $\theta > \rho - 1$ in our setup (see footnote 12).

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 $\mathcal{J}_i(\varphi) = \{j : I_{ij}(\varphi) = 1\}$.

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 be clear that if countries differ only along one dimension (either sourcing potential or fixed costs), then they can be uniquely ranked in terms of their ‘sourcing appeal’ for all firms. We will illustrate this shortly for the case of common fixed costs. Conversely, 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 (as in Blaum et al. (2014)), 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, we have $i_1 \in \mathcal{J}_i(\varphi)$; 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.

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 that these sharp results only apply when holding market demand – of which the price index is a key component – fixed. In general equilibrium, however, these same parameter changes also affect the level of market demand. As we shall see in our counterfactual exercises in section 6, the endogenous response of market demand is quantitatively important in our estimation, and thus the implications we derive from changes in trade costs are much more nuanced than those discussed above (see Bache and Laugesen (2013)).

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 on 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. 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 thus can treat wages as exogenous in solving for the equilibrium in each country’s manufacturing sector.

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 in hand, the equilibrium number of active firms is simply given by $N_i [1 - G_i(\tilde{\varphi}_i)]$.¹⁵

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 :

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

¹⁵In 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 result is analogous to that derived in canonical models of selection into exporting (see, for instance, Arkolakis et al. (2012)), but note that it here applies in a setup with interdependent entry decisions.

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. 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 simplify the equilibrium values for B_i and N_i in (12) and (13) considerably. 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.¹⁶

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. 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 first to consider the knife-edge case in which $\sigma - 1 = \theta$, and thus the extensive margin of global sourcing can be studied

¹⁶The 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 $N_i = \eta L_i / \sigma f_{ei}$.

independently across countries.

Independent Entry Decisions Whenever $\sigma - 1 = \theta$, one can use the formula for χ_{ij} in (6) to simplify (14) to ¹⁷

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

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 still 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 (since $\kappa > \sigma - 1 = \theta$).

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), where an analogous result applies. 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 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).

¹⁷To 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 $\varphi > \tilde{\varphi}_{ij}$.

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 show 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 parameter values of 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

Before turning to the empirical evidence, 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 limit ourselves to studying the effects of these extensions on firm behavior and relegate the mathematical details to the Online Appendix.

Exporting and Importing We first consider a reduction of trade costs for final-good varieties from their prohibitive levels to a combination of (bounded) iceberg trade costs τ_{ij}^X and fixed costs f_{ij}^X of exporting. Firm behavior conditional on a sourcing strategy is largely analogous to that in section 3.1, except for the inclusion of a determination of the firm’s ‘exporting strategy’. In the Online Appendix we show that, 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 a reduction in any τ_{ij} or an increase in any T_j – will necessarily lead to a (weak) increase in the extensive margin of exporting. Furthermore, 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 global sourcing. Finally, as firm productivity increases, the participation of firms in both export and import markets increases, but at a faster rate than when either of those margins 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.

Endogenous Input Variety Our benchmark model assumes that all final good producers use a measure one of inputs. It is simple to generalize our framework 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)). In order to create a check on the optimal degree of complexity, we assume that firms face input-specific fixed costs (in addition to the country-specific costs in our benchmark model). 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

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.

patterns in Figure 1 (which persist even when controlling for the number of imported inputs) 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 on the products, values, and countries of firms' imports. We match these data at the firm level to the 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,

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.

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.

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. About one quarter of U.S. manufacturing firms have positive imports in 2007. Additional details on the sample and data construction are in the Online Data Appendix.

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 and wages, as well as other controls—from various sources. Country R&D data and the number of private firms in a country for 2007 are from the World Bank Development Indicators. Wage data are from the ILO data described by Oostendorp (2005). 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 described by Heston et al. (2011). Control of corruption is from the World Bank’s Worldwide Governance Indicators. We also obtain years of schooling and 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 document a hierarchical pattern in firms’ 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 41 for the number of imported products and 11 for the number of source countries.

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

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. Note that data confidentiality protection rules preclude us from disclosing exact percentiles. Statistics for all percentiles in the paper are therefore the average for all firms that are within +/- one percent of the reported percentile.

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. 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 close to 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 Online 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 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 from these top ten countries 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 idiosyncrasies 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.

5 Structural Analysis

In this section, we use the firm-level data in conjunction with country-level data to estimate the key parameters of the model. In doing so, we distinguish country sourcing potential from the fixed costs of sourcing and quantify the extent to which the latter depend upon distance and common

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.

language. 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 estimate 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 the method of simulated 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.

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 with at least 200 U.S. importing firms. This criterion leaves us with a total of 67 countries, including the U.S.

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}^n 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}^n , is simply that country’s contribution to the firm’s sourcing capability, Θ_i^n . 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

$$\log \chi_{ij}^n - \log \chi_{ii}^n = \log \xi_j + \log \epsilon_j^n, \quad (20)$$

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 .²¹

The dependent variable in equation (20) is the difference between a firm’s share of inputs sourced from country j and its share of inputs sourced domestically. We measure these shares using data on a firm’s total input use as well as the imports from each country from which it sources. We include firms that import from countries with fewer than 200 U.S. importers in the estimation, adjusting their total input usage by subtracting their imports from any of the excluded countries. Additional details on our measure of input shares are in the Estimation Appendix.²²

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 the regression. This condition will be satisfied if firms only learn their country-specific efficiency shocks, ϵ_j^n , after their sourcing strategy is selected, or if the term ϵ_j^n simply represents measurement error.²³ It is also consistent with firm-country-specific shocks to the fixed costs of sourcing.

Table 5 provides summary statistics from estimating equation (20) 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.

Figure 2 plots the estimated sourcing potential fixed effects against total input purchases (left panel) and against the number of firms importing from that country (right panel). Our parameter

²¹When normalizing by the domestic share, we set domestic sourcing potential to 1. All other country variables are then measured relative to the U.S. value.

²²We note that a very small fraction of firms has negative values for its implied 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. In an alternative approach, we have limited the sample of firms in the structural analysis to firms with at least fifty percent of their sales in manufacturing. This ensures better measurement of firms’ inputs and helps to address concerns about imports being true inputs into production. We do not report the numbers from this analysis here to avoid potential future disclosure issues in revisions of the paper. We do note that the results are qualitatively similar and we may report them in future work.

²³Note that this assumption rules out measurement error related to a firm’s global sourcing strategy. In other words, we assume that the set of countries from which the firm imports is correctly observed and that a firm has positive imports from all countries for which it has paid a fixed cost of sourcing.

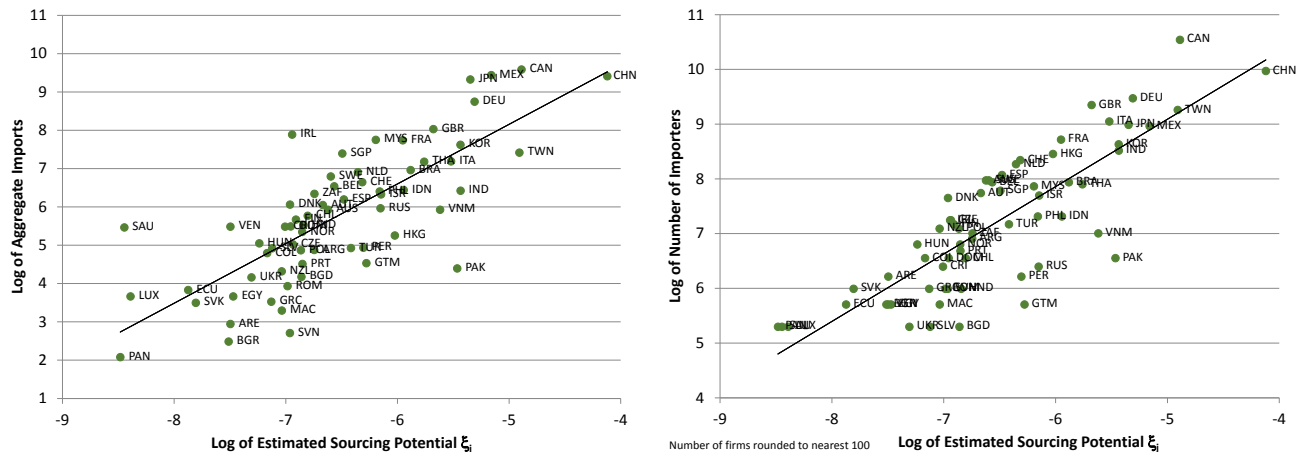
estimates suggest that China has the highest sourcing potential for U.S. firms, followed by Canada and Taiwan. More firms import from Germany and United Kingdom than from Taiwan, however, and more firms import from Canada than from China, suggesting that fixed costs of sourcing are likely to differ across source countries. Despite some heterogeneity, the number of firms and total import purchases are clearly positively associated with a country’s sourcing potential, with a tighter relationship between the sourcing potential and the number of firms sourcing from a country.

Table 5: Summary statistics for sourcing potential estimation

Number of observations	200,000
Mean Squared Error	2.64
Range of foreign $\log \xi_j$	- 4.12 to -8.42
Sum of foreign ξ_j	0.137

Notes: Summary statistics for regression based on equation (20). Estimated Fixed effects are displayed in Figure 2. Number of observations rounded for disclosure avoidance. The R^2 with a constant is 0.09.

Figure 2: Country sourcing potential parameters



Our estimates of the sourcing potential of a country enable us to calculate the extent to which the sourcing capability of a firm $\Theta^n = \sum_{j \in \mathcal{J}^n} \xi_j$ is higher if it imports from all countries as opposed to sourcing only domestically. Since the domestic sourcing potential is normalized to one and the summation of the foreign sourcing potential terms is 0.137, these results imply that the sourcing capability of a firm that sources from all 67 countries is 13.7 percent larger than that of a firm sourcing only domestically. The impact of a firm’s sourcing capability on its marginal cost in turn depends on the dispersion parameter θ of the intermediates productivities, as seen in equation (8). Equation (11) shows that the effect of sourcing capability on firm sales also depends on θ , as well

as on the elasticity of substitution, σ . We now turn to estimating 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. 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 2007 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 0.0005. This implies an estimate for the elasticity of demand, σ , of 3.85. Of course it is impossible to distinguish perfectly between fixed and variable costs in the data, and there may be certain costs that simply are not measured well, but we view this as a plausible estimate that is similar to previous estimates.²⁴ Given the potential issues that may affect the accuracy of our estimate of the demand elasticity, we include a sensitivity analysis in section 6.3.3 in which we consider alternative values for the elasticity of demand, as well as other parameters.

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 use data on wages to identify this elasticity. Recall that the sourcing potential ξ_j that we estimated in the previous section, is a function of a country’s technology parameter, trade costs, and wages. 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, and common language. Specifically, we estimate the following equation:

$$\begin{aligned} \log \hat{\xi}_j = & \beta_0 + \beta_r \log \text{R\&D}_j + \beta_k \log \text{capital}_j + \beta_C \text{control of corruption} + \beta_F \log \text{firms} \\ & - \theta \log w_j - \theta (\log \beta_c + \beta_d \log \text{distance}_{ij} + \text{language}_{ij} \log \beta_l) + \iota_j. \end{aligned} \quad (21)$$

Equation (21) shows that the parameter θ can be recovered from the estimated coefficient on wages. In theory, one could also identify θ using tariffs, but, as we show in the Estimation Appendix, there is not enough variation in U.S. tariffs to do so. 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

²⁴For example, Broda and Weinstein (2006) estimate a mean elasticity of 4 and a median of 2.2 at the SITC-3 level for 1990-2001. At the SITC-4 level, Feenstra and Romalis (2014) estimate a higher median elasticity of 6.2 for 1984-2011. Our estimate falls within this range.

of labor, we follow Eaton and Kortum (2002) and use a human capital-adjusted wage. Details of this adjustment are in the Estimation Appendix. 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. One concern with using population as an instrument is that it may violate the exclusion restriction if high population countries are also technologically advanced countries. With that in mind, we include country R&D stock, a level measure of technology, in all specifications. Country population may also indirectly affect sourcing potential since it could be correlated with the number of potential suppliers in a country, a concern that leads us to control for the number of private firms in the economy. These country-level variables are available for 56 of the 66 foreign countries included in the structural estimation.

The first column of Table 6 presents the results from estimating equation (21) via OLS. Column 2 provides the analogous IV estimates, using population as an instrument for wages. As expected, the IV estimate for θ (1.71) is larger than the OLS estimate. In line with the 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 estimate a coefficient of 4.56 when using the same specification as in equation (21) but with aggregate imports as the left-hand-side variable. These results are displayed in columns 4 and 5. The Estimation Appendix presents the first stage regression, along with additional robustness tests in which we control for GDP and tariffs, and in which we constrain the coefficient on wages and tariffs to be the same.

With estimates for country sourcing potential, the firm-level trade elasticity, and the elasticity of demand in hand, we can calculate how global sourcing affects firm costs and size. Our estimates imply that a firm sourcing from all countries faces around seven percent ($1.137^{(-1/1.71)}$) lower input costs than a firm sourcing purely domestically, and consequently its sales are around 24 percent ($1.137^{(-2.85/1.71)}$) larger.

Across various specifications, including additional robustness tests included in the Estimation Appendix, 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. In the third step of our estimation, we exploit this feature of our model and the data to solve the firm’s problem numerically and thereby estimate the fixed costs of sourcing from different markets.

5.3 Step 3: Estimation of Fixed Costs of Sourcing

In this section, we estimate the fixed costs of sourcing via the method of simulated moments. As is common in the literature that estimates trade costs, we allow the fixed cost of sourcing from a country to depend on the gravity variables distance and language. Our model implies that the number of importing firms is identical to the number of firms that sources from the most popular

Table 6: Estimation of firm and aggregate trade elasticities

	log ξ		log aggregate imports	
	OLS	IV	OLS	IV
HC adjusted wage	-0.50*** (0.17)	-1.71** (0.68)	-0.61 (0.39)	-4.56** (1.92)
log distance	-0.31 (0.19)	-0.57** (0.28)	-0.88** (0.42)	-1.73** (0.79)
common language	0.13 (0.21)	0.19 (0.28)	0.29 (0.48)	0.48 (0.78)
log R&D	0.36*** (0.07)	0.54*** (0.13)	0.72*** (0.15)	1.30*** (0.36)
log KL	-0.13 (0.17)	0.44 (0.38)	-0.23 (0.37)	1.64 (1.05)
Control of corruption	0.14 (0.14)	0.59* (0.31)	0.34 (0.32)	1.82** (0.86)
log no. of firms	0.13 (0.09)	-0.01 (0.14)	0.10 (0.20)	-0.37 (0.39)
Constant	-4.71*** (1.07)	-5.97*** (1.55)	11.07*** (2.41)	6.98 (4.35)
Observations	56	56	56	56

Notes: Standard errors in parentheses. In the IV specifications, the human capital-adjusted wage is instrumented by population.

country. This is an extreme prediction driven by the assumption that firms face an identical fixed cost to source from a particular market. In contrast, the data show that around 60,000 firms import, but only about 38,700 import from Canada – the most popular sourcing country. For the remainder of the structural analysis, we therefore enrich the model by allowing the fixed costs of sourcing to vary by firm-country combinations. Specifically, we model firm-country-specific fixed costs, f_{ij}^n , that are drawn from a lognormal distribution with dispersion parameter β_{disp}^f and scale parameter $\log \beta_c^f + \beta_d^f \log \text{distance}_{ij} + \log \beta_l^f \text{language}_{ij}$.

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 67 countries, the firm selects between 2^{67} , which is roughly 10^{20} , possible sourcing strategies. Clearly, calculating the profits of each of these strategies for every firm is infeasible. To reduce the dimensionality of the problem, we use the empirical evidence that $\sigma - 1 > \theta$ to infer that firms’ profit functions feature increasing differences in their sourcing strategies. This allows us to adapt 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 guess \mathcal{J} for the firm’s sourcing strategy, \mathcal{J}^n , the marginal benefit of including country j in the sourcing strategy \mathcal{J} is:

$$\begin{cases} \varphi^{\sigma-1} \gamma^{(\sigma-1)/\theta} B \left(\Theta_i(\mathcal{J} \cup j)^{(\sigma-1)/\theta} - \Theta_i(\mathcal{J})^{(\sigma-1)/\theta} \right) - f_{ij}^n & \text{if } j \notin \mathcal{J} \\ \varphi^{\sigma-1} \gamma^{(\sigma-1)/\theta} B \left(\Theta_i(\mathcal{J})^{(\sigma-1)/\theta} - \Theta_i(\mathcal{J} \setminus j)^{(\sigma-1)/\theta} \right) - f_{ij}^n & \text{if } j \in \mathcal{J}. \end{cases}$$

We define a mapping, $V_j(\mathcal{J})$ that 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 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), an iterative application of the V-operator that adds each country to the set one-by-one leads to a lower bound of the firm’s sourcing strategy. That is, the optimal sourcing strategy contains at least those countries for which the marginal benefit of adding a country is positive when that country is added individually. Similarly, when starting from the set \mathcal{J}^1 (which contains all countries), and removing individual countries one-by-one, the iterative application of the V-operator leads to an upper bound for the optimal sourcing strategy. Should the two sets not perfectly overlap, it is only necessary to evaluate the profits resulting from all possible combinations contained in the upper but not the lower bound set. 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 bound sets in the vast majority of simulations. In addition, the two sets only differ by a small number of countries in those cases in which the sets do not completely overlap (see Table B.7 in the Estimation Appendix). 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.

As in Melitz and Redding (2013), we fix the shape parameter of the Pareto distribution. We set $\kappa = 4.5$, but sensitivity results with alternative values for κ are discussed in section 6.3.3. We are left with the following five parameters to estimate: $\delta = [B, \beta_c^f, \beta_d^f, \beta_l^f, \beta_{\text{disp}}^f]$.²⁵ To do so, we simulate a large number of U.S. firms. That is, for each firm we draw a core-productivity shock from a uniform distribution (which, given κ , can be inverted to yield the Pareto-distributed firm core productivity level), and a J -dimensional vector of fixed cost shocks from a standard normal distribution (which, given a parameter guess β^f , can be used to calculate the lognormal distributed firm-country specific fixed cost level).²⁶ 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

²⁵We set $\varphi_{\text{US}} = 1$, as it scales input purchases equivalently to an increase in B .

²⁶We use a stratified random sampling technique to simulate the Pareto draws, in which we simulate many more points in right tail of the distribution (in total 12 intervals with 10 random draws each). For the fixed cost draws we use a Hybrid-Quasi-Monte-Carlo procedure, in which we generate a vector of 18,000 quasi-random numbers from a van der Corput sequence in one dimension (which have better coverage properties than usual pseudo-random draws), and then for each country we use this vector, but with independent random permutations of elements of this vector. Each core productivity draw is then interacted with a vector of fixed cost draws, which together represent a firm. In total, the interaction of fixed cost and core productivity draws yields $S = 2,160,000$ simulated firms.

distributions.

We use the simulated firms to construct three sets of moments. The first moment is the share of importing firms (about 24 percent in the data). This is simply a scalar, which we label as m_1 in the actual data and as $\hat{m}_1(\delta)$ for the simulated data. The second set of moments is the share of firms that sources from each country. We label this $J \times 1$ vector of moments in the data as m_2 and the simulated moment vector as $\hat{m}_2(\delta)$. This set of moments is informative about the overall magnitude of the fixed costs of sourcing, as well as on how they vary with distance and language. In addition, the share of importing firms from the most popular country relative to the total share of importers is indicative about the fixed cost dispersion parameter. The intuition here is that if there were zero dispersion in fixed costs across firms, the share of importers would be identical to the share of importers from the most popular sourcing country. Finally, the third moment is the share of firms whose input purchases from the U.S. are less than the median U.S. input purchases in the data. This is also a scalar, and we label the moment in the data as m_3 and the simulated moment as $\hat{m}_3(\delta)$. The information from this moment helps pin down the scale parameter B .

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

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

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

$$E[\hat{y}(\delta_0)] = 0. \tag{22}$$

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

$$\hat{\delta} = \arg \min_{\delta} [\hat{y}(\delta)]^\top \mathbf{W} [\hat{y}(\delta)], \tag{23}$$

where \mathbf{W} is a weighting matrix. We weight the moments equally, hence the weighting matrix is the identity matrix.

The parameter estimates are displayed in Table 7 below. We find that the fixed costs of sourcing are increasing in distance with an elasticity of 0.34, and that sourcing from countries with a common language reduces fixed costs by about 40 percent. The fixed costs of sourcing also seem reasonable in magnitude. The median fixed cost estimate ranges from 9,000 to 29,000 USD, though the assumption of a lognormal distribution means they can be substantially larger for some individual firm-country combinations. We conclude this section by describing the fit of the data by the estimated model.

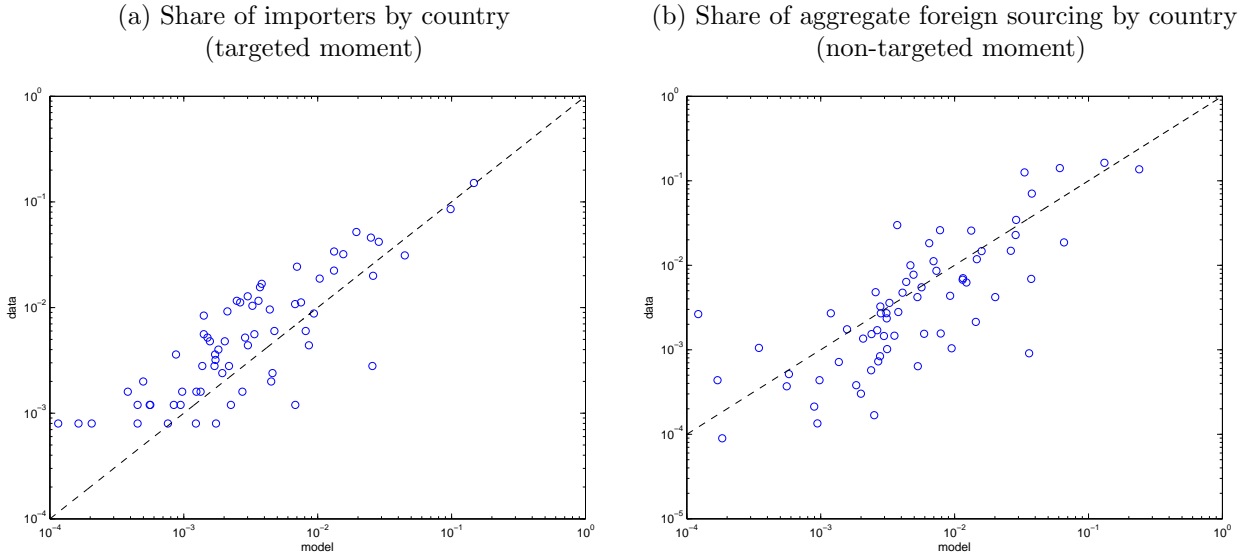
Table 7: Estimated parameters

B	β_c^f	β_d^f	β_l^f	β_{disp}^f
0.125	0.011	0.338	0.611	0.865

5.4 Fit of the Model

Overall, the model fits the data reasonably well. We start by comparing the predictions of the model for the moments it was targeted to fit. First, in both the actual and simulated data, approximately 24 percent of U.S. firms import (24.1 in the data and 24.4 in the model). For the second set of moments on the share of importing firms by country, the correlation coefficient between the actual and simulated data is 0.94. Figure 3a depicts this relationship by country. Finally, in both the data and in the parameterized model, the median firm’s input purchases from the U.S. are equal to approximately \$560,000 (562,000 in the data and 569,000 in the model). Importantly, the model also does a good job of matching the aggregate shares of sourcing by country. These shares were not targeted directly in the estimation, but still have a correlation coefficient of 0.78 between the actual and simulated data. This relationship is depicted in figure 3b.

Figure 3: Model fit: Share of firms and Aggregate sourcing



6 Counterfactual: An Increase of China’s Sourcing Potential

In this section, we use the parameter estimates from section 5 to assess how the sourcing decisions of different types of firms, the firm size distribution, and aggregate sourcing decisions by country are affected by a shock to China’s sourcing potential. We demonstrate that it is important to

distinguish between gross and net changes to sourcing, as some firms may increase their domestic purchases while others may reduce them. We pay particular attention to third-market effects – that is how a shock to one country affects sourcing from other countries (including sourcing from the U.S.). The 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 in producing goods for U.S. export. It is important to emphasize, however, that for tractability reasons, we consider a unilateral shock relative to the U.S. and not a multilateral shock that affects all countries (it is for this same reason that our model features a non-manufacturing sector that pins down wages). Throughout the counterfactual exercises we solve for the new equilibrium price index and let the mass of firms adjust such that the free entry condition is satisfied.

6.1 Baseline Predictions

We simulate a large increase (100 percent) of China’s sourcing potential. Clearly, new and continuing importers from China will benefit from this shock. These firms will sell their products at lower prices which leads the aggregate price index and the mass of active firms to adjust. We first document how these changes affect sourcing from third markets for different sets of firms. Table 8 shows that an additional 14.3 percent of firms start importing from China. Interestingly, these firms simultaneously increase their sourcing from the U.S. and from third countries, by one and 2.3 percent, respectively. Firms that continue sourcing from China comprise 9.8 percent of firms, and slightly decrease their domestic and third country sourcing. Firms that do not import from China face tougher competition as firms importing from China lower their costs, so their sourcing from the U.S. and other sourcing countries contracts by 1.1 and 3 percent, respectively. The responses to sourcing from other foreign sourcing countries tend to be larger as firms either begin or stop sourcing from these countries; this extensive margin effect is absent for domestic sourcing.

Table 8: THIRD COUNTRY SOURCING EFFECTS OF POSITIVE CHINESE SOURCING POTENTIAL SHOCK

Chinese import status	Change sourcing from US	Change Sourcing from other countries	Change Sourcing from China	Share of firms
Entrants	1.010	1.023	∞	0.143
Exiters	-	-	-	0
Continuers	0.999	0.998	1.998	0.098
Others	0.989	0.970	-	0.759

Notes: This table contains only surviving firms. Entrants (exiters) are those firms that begin (stop) offshoring to China. Columns 1, 2, and 3 contain the ratio of the total sourcing by this group of firms before and after.

New importers from China have the strongest relative increase in their sales, while firms that do not import from China contract. Figure 4a depicts the percent growth in sales by firm size

Table 9: GROSS AND NET US SOURCING EFFECTS

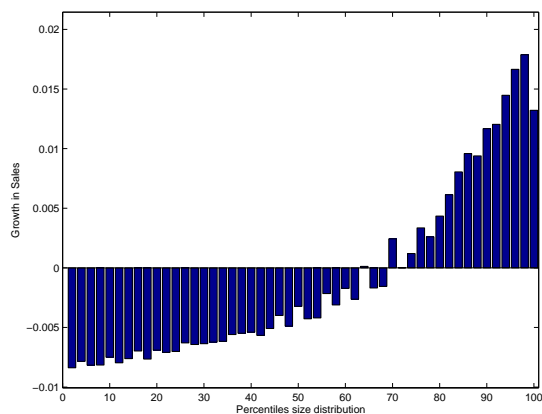
	Baseline		Fixed Sourcing strategies	
	Difference in sourcing in billion USD	Change in percent of total US sourcing	Difference in sourcing in billion USD	Change in percent of total US sourcing
Increase in domestic sourcing	6.17	0.16	1.53	0.04
Decrease in domestic sourcing by firms that continue to operate	-22.70	-0.59	-22.58	-0.59
Decrease in domestic sourcing by firms that shut down	-23.78	-0.62	0	0
Total	-40.31	-1.06	-21.05	-0.55

Notes: We use the model’s predictions for the U.S. sourcing pre and post the shock to China’s sourcing potential in order to calculate percentage changes in U.S. sourcing. We then use aggregate U.S. sourcing purchases in the data and the percentage differences predicted by the model to calculate the predicted USD change in sourcing in response to the shock.

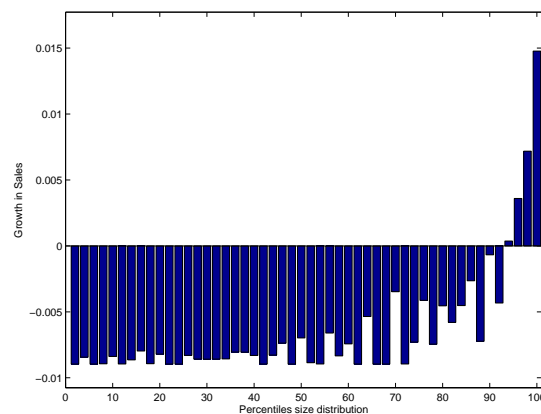
percentile. Firms are ranked according to their sales before the China shock. The relative changes in firm sales take an S-shape. Large firms increase their sales on average by 1-2 percent, with the largest 2 percent of firms growing slightly less than smaller firms. Note that while these magnitudes seem small, this is partly driven by the fact that fixed costs are firm-country-specific and hence not all firms within a certain size category import from China. Figure 4 therefore shows average size changes for firms that do and do not import from China. Firms below the 75th percentile of the size distribution see their sales shrink on average in response to the China shock.

Figure 4: Changes in the size of firms

(a) Baseline - Changes in the size of firms in response to 100% shock to China’s sourcing potential



(b) Fixed Sourcing Strategies - Changes in the size of firms in response to 100% shock to China’s sourcing potential



Aggregate imports from China increase by 195 percent while aggregate sourcing from the U.S. falls by 1.06 percent. The net change in U.S. sourcing masks a substantial amount of domestic churning (displayed in the left panel of Table 9). Calculating the effect of the shock separately for firms that expand versus those that contract or exit domestically shows that for every one dollar reduction in domestic sourcing, growing firms expend an additional 13 cents. Since in our model this spending represents jobs in the intermediate good sector, we predict a substantial amount of churning as a result of an increase of the Chinese sourcing potential. In terms of welfare, we find that an increase in China’s sourcing potential lowers the price index in the manufacturing good sector by 0.4 percent.

6.2 Fixed Sourcing Strategies

An important emphasis of our model is the role of the extensive margin in firms’ sourcing strategies. To assess the importance of this margin in aggregate trade patterns and welfare, we repeat the same counterfactual exercise but hold firms’ extensive margin sourcing strategies constant. In the exercise, we recompute the equilibrium price index and the mass of firms. As expected, when firms’ sourcing strategies are fixed, the aggregate response to the China shock is substantially smaller (see Table 9). U.S. sourcing decreases by 0.55 percent and aggregate sourcing from China increases by 100 percent, which is about half as much as the response under flexible strategies.

The micro effects by firm type are helpful to understand these differences. In the baseline, the large number of new importers from China drives most of the aggregate Chinese sourcing increase. This increase leads the price index to fall significantly, which reduces sourcing from both the U.S. and other countries. In addition, in the baseline analysis the increase in the firms’ extensive margin sourcing decisions results in higher expenditures on fixed costs. Given constant expenditure on the manufacturing sector, this leads *ceteris paribus* to lower expected profits, which in turn are brought up again by a decrease in the number of firms. In contrast, under fixed sourcing strategies, the number of firms remains the same in equilibrium and consequently there is no decrease in U.S. sourcing from firms that shut down (see Table 9). Firms that source from China slightly increase their sourcing from the U.S. and other countries by 0.1 percent, while all other firms reduce their sourcing from the U.S. and other countries by 0.9 percent (see Table B.3 in the Appendix).

The responses of the size distribution are also quite different (see Figure 4b). When firms’ sourcing strategies are fixed, the only firms that grow are those that previously imported from China, which on average tend to be the larger firms. Firms of all other sizes contract. We find that under fixed sourcing strategies the price index declines by .32 percent, which is about 80 percent of the impact of the shock when firms can adjust their extensive margin decisions.

6.3 Alternative Parameter Values

As discussed in Section 3.4 when describing the model’s predictions on aggregate trade flows, our framework nests two canonical frameworks of the last decade: the Melitz (2003) heterogeneous firms model and the Eaton and Kortum (2002) Ricardian trade model. In order to demonstrate

the qualitative differences between our baseline model and these other models, we modify the parameters such that these canonical models emerge as special cases. An important caveat is that, because these exercises require changing parameter values, it is clearly not possible to hold all else equal. We therefore focus on the qualitative differences in the predictions and caution that quantitative differences may reflect several changing factors. We report the headline numbers here and provide complete tables for these exercises in the Online Appendix.

6.3.1 Independent Entry Decisions

First, we consider a scenario similar to Chaney (2008) in which firms' import decisions are independent across markets, which occurs in the knife-edge case of $\theta = \sigma - 1$ in our model. To do so, we increase the value of θ to 2.85 so that it is equal to $\sigma - 1$. We re-do the estimation of the fixed costs under these parameter values in Step 3, leading to a similar model fit as in the baseline. Note that because we deviate from the estimated value of θ in Step 2, this parameter value fits the firm-level trade elasticity part less well. To minimize the number of moving parts, and to account for the empirical fact that the number of importing firms is much larger than the number of firms importing from the most popular sourcing country, we continue to include firm-country-specific costs, although these are not present in the original Chaney (2008) paper.

The counterfactual predictions from the knife-edge version of our model in which firms' import entry decisions are independent across markets are quite different from the baseline results. Our baseline model predicts substantial differences between net and gross changes in sourcing, with some firms expanding their sourcing from other countries and the U.S. and other firms contracting. In contrast, under independent entry decisions the only effect on sourcing from other countries comes from a general equilibrium effect (a fall in the aggregate price index), that leads all firms to reduce their sourcing from the U.S. and other countries except China. All active firms reduce their sourcing from the U.S. by the same amount: -0.4 percent (see Table B.4 in the Appendix). Sourcing from third markets also declines. Most notably, the aggregate net effects are exactly equal to the gross effects under these parameter values. There is no scope for churn and reallocation under the assumption of independent entry into sourcing markets.

6.3.2 Universal Importing

Second, we follow Eaton and Kortum (2002) and assume that there are no fixed costs to import, which results in universal importing by all firms. In this exercise, we re-do the quantification of the scale parameter, B , in Step 3 of the estimation, while restricting the fixed cost parameters to be zero. All firms decrease their sourcing from the U.S. and from all other countries except China to 0.986 of the original level (see Table B.5 in the Appendix). Furthermore, as in the case with no interdependencies in firms' entry decisions across markets, net changes in import flows are identical to gross changes in flows.

6.3.3 Sensitivity Analysis

To test the sensitivity of our results to alternative values of the shape parameter of the core productivity distribution, κ , the elasticity of demand, σ , and the dispersion parameter of intermediate input efficiencies, θ , we re-do the estimation of step 3 for alternative values of these parameters. The full results are displayed in table B.6 in the Estimation Appendix, however, we provide a brief summary here. The estimates for the effects of language and distance on the fixed costs of sourcing are remarkably robust across alternative parameters. The estimated range of fixed costs across countries varies, but not dramatically. With respect to the counterfactual predictions, as expected the price index changes are falling in the firm-level trade elasticity θ . The amount of churning attributable to changes in sourcing from the U.S. increases for a lower value of $\theta = 1.3$. In that case, for every three dollar decrease in U.S. sourcing, a one dollar increase in U.S. sourcing takes place. The amount of churning is lower for larger values of θ and for lower values of κ , since more weight is given to the large firms that are already incumbents of sourcing from China. The amount of churning is roughly unchanged when the elasticity of demand increases to $\sigma = 5$.

7 Conclusion

This paper provides a new framework in which to analyze the global sourcing decisions of firms in a multi-country world in which production combines multiple inputs. Our model nests the two canonical models of the extensive margin of exporting, the Eaton and Kortum (2002) Ricardian competitive model and the Melitz (2003) monopolistic competition with heterogeneous firms. Nevertheless, because sourcing decisions generally interact through the cost function, we show that the extensive margin of importing typically involves solving a 2^J -dimensional discrete choice problem (where J is the number of countries) rather than solving J binary problems as in canonical models of exporting. We show, however, that under empirically plausible conditions, we can exploit the existence of complementarities in this large discrete-choice problem to partially characterize the qualitative features of the extensive margin of sourcing. These same complementarities facilitate the implementation of an estimation approach – pioneered by Jia (2008) – that allows us to recover the key parameters of our model from confidential U.S. firm-level data on the sourcing decisions of U.S. firms in 67 countries. Armed with these estimates, we explore the quantitative bite of the key novel features of our framework by studying how a shock to the potential benefits of sourcing from a country (namely, China) differentially affects the sourcing decisions of U.S. firms depending on their core productivity and their pre-shock sourcing strategies. A distinctive characteristic of our framework is that a sectoral import competition shock that does not simultaneously increase export opportunities may still lead to intraindustry reallocation effects by which firms sourcing from these shocked countries may expand, while firms not sourcing from those countries shrink. These effects are not present in the canonical heterogeneous firms trade models, and have not been analyzed in many influential empirical studies on the net implications of import competition shocks (e.g., Autor et al., 2013).

Our theoretical framework is admittedly stylized but it can flexibly accommodate various extensions. We have shown how to incorporate fixed sourcing costs at the input level, and also to characterize firms' importing and exporting decisions jointly. These extensions could serve as the basis for a richer quantitative exercise than the one conducted in the paper. Similarly, we have abstracted from the type of contractual frictions inherent in global sourcing transactions, but as outlined in Antràs (2014), these contractual aspects can also be incorporated in our framework, thus permitting a multi-country analysis of the choice between intrafirm versus arm's-length global sourcing, along the lines of Antràs and Helpman (2004). Finally, we believe that the methodological tools we have developed in this paper, and particularly our application of Jia's (2008) iterative algorithm for solving single-agent entry decisions with interdependencies across markets, could be fruitfully adopted in alternative environments, such as in exporting models with non-constant marginal costs or demand linkages across markets.

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

$$w_i f_e = B_i \int_{\tilde{\varphi}_{i\vartheta(i)}}^{\infty} (\gamma \Theta_i(\varphi))^{(\sigma-1)/\theta} \varphi^{\sigma-1} dG_i(\varphi) - w_i \int_{\tilde{\varphi}_{i\vartheta(i)}}^{\infty} \sum_{j \in \mathcal{J}_i(\varphi)} f_{ij} dG_i(\varphi), \quad (24)$$

where $\vartheta(i)$ denotes the location from which the least productive active firm in country i sources its inputs,

or formally, $\vartheta(i) = \{j \in J : \tilde{\varphi}_{ij} \leq \tilde{\varphi}_{ik} \text{ for all } k \in J\}$. Note that $\vartheta(i)$ satisfies

$$(\tilde{\varphi}_{i\vartheta(i)})^{\sigma-1} B_i \left(\gamma T_{\vartheta(i)} (\tau_{i\vartheta(i)} w_{\vartheta(i)})^{-\theta} \right)^{(\sigma-1)/\theta} = w_i f_{i\vartheta(i)}. \quad (25)$$

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 $w_i f_{ij}$.²⁷

Computing the derivative of the right-hand-side of (24) with respect to B_i , and using (25) to eliminate the effects working through changes in $\tilde{\varphi}_{i\vartheta(i)}$, we can write this derivative as simply

$$\int_{\tilde{\varphi}_{i\vartheta(i)}}^{\infty} \frac{\partial \left(\varphi^{\sigma-1} (\gamma \Theta_i(\varphi))^{(\sigma-1)/\theta} B_i - w_i \sum_{j \in \mathcal{J}_i(\varphi)} f_{ij} \right)}{\partial B_i} dG_i(\varphi) > 0. \quad (26)$$

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.²⁸ We can thus conclude that the right-hand-side of (24) 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 (24) 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 (24) goes to 0.

It thus only remains to show that the right-hand-side of (24) is a *continuously* non-decreasing function of B_i . This may not seem immediate because firm-level profits jump discontinuously with B_i whenever such changes in B_i lead to changes in the global sourcing strategy of firms. It can be shown, however, that

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

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

²⁷To be precise, it could be the case that the least productive active firm in country i might source inputs from more than one location. In such a case, the left-hand-side of equation (25) would incorporate the other location's sourcing potential, but equation (26) below would remain unaltered.

²⁸Following 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 prove the same result for the case of an exogenous number of firms N_i .

at different levels of $B_i\varphi^{\sigma-1}$. This is analogous to writing

$$(\Theta_i(\varphi))^{(\sigma-1)/\theta} B_i\varphi^{\sigma-1} = \begin{cases} \theta_1 B_i\varphi^{\sigma-1} & \text{if } \varphi < a_1/B_i^{1/(\sigma-1)} \\ \theta_2 B_i\varphi^{\sigma-1} & \text{if } a_1/B_i^{1/(\sigma-1)} \leq \varphi < a_2/B_i^{1/(\sigma-1)} \\ \vdots & \vdots \\ \theta_J B_i\varphi^{\sigma-1} & \text{if } a_{J-1}/B_i^{1/(\sigma-1)} \leq \varphi \end{cases}. \quad (27)$$

Hence, we have

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

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.²⁹

Using similar arguments we can next show that

$$\int_{\tilde{\varphi}_{i\theta(i)}}^{\infty} \frac{\partial \left(w_i \sum_{j \in \mathcal{J}_i(\varphi)} f_{ij} \right)}{\partial B_i} dG_i(\varphi) \quad (28)$$

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 (27), in which the position of the steps is continuously differentiable in B_i . This in turn ensures that (24) is continuous in B_i and concludes the proof that there exists a unique B_i that solves equation (12).

²⁹The two last expressions assume that there are $J-1$ jumps, implicitly assuming that at each jump, only one country is added to the sourcing strategy. Given the complementarities in our model, and as pointed out in footnote 27, an increase in B_i might well lead to the simultaneous inclusion of two or more locations. In such a case, there would be less than $J-1$ jumps, but the continuous differentiability of (27) would clearly be preserved.

B Estimation Appendix

B.1 Measuring firm-level offshoring shares

Equation (6) provides a clear formulation of the share of intermediate inputs sourced by a firm in country i from country j . 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.³⁰ 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 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.

B.2 Estimation of the Trade Elasticity

We identify the firm-level trade elasticity, θ , using variation in country wages. The wage data are from the International Labor Organization reported average nominal monthly wages in local currencies for 2007. These wages were converted to USD using exchange rates from the World Bank. When data for 2007 were missing, we used data for the next closest year within a two-year range. To address the fact that skill levels differ across countries, we follow Eaton and Kortum (2002) and use human capital adjusted wages, $w_i^{HCadj} = (w_i)e^{-0.06H_i}$, where H_i is the years of schooling from Barro and Lee (2010) and 0.06 represents the return to education estimated in Bilal and Klenow (2000).

Table B.1 presents several robustness tests for estimation of the firm-level trade elasticity. We use data on country GDP from the World Bank Development Indicators (WDI) and country tariffs are the simple average of country tariffs from the World Bank WITS database. Column 1 shows that our estimate of θ is somewhat smaller, suggesting greater complementarity, when controlling for GDP. Column 2 shows that

³⁰The 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.

it is virtually unchanged when controlling for tariffs. In column 3 we report the results of constraining the coefficient on tariffs and wages to be the same. Finally, column 4 shows results when we do not control for the number of firms in a country.

Table B.1: Robustness estimates for the firm-level trade elasticity

Dependent variable is $\log \xi$				
	IV R1	IV R2	IV R3	IV R4
HC adjusted wage	-1.20*** (0.26)	-1.72** (0.69)		-2.01*** (0.72)
$\log(1+\text{tariff}) + \log \text{wage}$			-1.46*** (0.56)	
$\log \text{distance}$	-0.44** (0.20)	-0.67** (0.34)	-0.50** (0.25)	-0.83** (0.35)
common language	0.21 (0.23)	0.28 (0.31)	0.16 (0.25)	0.20 (0.31)
$\log \text{R\&D}$	0.35*** (0.13)	0.54*** (0.13)	0.50*** (0.11)	0.52*** (0.09)
$\log \text{KL}$	0.24 (0.20)	0.46 (0.39)	0.32 (0.32)	0.61 (0.43)
Control of corruption	0.44** (0.17)	0.58* (0.30)	0.50* (0.26)	0.70** (0.34)
$\log \text{GDP}$	0.19 (0.17)			
$\log \text{tariff}$		6.51 (9.37)		
$\log \text{no. of firms}$	0.02 (0.10)	-0.02 (0.14)	0.02 (0.12)	
Constant	-11.04** (5.20)	-6.03*** (1.57)	-5.70*** (1.38)	-6.60*** (1.87)
Observations	56	56	56	58

Notes: Standard errors in parentheses. The human capital-adjusted wage is instrumented by population. In IV R3, the coefficient on wages and tariffs is constrained to be identical. For this specification, both population and tariffs are the instruments.

Table B.2 presents the first stage regressions for the IV estimates of θ . The first column corresponds to our baseline IV estimate presented in the paper, in which we instrument for country wages using population. As expected, the estimated coefficient on population is negative and significant at the 95 percent level. In column 2, we also report the first stage regression when instrumenting for the sum of wages and tariffs. This first stage corresponds to the IV specification reported in column 3 of Table B.1 above. For this specification, we use both population and tariffs as excluded instruments. Although the coefficient on population remains negative and significant, the coefficient on tariffs is not significant. This result is not surprising since there is very little variation in U.S. import tariffs across countries. For these reasons, we do not rely on tariffs as an empirically reliable way to estimate θ .

Table B.2: First stage regressions for trade elasticity estimates

Dependent variable is	log HC adj. wage	log HC adj. wage + log tariff
log population	-0.31** (0.13)	-0.31** (0.13)
log distance	-0.20 (0.14)	-0.27* (0.16)
common language	-0.05 (0.17)	0.02 (0.18)
log R&D	0.32*** (0.09)	0.32*** (0.09)
log KL	0.14 (0.18)	0.16 (0.18)
Control of corruption	0.22* (0.12)	0.21* (0.12)
log no. of firms	-0.04 (0.07)	-0.05 (0.07)
log tariff		5.81 (5.15)
Constant	0.02 (0.93)	-0.02 (0.93)
R-squared	.88	.88
Observations	56	56

Notes: First stage regressions for the IV estimates presented in Table 6. Wage is the log of the human capital-adjusted wage. Column 1 is the first stage for IV regressions that instrument for wage and in which population is the excluded instrument. Column 2 is the first stage for the IV regression in which the coefficients on wage and tariff are constrained to be equal and in which population and tariffs are the excluded instruments.

B.3 Estimation Results and Counterfactual Predictions

In this Appendix, we first provide Tables analogous to Table 8 in the main text for the counterfactual exercises with (i) fixed sourcing strategies, (ii) independent entry decisions, and (iii) universal importing. We next provide a more extensive table reporting certain estimation and counterfactual results for various alternative parameter values. Finally, we include a table with some details related to the performance of our application of Jia's (2008) algorithm.

Table B.3: FIXED SOURCING STRATEGIES: THIRD COUNTRY SOURCING EFFECTS OF POSITIVE CHINESE SOURCING POTENTIAL SHOCK

Chinese import status	Change sourcing from US	Change Sourcing from other countries	Change Sourcing from China	Share of firms
Entrants	-	-	-	0
Exiters	-	-	-	0
Continuers	1.001	1.001	2.002	0.098
Others	0.991	0.991	-	0.902

Notes: Entrants (exiters) are those firms that begin (stop) sourcing from China. Columns 1, 2, and 3 contain the ratio of the total sourcing by this group of firms before and after

Table B.4: INDEPENDENT ENTRY DECISIONS: THIRD COUNTRY SOURCING EFFECTS OF POSITIVE CHINESE SOURCING POTENTIAL SHOCK

Chinese import status	Change sourcing from U.S.	Change Sourcing from other countries	Change Sourcing from China	Share of firms
Entrants	0.996	0.991	∞	0.158
Exiters	-	-	-	0
Continuers	0.996	0.992	1.993	0.078
Others	0.996	0.989	-	0.763

Notes: This table contains only surviving firms. Entrants (exiters) are those firms that begin (stop) sourcing from China. Columns 1, 2 and 3 contain the ratio of the total sourcing by this group of firms before and after. Independent entry decisions arise of $\sigma = 2.71$.

Table B.5: UNIVERSAL IMPORTING: THIRD COUNTRY SOURCING EFFECTS OF POSITIVE CHINESE SOURCING POTENTIAL SHOCK

Chinese import status	Change sourcing from U.S.	Change Sourcing from other countries	Change Sourcing from China	Share of firms
Entrants	-	-	-	0
Exiters	-	-	-	0
Continuers	0.986	0.986	1.972	1
Others	-	-	-	0

Notes: This table contains only surviving firms. Entrants (exiters) are those firms that begin (stop) sourcing from China. Columns 1, 2 and 3 contain the ratio of the total sourcing by this group of firms before and after. Universal importing arises if the fixed costs of foreign sourcing are zero.

Table B.6: Sensitivity of parameter estimates and counterfactual predictions to alternative values for θ , κ , and σ

	Variation in θ				Variation in κ			Variation in σ	
	Baseline	$\theta = 1.3$	$\theta = 2$	$\theta = 2.85$	$\kappa = 3$	$\kappa = 4$	$\kappa = 5$	$\sigma = 2.71$	$\sigma = 5$
Parameter Estimates									
B	0.125	0.124	0.124	0.128	0.098	0.117	0.130	0.249	0.074
β_c^f	0.011	0.015	0.011	0.006	0.018	0.012	0.011	0.012	0.016
β_d^f	0.339	0.341	0.369	0.423	0.492	0.368	0.324	0.261	0.463
β_l^f	0.612	0.613	0.620	0.676	0.585	0.613	0.607	0.505	0.596
β_{disp}^f	0.865	0.862	0.994	0.940	1.136	0.917	0.845	0.876	1.086
<i>Median fixed cost range across countries (in tsd USD)</i>									
Minimum	8.933	11.603	8.821	5.927	14.856	9.865	8.323	7.527	13.742
Maximum	28.880	37.622	29.894	20.514	68.376	33.763	26.292	25.205	58.567
Counterfactual predictions									
Price index change (in percent)	-0.395	-0.529	-0.305	-0.218	-0.610	-0.445	-0.354	-0.215	-0.576
Change in third country sourcing by new importers from China (in percent)	2.280	5.348	0.953	-1.371	0.603	1.806	2.624	-0.938	3.667
Gross increase in U.S. sourcing (in percent of total U.S. sourcing)	0.162	0.476	0.067	0.000	0.026	0.118	0.205	0.000	0.166
Net change in U.S. sourcing (in percent of total U.S. sourcing)	-1.056	-1.081	-0.957	-0.968	-1.203	-1.082	-1.028	-0.798	-1.192

Notes: This table presents Step 3 estimation results and counterfactual predictions for alternative parameter values for the firm-level trade elasticity θ , the shape parameter of the Pareto distribution κ , and the elasticity of substitution σ .

Table B.7: Cardinality of differences in bounds

Cardinality of differences in bounds		0	1	2	3	4	5	6	7	8	9 - 25	≥ 26
	Baseline	8,425,677,988	0	448,720	31,297	1,901	92	2	0	0	0	0
Variation in θ	$\theta = 1.3$	9,494,065,126	0	1,155,451	124,863	12,966	1,410	179	5	0	0	0
	$\theta = 2$	9,510,345,043	0	129,760	4,912	243	42	0	0	0	0	0
	$\theta = 2.85$	9,473,760,000	0	0	0	0	0	0	0	0	0	0
Variation in κ	$\kappa = 3$	9,628,850,232	0	401,259	26,669	1,652	183	5	0	0	0	0
	$\kappa = 4$	8,350,060,201	0	465,411	32,291	1,986	106	5	0	0	0	0
	$\kappa = 5$	7,846,965,683	0	295,886	17,468	913	48	1	1	0	0	0
Variation in σ	$\sigma = 2.7$	9,817,200,000	0	0	0	0	0	0	0	0	0	0
	$\sigma = 5$	9,729,196,313	0	1,398,164	180,606	22,259	2,321	314	18	4	1	0

Notes: This table displays the number of firm (productivity and fixed cost draws) and parameter combinations for which the cardinality of the differences in the bounds reached a particular value. While the productivity and fixed cost draws are held fixed during the estimation process, the parameter vector, δ , varies through the iterations of the estimation process. We allow the differences of the bounds to be less than 26, before we would revert to evaluating the objective value of the firm's problem at the bounds and a small number of random values for the countries in the bound. Since this cardinality of the differences in the bounds is never very high, we always solve accurately the problem of the firm.

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: Firms, Contracts and Trade Structure* 2014. unpublished 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.
- , **Luis Garicano, and Esteban Rossi-Hansberg**, “Offshoring in a Knowledge Economy,” *The Quarterly Journal of Economics*, 2006, *121* (1), 31–77.
- Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare**, “New Trade Models, Same Old Gains?,” *American Economic Review*, February 2012, *102* (1), 94–130.
- , **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.
- Autor, David H., David Dorn, and Gordon H. Hanson**, “The China Syndrome: Local Labor Market Effects of Import Competition in the United States,” *American Economic Review*, 2013, *103* (6), 2121–68.
- Bache, Peter and Anders Laugesen**, “Monotone Comparative Statics for the Industry Composition,” *Working Paper, Aarhus University*, 2013.
- Barro, Robert and Jong Wha Lee**, “A New Data Set of Educational Attainment in the World, 1950–2010,” 2010. NBER Working Paper No. 15902.
- Bernard, Andrew B. and Teresa C. Fort**, “Factorlyess 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.
- Bils, Mark and Peter J. Klenow**, “Does Schooling Cause Growth?,” *American Economic Review*, 2000, *90*, 1160–1183.
- Blaum, Joaquin, Claire Lelarge, and Michael Peters**, “Non-Homothetic Import Demand: Firm Productivity and Quality Bias,” *Unpublished paper*, 2013.
- , —, and —, “Estimating the Productivity Gains from Importing,” *Working Paper, Brown University*, 2014.
- Blum, Bernardo S, Sebastian Claro, and Ignatius Horstmann**, “Facts and Figures on Intermediated Trade,” *American Economic Review*, May 2010, *100* (2), 419–423.
- Broda, Christian and David E. Weinstein**, “Globalization and the Gains From Variety,” *The Quarterly Journal of Economics*, May 2006, *121* (2), 541–585.

- 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 Jenkins, 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.
- and **John Romalis**, “International Prices and Endogenous Quality,” *The Quarterly Journal of Economics*, May 2014, *129* (2), 477–527.
- Fort, Teresa C.**, “Technology and Production Fragmentation: Domestic versus Foreign Sourcing,” Working Paper 2014.
- 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.
- Heston, Alan, Robert Summers, and Bettina Aten**, “Penn World Table Version 7.0,” 2011. Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania.
- 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 and Pinelopi Koujianou Goldberg**, “Firm Performance in a Global Market,” *Annual Review of Economics*, 2014, 6 (1), 201–227.
- , **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.
- **and Stephen Redding**, “Heterogeneous Firms and Trade,” in “Handbook of International Economics,” Vol. Volume 4, Elsevier, 2013.
- 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.
- Ramanarayanan, Ananth**, “Imported Inputs and the Gains from Trade,” *Working Paper, Western Ontario*, 2014.
- 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.
- Tintelnot, Felix**, “Global Production with Export Platforms,” 2014. mimeo University of Chicago.
- 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.